DOCTOR OF PHILOSOPHY

Posture classification for real life, real time applications

Rednic, Ramona

Award date:
2012

Awarding institution:
Coventry University

Posture Classification for Real life, Real time Applications

Ramona Arabela Rednic

A thesis submitted in partial fulfilment
of the University's requirements for the Degree of
Doctor of Philosophy

April 2012

Coventry University
Faculty of Engineering and Computing

Abstract

In recent years, Body Sensor Networks (BSNs) have been used as the basis of many systems aimed at monitoring bodily parameters ranging from skin temperature, to breathing, to motion. These measurements can then be used to generate additional information related to the monitored subject, such as for heat stress prediction or fall detection.

This thesis is concerned with the design, development and realistic evaluation of a BSN-based end-to-end posture classification platform using on-body accelerometers. The work is motivated by applications that require stable, sub-second, end-to-end classification of postures, as well as dynamically configurable operation to support exploratory data collection. Classification is performed on-node, thus reducing the amount of data/information transmitted from the wearable nodes to a data sink. The work is experimentally-led, and uses an application case study—on-body monitoring of Explosive Ordnance Disposal (EOD) operatives—to provide context for system requirements and experimentation performed.

This thesis provides three main contributions:

First, the design of a platform that allows real-time on-body classification of static and dynamic postures—a capability not present in existing work. The specific posture set consists of six static postures (sitting, standing, kneeling, and lying on back, front and one side) and two dynamic postures (walking and crawling), of which kneeling and crawling are not commonly considered in the literature. Classification is performed on a small, light embedded device using a simple easy-to-implement algorithm. The classification algorithm used is a C4.5 decision tree, with a temporal feature (windowed variance) to aid in distinguishing dynamic and static postures. Offline classification accuracy is shown to be 96.3% based on data gathered from subjects in a laboratory environment, and real-time on-node classification accuracy is shown to be comparable to this figure (95.5%).

Second, further advance beyond the state of the art is presented through an investigation into posture transitions. Posture transitions cause transient (<1s) posture changes in the classifier output and are shown to reduce classifier accuracy by 2% for every transition / minute for classifiers not specifically designed to handle transitions. Three posture filters that remove such transient posture changes are designed, implemented and tested on experimental data. The best performing filter, Exponentially Weighted Voting (EWV), is shown to reduce posture change events by 75.2% and increase accuracy by 1% (over unfiltered results). Compared to streaming raw data, an event-based posture classification system is shown to reduce transmissions by 98.5% (66-fold reduction).

Finally, a broad investigation is presented into the effect of both system-related and training-process factors on the accuracy of machine learning-based posture classifiers. The factors analysed include i) temporal and feature parameters, ii) sensor sampling rate, iii) number of sensors used, iv) posture class aggregation and v) number of subjects used for training. Optimal parameters are determined for the motivating EOD application, with a range of parameter values shown to guide development of other classifiers.

Through the novel contributions presented, this thesis provides a solid groundwork for further research in BSN based posture classification systems and simplifies optimisation of machine-learning classifiers for specific posture classification applications.

Acknowledgements

Without a doubt, I would not be able to accomplish the achievements within this thesis without support.

I would like to express my sincere gratitude to my Director of Studies, Professor Elena Gaura, for her advice and encouragement throughout the work. I could not have asked for a better supervisor to guide me on the road to getting a Ph.D. As a female researcher I found it greatly motivating to have such an accomplished Director of Studies. As a supervisor and as friend she was always looking out for me. She finds ways to make you want more from yourself and supports you in achieving your goals.

I would like to thank my second supervisor, Dr. James Brusey, who made his support available in a number of ways. James always had something new to teach me and he always tried to set me on the right path. Throughout my Ph.D. he provided good teaching, good company, sound advice and lots of good ideas.

I would also like to thank Matt Welsh for giving me the opportunity to work alongside his students at Harvard University on a very exciting project. My time spent at Harvard was also made enjoyable by Bor-rong Chen, and his support was greatly appreciated.

I wish to thank my parents for providing their love and moral support over the years. Their encouragement kept me going and made me believe I can make it.

I would like to thank the staff and students at Cogent Computing for their friendship and support. Special thanks go to Dr. Michael Allen for his support in the final stages of my Ph.D.

Words fail me to express my appreciation to my colleague and partner John, whose dedication, love and persistent confidence in me have helped me through these years. I owe him for all the support and patience during the final stages of my Ph.D.

Contents

List of Figures

List of Tables

Acronyms

BAN	Body Area Network
BSN	Body Sensor Network
COPD	Chronic Obstructive Pulmonary Disease
CO_2	Carbon Dioxide
EKG	Electrocardiograph
EOD	Explosive Ordnance Disposal
EWMA	Exponentially Weighted Moving Average
EWV	Exponentially Weighted Voting
FFT	Fast Fourier Transform
GUI	Graphical User Interface
HMM	Hidden Markov Model
LOSOXV	Leave-One-Subject-Out Cross-Validation
NTP	Network Time Protocol
RMS	Root Mean Square
SMA	Signal Magnitude Area
SVM	Signal Vector Magnitude
SHIMMER	Sensing Health with Intelligence, Modularity, Mobility and Experimental Reusability
UHS	Uncompensable Heat Stress
WEKA	Waikato Environment for Knowledge Analysis
WM	Windowed Mean
WVar	Windowed Variance

Chapter 1

Introduction

In recent years, the increased availability of small, low power, high capability personal computing devices and inexpensive miniaturised sensors has made non-intrusive pervasive sensing of people and their environment a practical possibility. The provision of such capabilities has applications ranging from workplace safety to healthcare to the military, based on diverse sensed parameters such as air pollutants, worker activity level, or structural stresses.

This thesis is concerned with the use of machine learning for real-time postural activity classification based on acceleration data. Posture classification determines the posture of the subject (assigning a label such as "sitting" or "standing"), as opposed to motion capture which focuses on determining the relative position of each segment of the body. The technology available to support real-time remote posture monitoring has evolved over time: advances in micro-sensors have resulted in smaller and lighter accelerometers [13, 115], the adoption of embedded computing into common household and personal devices has driven the development of microcontrollers with high processing capability and low power draw in small packages, and radio-based communication has replaced wired links in many applications [91]. In combination, these technologies can form the basis of a compact and ubiquitous Body Sensor Network (BSN) solution for human posture/activity monitoring [14, 74, 86, 116]. BSN-based systems using these technologies have the advantage that the equipment required to capture motion or posture is much smaller and lighter than other solutions, (such as video-based capture) and can be carried by the subject—reducing the need for equipment to be installed in the subject's environment. The use of radio-based communication means that the subject is not tethered to a specific location, allowing for natural movement and the ability to carry out tasks away from a monitoring base station. The ability to automatically recognise posture can facilitate the provision of personalised computer-based support in areas such as medicine [12, 63, 85, 93, 96], workplaces [81, 105], and sports [10, 79].

The goal of the work in this thesis is to demonstrate that posture classification can be accurately performed in real-time using a wearable monitoring system. Such postural information can be displayed to an observer or used for further autonomous modelling/prediction. This is of particular benefit where the monitored subject is required to work away from a support team in dangerous conditions, such as for firefighters or some military applications. The work in this thesis, therefore, focuses on three main topics or strands: 1) real-time on-body classification of posture, 2) reducing the impact of posture transitions on classifier performance, and 3) an investigation of the design space and selection of optimal system design parameters for posture classification systems.

This chapter provides an overview of the motivation for the work and the methods used, structured as follows: Section 1.1 provides a broad overview of the requirements common to the applications benefiting from posture monitoring. Section 1.2 presents and justifies the case study that provides context for the work in this thesis. Section 1.3 presents the research questions that drove the work. Section 1.4 details the approach to research, data gathering and classifier assessment methodology. Section 1.5 presents the contributions to knowledge. Section 1.6 lists publications by the author resulting from the work presented in this thesis. Section 1.7 acknowledges the contributed work. Finally, Section 1.8 describes the structure of this thesis.

1.1 General requirements of posture/activity monitoring applications

This work focuses on real life applications as these provide concrete requirements and constraints to guide the development of usable monitoring systems. Knowledge of posture is an important source of information in a diverse range of applications, including health monitoring, work safety, dance, sports, and video games. Specific applications within these categories include monitoring of patients undergoing physical rehabilitation, detection of falls in elderly people, monitoring the activity of workers in dangerous environments, and detecting deviation from daily routine. Section 2.1 on page 16 provides details of a number of such applications.

Generally, a posture monitoring deployment has one of two goals: 1) observation, decision making, and application of corrective actions, or 2) profiling of the subject's routine over a period of time. Within the health monitoring domain, for example:

- Long-term postural monitoring of patients with Parkinson's disease can provide an indication of the progression of the condition.
- Monitoring of patients undergoing an at-home physical rehabilitation regime can show whether they are performing the recommended exercises and maintaining an appropriate level of general activity.

Such applications have a set of generic requirements that are commonly applicable, a subset of which must be met for each individual application. Requirements that are often specified for posture classification applications are:

- The need for high accuracy. This is a fundamental requirement for posture classification systems targeted at any application. The classifier accuracy should at least be as high as the state of the art, which in related work appears to be around 90%.

- The need for a small, lightweight system that can be worn on the body. On-body systems have become popular as they avoid the need to install monitoring equipment such as video cameras in the subject's environment. Considerations for unobtrusive wearable systems include the size and weight of the on-body components, the quantity of cabling, and the method of affixing sensors to the body.

- The need to monitor subjects performing free-form activities in an environment such as a hospital or the home. While some applications (monitoring patients undergoing physical rehabilitation for example) may have relatively constrained or predictable activities, many other real-world deployments will not. When monitoring the daily activities of a subject, for example, such activities could include any form of free movement, including static and dynamic postures along with transitions between them.

- Provision of postural information in real-time. In general terms, this is motivated by the need to provide timely input to another system or to monitor in real-time the evolution of a data stream. The exact definition of real-time in terms of latency will be motivated by an application's specific needs. In this thesis, a *real-time* posture classification system is defined as one in which the nominal end-to-end latency of the system (including classification and any wireless communication) is the same order of magnitude as the sampling period. Transmission latency may increase due to external influences (such as radio interference) but these are expected to be transient in most cases. (Note that this is a soft real-time requirement as a delayed posture result is not catastrophic for the class of applications considered here, but does degrade the usefulness of the postural information delivered by the system.)

- On-body processing of sensor data in cases where 1) autonomous operation is required, or 2) battery life is affected by the quantity of data transmitted. Police and firefighters, for example, are required to work in areas where permanent monitoring equipment is not installed. In these applications, the

These images have been removed

Figure 1.1: Subject wearing an EOD suit while walking and kneeling. Reprinted from Kemp [66] with permission.

subjects are mobile and are responding to events as they occur, meaning that communication with a base station cannot necessarily be relied upon.

The case study application focused on in the work here is related to work safety, specifically the monitoring of Explosive Ordnance Disposal (EOD) operatives during missions. The requirements of this application include all five of those described above, meaning that development work towards designing and implementing a system for EOD operative monitoring can feed back into the design of posture monitoring systems targeted at other applications. In particular, the EOD application shares characteristics with similar applications such as the monitoring of infantry or firefighters. Furthermore, the case study application guides the data gathering and classifier evaluation process through the defining of realistic constraints and requirements. The EOD operative monitoring application is described in the next section.

1.2 Case study: posture classification of EOD operatives during missions

The specific application chosen to guide the constraints and requirements for this work was that of monitoring the postures/activities of EOD operatives during missions.

Particular parallels between the EOD operative monitoring case study and the work here were: 1) the typical postures and activities encountered during EOD missions map to the postures considered in this work and 2) operative health considerations (described later in this section) require the provision of real-time on-body classification of posture. As noted in the previous section, the EOD application guided the data gathering and classifier evaluation methodology design in Chapter 4, prompting the use of specific activity routines for example.

The EOD application shares characteristics with a wider class of applications that includes monitoring of personnel such as firefighters and infantry. The similarities between the EOD application and other personnel monitoring applications include the use of protective clothing, the typical postures and activities encountered (including some not typically found in healthcare related or daily activity monitoring), and the need for real-time physiological monitoring due to the harsh environments that may be endured by the monitored subject. While EOD operative monitoring is the case study application within this thesis, the similarities to the applications in the wider class will allow the work here to be generalised and applied to those applications.

During a typical mission, the EOD operative has to wear a protective suit and helmet (which together weigh over 40 kg) and carry a tool box of equipment the 100 or so meters to the site. To reach the bomb's location and fulfil the mission, it may be necessary to climb stairs, crawl through passageways, kneel, use specialist equipment, or lie down. Examples of a subject wearing an EOD suit are shown in Figure 1.1 on the preceding page.

Within the enclosed suit micro-climate, evaporative cooling through perspiration is less effective. This can lead to Uncompensable Heat Stress (UHS), which occurs when the body cannot cool itself as fast as heat is being generated due to muscular exercise [33, 75, 117]. The result of UHS is that core temperature increases beyond the safe range, leading to health problems and potentially death.

An additional problem is that of potentially dangerous build-up of Carbon Dioxide (CO_2) within the EOD operative's enclosed helmet—it has been shown that, even when below toxic levels, excessive inhaled CO_2 concentrations combined with high temperatures may lead to cognitive impairment [25]. Within the helmet, the CO_2 concentration can increase to as much as a factor of 60 over the ambient level.

The suit manufacturer's solution to these hazards is to integrate a cooling system within the suit that blows cooled air into the helmet and onto the operative's back. In theory this serves the dual purposes of reducing the temperature within the suit and the CO_2 concentration within the helmet. However, the cooling system's battery life is not sufficient to last for the entire mission duration (performed in segments of approximately 1 hour) if the operative sets the fan to a high speed and then performs no other control adjustments. The operative is likely to do this as their primary concern during the mission is on disabling the explosive device.

This graph has been removed

Figure 1.2: Mean skin temperature against $P(T_{sk,u} \geq T_d|T_{sk,t})$ with $T_d = 36.5$ °C and chest cooling applied. Curves are shown for individual activities and for an aggregate if activity information is not known and all activities are equally likely. Reprinted from Kemp [66] with permission.

The Medusa2 system [66] provides a solution for effective cooling system control via automatic actuation of the cooling fans based on sensors integrated within the EOD suit. Medusa2 senses skin temperature at multiple locations and the CO_2 concentration within the helmet. These measurements are used to provide automatic control of the cooling via predictive modelling to determine the risk of UHS occurring and to detect excessive CO_2 concentration in the helmet.

A requirement of both the predictive health risk modelling and the air quality control implemented within the Medusa2 system is the availability of real-time postural information. Figure 1.2 demonstrates the impact of posture on the UHS risk prediction modelling performed by the Medusa2 system. It can be seen that the probability of the *future* average skin temperature (that is, the skin temperature after 5 minutes) exceeding a critical threshold, T_d, of 36.5 °C varies for any given *current* skin temperature based on the subject's posture/activity. For example, if the subject is performing the "Weights" activity then the risk is significantly higher than for the "Treadmill" activity, particularly for skin temperatures below 36 °C. Knowledge of posture within the prediction algorithm thus not only allows potentially life-saving warning of dangerous conditions but could also allow energy saving through providing a lower level of cooling in some cases.

Two possible explanations for the demonstrated effect that posture has on heat stress risk are that 1) movement, such as walking, has a tendency to force air to circulate within the suit, whereas certain postures such as kneeling will restrict circulation of air and 2) each activity has an associated level of energy use, and thus heat production, within the body. Kemp [66] shows a similar posture-dependency when investigating CO_2 concentration within the EOD suit helmet. It is thus clear that postural information is an essential parameter to enable accurate prediction of health risks such as UHS.

In generic terms, the EOD application is considered to be a good case study for the posture classification research presented in this thesis as it has three traits that generalise well to other applications:

1. A need to classify both static (such as sitting) and dynamic (such as walking) postures while the operative is performing free-form activities. The operative is not being instructed or constrained as they would be in a laboratory environment. For example, they may kneel while also moving objects out of their path or lie down while also inspecting a suspicious device.

2. A need for real-time operation. Real-time posture classification is required as postural information is supplied to an existing system (Medusa2) as an input to support modelling towards ensuring the operative's safety.

3. A need for local on-body processing. Classification of posture must be performed on-body (rather than utilising a more powerful base station computer) as 1) a radio link to the base station cannot be guaranteed, due to environmental obstructions such as buildings and 2) the unpredictable latency of wireless links impacts on the required real-time operation specified in (2) above.

1.3 Research questions

The motivation for the work in this thesis is the development a real-time wearable BSN-based instrument capable of classifying static and dynamic postures. As the instrument is targeted at a realistic application, there are issues that must be considered that would not necessarily be encountered in a laboratory environment or in a theoretical case study. The contributions in this thesis are largely around providing solutions to these issues. The research questions answered in this thesis are:

1. Can the defined set of postures, namely sitting, kneeling, crawling, standing, walking, and lying on front, back and one side, be accurately classified in real-time using an on-body wearable sensor-based system?

2. How do transitions between postures affect classifier accuracy and can any negative impact be reduced or eliminated?

3. What is the design space for a posture classifier targeted at specific application requirements?

These questions are described in more detail as follows:

1.3.1 Can the defined set of postures, namely sitting, kneeling, crawling, standing, walking, and lying on front, back and one side, be accurately classified in real-time using an on-body wearable sensor-based system?

This question breaks down into two main sub-questions. First, can the given combination of static and dynamic postures be classified with a high accuracy? Second, can that classification be performed in real-time on a lightweight on-body device? These questions are addressed in Chapters 3 and 4.

1.3.2 How do transitions between postures affect classifier accuracy and can any negative impact be reduced or eliminated?

In the grammar of human movement, a postural transition separates two distinct postural phases, such as sitting and standing. Posture classification research has mostly ignored the problem of transitions, however they are an important part of normal human movement and a successful classification system must therefore handle them as correctly as possible. Understanding and addressing the problem of postural transitions is the subject of Chapter 5.

1.3.3 What is the design space for a posture classifier targeted at specific application requirements?

The term “design space” refers here to the parameters that must be chosen during the development of a posture classifier in order to ensure that it is capable of a high classification accuracy. This question thus breaks down into several sub-questions such as 1) how many sensors are required, 2) which data feature best allows postures such as standing and walking to be distinguished, and so on. The factors considered in the work here are: 1) extracted data feature choice, 2) data feature window size, 3) number and location of on-body sensors, 4) training set size, 5) sampling rate, and 6) targeting of individual postures. This investigation is described in Chapter 6.

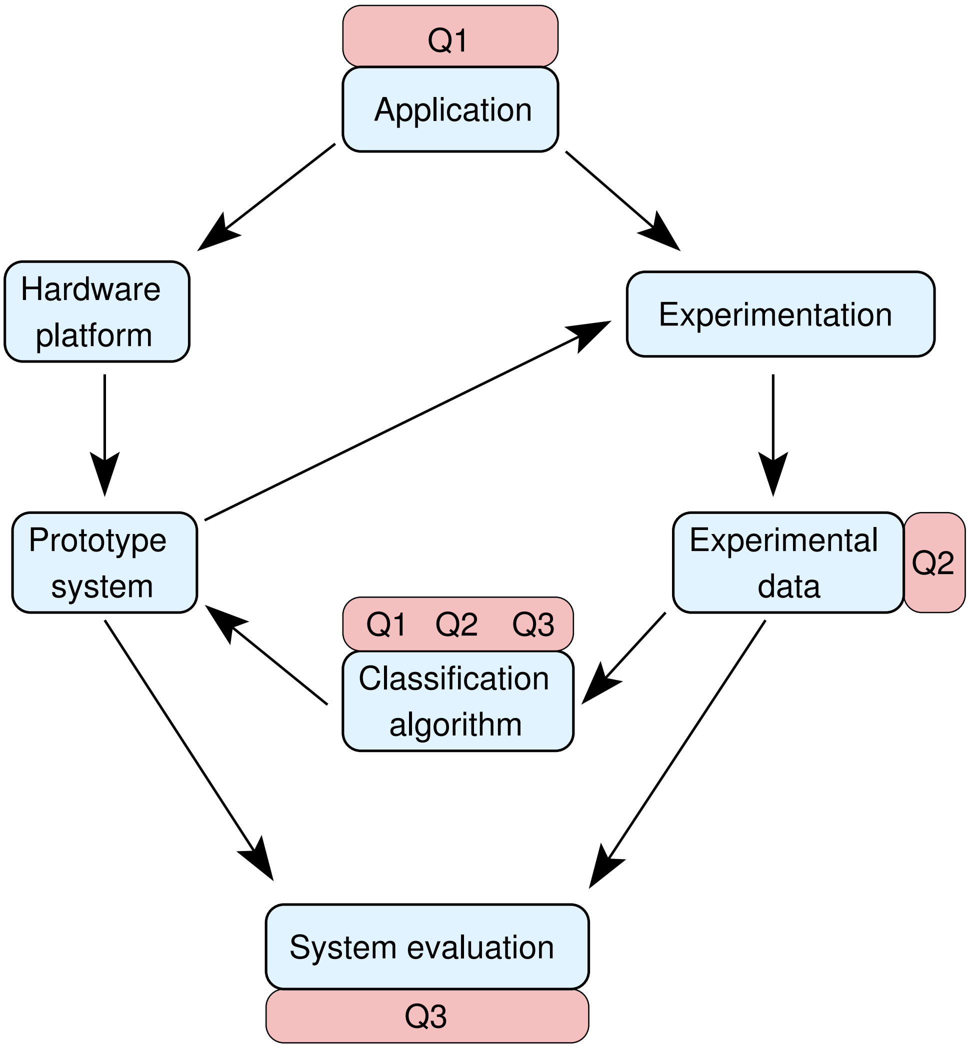

Figure 1.3: Overall structure of the work in this thesis, showing how each aspect contributed towards the others.

1.4 Approach to research

Much of the discovery in this thesis was enabled by an iterative "prototyping-deployment-data analysis" approach. This practical investigation took advantage of the driving application to provide realistic constraints and requirements for the systems and algorithms developed. The experimentally-led investigation meant that results found throughout the work could feed back into the development of the system towards suitability for real-world deployments in the given application scenario.

Figure 1.3 summarises the overall structure of the work performed, along with the elements that contributed to answering each of the research questions given in Section 1.3. The following subsections describe the role of experimentation in the work performed (Section 1.4.1) and the role of iterations in the system prototyping (Section 1.4.2).

1.4.1 Experimentally led investigation

The work presented within this thesis is entirely experimentally focused. Data was gathered from human subjects performing a series of experimental activity regimes by use of a prototype body sensing system based on aspects of a posture classification platform developed in the course of this work. The activity regimes were defined based on the needs of the case study application, reflecting the types of activities that an EOD operative is expected to perform. The gathered data was used in training and testing posture classifiers and so it was important that data was gathered from enough subjects and that the regimes were implemented consistently across trials. The final implemented software algorithms were implemented on a system derived from the platform described in Chapter 3, forming a full end-to-end classification system. This allowed them to be experimentally evaluated to demonstrate the functionality of the complete system and confirm that the overall accuracy matched the results found during offline testing.

A key aspect of this experimental work is to ensure that results are likely to match those that would be found when the system is deployed. There are several elements to this:

1. The use of Leave-One-Subject-Out Cross-Validation (LOSOXV), described in Section 4.3 on page 61, ensures that the estimated system performance is not specific to any group of human subjects. Human subject specificity is a key problem in this domain. Ordinary stratified cross validation will tend to overestimate the true performance on unseen human subjects and is thus inappropriate.

2. The experimental regimes are designed to be realistic and involve natural movements. Regimes are further described in Section 4.4 on page 63.

3. Human subjects with a range of heights and body builds were used, as described in Section 4.6 on page 67.

4. All the proposed algorithms were further evaluated through deployment on wearable hardware and used in realistic scenarios involving tasked activities, as described in Sections 4.12 on page 76 and 5.5 on page 94.

1.4.2 Iterative system prototyping

The investigations presented in this thesis required the implementation of a body sensing system for data gathering purposes to enable the training and testing of posture classifiers. Furthermore, the full end-to-end posture classification platform required evaluation through implementation as a prototype on-body system. This led to an iterative approach to prototyping, starting with a system focused on gathering the needed data and progressing through iterations as new features and software algorithms were implemented over the course of the work. This progression allowed the real-time testing and evaluation of the algorithms.

The hardware and software forming the basis for the prototype instrumentation system are described in Chapters 3 and 4, with further hardware details in Appendix A.

1.5 Contributions to knowledge

In answering the research questions listed in Section 1.3, this thesis provides several contributions to knowledge, as follows:

- The design of a platform that allows real-time on-body classification of static and dynamic postures—a capability not present in existing work. The specific posture set consists of six static postures (sitting, standing, kneeling, and lying on back, front and one side) and two dynamic postures (walking and crawling), of which kneeling and crawling are not commonly considered in the literature. Classification is performed on a small, light embedded device using a simple easy-to-implement algorithm. The classification algorithm used is a C4.5 decision tree, with a temporal feature (windowed variance) to aid in distinguishing dynamic and static postures. This contribution is presented in the work in Chapters 3 and 4.

- The design and implementation of several posture filters to prevent rapid (>1 Hz) classifier output changes during posture transitions. The impact of transitions is often not considered in the posture classification literature, as described in Section 2.5. The filters are evaluated in terms of their effect on 1) classification accuracy and 2) the number of posture change events generated. They provide a solution to handling postural transitions targeted at the case-study scenario but is also applicable more generally. This contribution is presented in Chapter 5.

- An evaluation of factors that affect posture classification accuracy in a deployed system. The factors considered are: 1) extracted data feature choice, 2) feature window size, 3) number of sensors,

4) training set size, 5) sampling rate, and 6) targeting of individual postures. A methodical investigation of parameters such as this is absent in the existing literature. Where investigation is present, it focuses on a subset of the factors and presents results specific to the system implementation used (an overview of existing investigation is given in Section 2.4). Optimal parameters are selected for the application scenario targeted here, but the results and discussion provide more general applicability for similar decision tree based posture classification systems. This investigation is presented in Chapter 6.

1.6 Publications resulting from this work

This section lists the academic publications by the author resulting from 1) the work presented in this thesis and 2) related work within a separate research project. Appendix B gives a full list of publications along with their abstracts.

1. Elena Gaura, James Brusey, Michael Allen, Ross Wilkins, Dan Goldsmith, **Ramona Rednic**. **Edge mining the internet of things**. IEEE Sensors Journal May 2013 , To appear

2. John Kemp, Elena Gaura, **Ramona Rednic**, James Brusey, **Long-term behavioural change detection through pervasive sensing**. In Proceedings of the 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2013), Honolulu Hawai, U.S.A., 1–3 July 2013.

3. **Ramona Rednic**, John Kemp, Elena Gaura, James Brusey. **Fielded autonomous posture classification systems: Design and realistic evaluation.** In Proceedings of the 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2013), Honolulu Hawai, U.S.A., 1–3 July 2013.

4. **Ramona Rednic**, Elena Gaura, James Brusey, and John Kemp. **Wearable posture recognition systems: factors affecting performance.** In Proceedings of the IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI 2012), pages 200–203, Shenzhen, China, 5–7 January 2012.

5. **Ramona Rednic**, John Kemp, Elena Gaura, and James Brusey. **Networked body sensing: Enabling real-time decisions in health and defence applications.** In Proceedings of the Annual International Conference on Advance Computer Science and Information Systems 2011 (ICACSIS 2011), pages 17–24, Jakarta, Indonesia, 17–18 December 2011

6. Bor-rong Chen, Shyamal Patel, Thomas Buckley, **Ramona Rednic**, Doug McClure, Ludy Shih, Daniel Tarsy, Matt Welsh, and Paolo Bonato. **A Web-Based System for Home Monitoring of Patients With Parkinson's Disease Using Wearable Sensors.** IEEE Transactions on Biomedical Engineering (TBME) Letters Special Issue on Emerging Technologies in Point-of-Care Health Care, 58(3):831–836, March 2011.

7. Shyamal Patel, Bor-rong Chen, Thomas Buckley, **Ramona Rednic**, Doug McClure, Daniel Tarsy, Ludy Shih, Jennifer Dy, Matt Welsh, and Paolo Bonato. **Home monitoring of patients with Parkinson's disease via wearable technology and a web-based application.** In Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pages 4411–4414, Buenos Aires, 31 August–4 September 2010. IEEE.

8. Bor-rong Chen, Thomas Buckley, **Ramona Rednic**, Shyamal Patel, Paolo Bonato, and Matt Welsh. **MercuryLive: A Web-Enhanced Platform for Long-Term High Fidelity Motion Analysis.** In Proceedings of the 7th Annual IEEE Communications Society Conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON), pages 1–2, Boston, USA, 21–25 June 2010.

9. James Brusey, **Ramona Rednic**, and Elena Gaura. **Classifying transition behaviour in postural activity monitoring.** Sensor & Transducers Special Issue, 17(10):213–223, October 2009.

10. James Brusey, **Ramona Rednic**, Elena I. Gaura, John Kemp, and Nigel Poole. **Postural activity monitoring for increasing safety in bomb disposal missions.** *Measurement Science and Technology*, 20(7):075204 (11pp), July 2009.

11. **Ramona Rednic**, Elena Gaura, and James Brusey. **ClassAct: Accelerometer-based real-time activity classifier.** In Sensors & Instrumentation KTN: Wireless Sensing Demonstrator Showcase (WiSIG), 2 July 2009.

12. **Ramona Rednic**, Elena Gaura, James Brusey, and John Kemp. **Wireless sensor networks for activity monitoring in safety critical applications.** In *Proceedings of NSTI Nanotech 2009*, volume 1—Fabrication, Particles, Characterization, MEMS, Electronics and Photonics, pages 521–525, Houston, Texas, USA, May 3–7 2009. ISBN: 978-1-4398-1782-7.

13. **Ramona Rednic**, John Kemp, Elena Gaura, James Brusey. **Posture Determination Using a Body Sensor Network.** Technical Report COGENT.006, Coventry University, 2008.

1.7 Acknowledgement of contributed work

The following is the contribution made by other researchers which has aided the progress of the work presented in this thesis:

- Dr John Kemp and Dr Xiang Fei developed the accelerometer sensor board described in Section A.1.2 on page 146. Testing of the boards and integration with the posture classification systems was performed by the author.

1.8 Thesis structure

This thesis is structured as follows: This chapter described the motivation and aims of the work in this thesis and presented the research questions and the resulting contributions. Chapter 2 provides a review of the literature surrounding posture classification applications and techniques. Chapter 3 describes the design of a real-time end-to-end posture classification platform along with two example usage scenarios and an implementation used for algorithm evaluation in this work. Chapter 4 presents the algorithm selected for posture classification, describes the data gathering methodology supporting the investigations in this work, and assesses the suitability of the generated classifiers for the case study application. Chapter 5 provides an investigation into the effect of transitions and demonstrates the method of handling them chosen here. Chapter 6 presents an investigation into the design space for a supervised machine learning based posture classifier, with a focus on C4.5 decision trees. Finally, Chapter 7 concludes on the work and provides the answers to the research questions outlined in Section 1.3.

Chapter 2

Literature Review

The work in this thesis addresses problems encountered when designing and deploying BSN-based posture classification systems in realistic (real-life) application scenarios. This chapter thus provides a discussion of the literature in four main areas: 1) existing applications of BSN-based posture classification, 2) the hardware platforms used and end-to-end posture classification systems developed, 3) data processing and the system design space for posture classification algorithms, and 4) methods of approaching the real-life issue of transitions between postures.

The aim of the literature review is to inform the thesis work, provide background and support for the developments proposed by the author and reveal the gaps in knowledge and practice in the field. Thus, the literature is drawn on as follows:

- Chapter 3: primarily, the review demonstrates that wearable systems capable of performing real-time on-node posture classification are not found in published works. Furthermore, posture classification for the application class including military and emergency personnel such as infantry and firefighters does not have a large body of published work available. Finally, the review aided in determining the common requirements and constraints imposed on BSN-based posture classification systems.

- Chapter 4: the review establishes a view of the common classifier algorithms used in the literature, thus informing the author's choices. The review also examines issues surrounding data gathering and reporting on classifier training and testing.

- Chapter 5: the review reveals a gap in existing works regarding the handling of transitions in applications faced with natural movement.

- Chapter 6: the review informs the investigation of the design space for on-body posture classification systems. Particularly it highlights the range of system parameters used in prior work.

This chapter is structured as follows. Section 2.1 provides an overview of applications that benefit from or require posture monitoring. Section 2.2 looks into BSN wearability. Section 2.3 presents an overview of the BSN platforms that have been used for building posture classification systems and investigates the choices of system parameters such as sensor positioning. Section 2.4 describes data processing techniques and posture classification algorithms used in the literature as gives an overview of the posture classification system design space. Finally, Section 2.5 describes the techniques used to classify or otherwise handle posture transitions.

2.1 Applications of posture and activity monitoring

High accuracy, autonomous, easy-to-wear, real-time BSNs for posture classification can benefit a number of scenarios in healthcare, leisure activities, and the workplace. The variety of systems and applications reported in the literature shows that posture monitoring is a relatively well covered research subject with a number of branches and applications from classification of daily activities [14, 42, 43, 79] to rehabilitation [125] to real time movement recognition for martial arts [54] and manufacturing environments [111].

This section aims to give a broad overview of the types of applications targeted in the literature, along with the requirements and constraints associated with them. The applications have been broadly categorised as: daily activities, dance and sport, virtual reality, healthcare, fall detection, and work-related. The issue of wearability of on-body systems is also discussed.

2.1.1 Daily activities

A common target application in the posture and activity classification domain is that of monitoring everyday activities. This is a broad application area, potentially encompassing all activities that are possible for a subject to perform in their day-to-day life. As such, the requirements and constraints placed on BSNs for this applications tend to be relatively generic, for example focusing on the use of small lightweight components. Often the primary goal of the research work is to show that the activities *can* be classified, and this is investigated using an offline classifier rather than producing a *deployable* real-time system. The large amount of freedom of movement and the unpredictable behaviour of humans make classification of daily activities a challenging goal.

The activities considered are generally taken from one of three main groups: activities within the home, leisure/fitness activities outside of the home, and office activities. Bao and Intille [14], for example, developed a classifier targeting twenty different home and leisure activities including walking, sitting while folding laundry, bicycling, and vacuuming. Huynh *et al.* [56] also focused on home activities, drawing a difference between low-level activities (such as walking, sitting, standing, eating, and washing dishes—

usually lasting up to several minutes) and high-level activities (such as cleaning the house—composed of multiple low-level activities and lasting as long as a few hours). Bharatula *et al.* [18] presented a system aimed at classifying daily office activities such as fast typing on a keyboard, moving a computer mouse, writing on a whiteboard, and opening a cupboard.

The system developed by Pansiot *et al.* [96] integrates an ear-worn activity recognition sensor (e-AR, which senses tilt and movement frequency spectrum) and ambient blob sensors that process a video signal to identify blobs or silhouettes and their motion based on optical flow. The system is capable of differentiating between sitting, sitting (sofa), standing, standing (head tilted), reading, eating, lounging, walking, and lying down. Although Pansiot's proposed system can be installed in a home environment, it is not suitable for deployment in situations where the subject would move to other unplanned locations due to the dependency on the subject being visible to the blob sensor. This problem motivates the benefits of developing systems that consist only of body-worn sensors: to enable mobility and functionality in an unconstrained environment.

Tapia [116] considered three activity categories: 1) static postures (such as lying down, standing, and sitting), 2) activities with multiple intensities (such as walking, rowing/arm ergometry, and cycling), and 3) other activities (such as running, calisthenics, and moving weights). Similarly, Ermes [42] targeted a mixture of indoor activities (such as lying, working on a computer, and standing reading a paper) and outdoor activities (such as playing football, running, rowing, and cycling). Data processing and postural information extraction is performed offline (rather than in real time during use of the system) which is common in studies/research in this domain, although Ermes suggests real-time operation as a future direction.

Laerhoven *et al.* [74] expanded on classification of daily activities by introducing a rhythm model that captures the user's normal daily pattern of behaviour. Activities included having breakfast, relaxing in the sauna, and watching TV. The rhythm model allows the system to perform classification of otherwise ambiguous sensor data. An example given in the work is of a user who gives a lecture every Tuesday afternoon. If the result of activity classification was inconclusive at that time then the rhythm model could improve the estimate.

To summarise, the primary motivation in this application area is the classification of a large number of postures and activities. The focus is on the performance of the classifier itself rather than on a deployable, real-time classification system.

2.1.2 Dance and sports

Dance and sport often employ motion caption systems to either 1) provide an interactive output based on the subject's movements or 2) provide feedback as to whether specific movements have been performed correctly (usually compared against a professional performing the same movements).

Bellis *et al.* [16] and Lynch *et al.* [83] designed wearable systems for an interactive dance environment. Bellis *et al.* present a node design based on stackable 25 mm boards, each fulfilling a specific purpose (such as sensing, power, and processing). Their aim was to produce small, modular, wireless devices with integrated signal processing to allow the implementation of data processing algorithms to reduce data transmission. Lynch *et al.* extend the concept put forth by Bellis *et al.*, designing wireless nodes based around 10 mm cubes incorporating the same capabilities as the 25 mm system. The primary consideration is further reduction in size. The inclusion of wireless communication is presented as an advantage as it avoids tethering the nodes (and thus subject) to a fixed location. They suggest that these same goals would also make the nodes suitable for site monitoring deployments in industrial plants.

Another system based around motion capture, is that developed by Young *et al.* [128] (also described by Arvind *et al.* [6]). This system is based on the Orient inertial sensor device which is dedicated to motion capture in fast-movement applications. An on-body network of fifteen devices are used to capture full body 3D movement in real time, which is translated into a real-time 3D model of the subject's motion using a rigid-body model. The on-body devices were designed with the goal of small size and low weight to increase the wearability, with wireless communication to allow the subject to move freely. Orientation estimates are calculated by each node in order to reduce the bandwidth requirements of the wireless communication links. The intended battery life is 1–2 hours to accommodate stage performance or multiple shorter takes in an animation studio.

For these systems it can be seen that the provision of small, light, wireless nodes is important to support the application in an unconstrained environment. They increase the wearability of the system, minimise deployment time, and allow the subject to move freely while wearing the system. Further, given the high sampling rates and potentially large number of sensors in motion capture applications, on-node data processing is used to reduce the bandwidth requirements.

2.1.3 Healthcare

Many examples of systems for posture classification exist in the patient care application area. Monitoring of patients undergoing physical rehabilitation is a common application, as the information required for assessment of progress is primarily motion-based. Other applications involve monitoring the status of

Table 2.1: Targeted postures and activities in a sample of healthcare related literature.

Authors	Application	Targeted postures/activities
Long *et al.* [79]	Healthy life styles	Walking, running, cycling, driving, and various sports
Pansiot *et al.* [96]	Monitoring elderly people	Walking, standing, reading, eating, sitting, lounging, lying down
Ying *et al.* [127]	Monitoring of people with Parkinson's disease	Step detection
Motoi *et al.* [93]	Rehabilitation	Walking speed, posture changes
Mathie and Celler [85]	Patients with congestive heart failure or COPD	Walking, falling, sit-to-stand, sit-to-lie, standing, sitting, and lying on back, front, and side
Zhang *et al.* [129]	Monitoring of correct posture	Sitting (back arched, leaning right, leaning left, normal), standing (upright, leaning forward), lying (right side, on back, face down)
Steele *et al.* [110]	Patients with COPD	Walking

patients with conditions that either impair their movement or cause involuntary movement. Continuous monitoring of patients at home is another important application area. Mathie *et al.* [85] point out that monitoring a patient at home allows early detection and treatment of health status changes and that "when monitoring the condition of patients with neurodegenerative or chronic diseases, a knowledge of their body movement and physical activity levels during the day is important". Table 2.1 gives a sample of common applications and targeted postures/activities in the literature, where it can be seen that the applications are broadly split into two categories: those that require classification of a wide range of daily activities and those that require classification of a specific subset.

Ying *et al.* [127] implemented a system that provides automatic step detection for patients with Parkinson's disease. The system implemented consisted of two dual axis accelerometers mounted on the patients' feet. Other work by Bamberg *et al.* [12] describes a wireless system for performing gait analysis (pattern of movement during locomotion). The platform includes two dual-axis accelerometers, three gyroscopes, four force sensors, and two bidirectional bend sensors integrated into a shoe. While their system was mounted within a shoe, one of the requirements was that it should not affect the gait of the subject. They present their system as an alternative to current methods of gait analysis. Gait analysis is generally carried out in a motion laboratory using expensive computer-based force and optical tracking sensors that must be attached to the patient or via visual observation by a clinician wherein the results are qualitative, unreliable, and difficult to compare across multiple visits. Their system provides repeatable quantitative results for longer periods of monitoring and allows gathering of data from subjects in their home (it is noted that patients often perform better in laboratory tests than in their "natural" environment, making such tests unreliable indicators of status).

Jovanov *et al.* [63] developed the ActiS sensor node, designed to be used as part of a wireless Body Area Network (BAN). This node incorporates a bio-amplifier and two accelerometers, allowing the monitoring of heart activity as well as the position and activity of body segments. The main focus is the node's use for monitoring the activity of physiotherapy patients outside of the laboratory. Jovanov *et al.* describe in depth the limitations of existing monitoring systems and the requirements for replacement systems. They note that current systems are not widely accepted for continuous monitoring primarily because of the amount of equipment required, the unwieldy wiring between individual components, and the lack of support for analysis of large banks of gathered data. Wiring has a negative effect on the patient's comfort and level of activity. Furthermore, the time taken to deploy such a system on the patient impacts on each individual monitoring session by adding non-productive time (i.e. time spent not gathering data) to the session. Clearly, wearable wireless technology could provide a solution to speed up deployment compared to wired monitoring, and allow attachment to the patient for a prolonged period in an unconstrained environment.

Motoi *et al.* [93] presented a method for monitoring posture and walking speed in the sagittal plane (the vertical plane from front to back dividing the body into left and right halves). The system integrates a trunk unit (with a sensor unit for measuring trunk angle) and lower body limb sensors (two sensor units with an accelerometer and gyroscope), and is based on an earlier system that suffered from several drawbacks (including the subject having to carry multiple pieces of equipment and a large quantity of cabling). The new revision of the system reduced the number of on-body units from four to two, simplified the wiring arrangements, and added additional sensors. This resulted in a more comfortable system for the subject to wear, along with better results due to the additional data made available from the increased sensor load.

Monitoring of patients with congestive heart failure and Chronic Obstructive Pulmonary Disease (COPD) was studied by Mathie *et al.* [85]. A single triaxial accelerometer attached to a belt placed in a pager casing was used for monitoring postures (such as standing, sitting, and lying on the left side, right side, front and back), metabolic energy expenditure and movement. It is noted that the placement of the sensor was not optimal, trading off some clarity in sensor data to improve the comfort and ease of attachment of the device (as reported by test subjects). Data was transmitted from the sensor node to a computer with no on-body processing.

In healthcare monitoring applications the wearability is vitally important, even at the cost of deliberately selecting a non-optimal sensing location in order to improve the subject's experience (Mathie *et al.* [85]). In some cases, existing systems are bulky or otherwise uncomfortable, limiting user satisfaction and the potential for long-term monitoring deployments. The need to address a subject's comfort becomes more apparent when home monitoring is considered as an alternative to short-duration laboratory or clinical monitoring. In these cases, the subject must wear the system for some large proportion of the day. Bao and Intille [14] note that subjects often feel self-conscious if the on-body system involves wiring that may be seen by others.

2.1.4 Fall detection

The area of fall detection is, in many ways, related to the healthcare application area. The primary goal is usually to detect falls and near falls either as a means of monitoring the progression of an existing health condition or, particularly with near falls, as a preventative measure for a subject who is suspected to be at risk of falling (for example elderly patients).

Li *et al.* [76] present a fall detection system using two sensor nodes (consisting of an accelerometer, dual-axis gyroscope, and single-axis gyroscope) placed on the chest and thigh. The system is aimed at classifying daily activities (walking, sitting, jumping, lying down, running, walking on stairs, and running on stairs), fall-like motions (quickly sitting), flat surface falls (falling forward, backwards, right, and left), and inclined falls (falling on stairs).

Jeon *et al.* [61] conducted three different studies looking into posture changes, falls, and daily activities. An accelerometer was placed on the chest and data was transmitted through Bluetooth to a PDA. The intention of the system is that when a fall is detected the system will display an alert on the PDA. If the alert is not responded to then an emergency centre will be automatically contacted. The advantages of such a system include portability, convenience, and low cost. In addition, the user interaction in normal situations is minimal, meaning that the device does not intrude on the subject's activities.

Nyan *et al.* [94] designed a system for classifying walking, sitting down, standing up, lying down, getting up, ascending stairs, and descending stairs, along with transitions between the postures (such as sitting to standing). An accelerometer was placed on the subject's shoulder and data was transmitted via Bluetooth to a laptop or a phone. An SMS is sent to a pre-defined phone number if the person falls or if an emergency button on the sensor node is pressed. Data was recorded from six subjects performing a predefined set of activities over a period of five hours and the system achieved an overall sensitivity of

98.83% and specificity of 94.98%. (Sensitivity is calculated as true positives divided by the sum of true positives and false negatives or $tp/(tp + fn)$ while specificity is true negatives divided by the sum of true negatives and false positives or $tn/(tn + fp)$).

Jafare *et al.* [60] proposed a methodology to classify four transition movements—sit-to-stand, stand-to-sit, lie-to-stand and stand-to-lie. The sensing system consisted of a sensor board incorporating a three-axis accelerometer and a GPS, transmitting the sensed data via Bluetooth to a medical centre via a laptop and a mobile phone. Two sets of experimentation were performed, the first with two subjects imitating 68 types of falls and the second with eleven subjects performing a prescribed regime of walking, sitting down and lying down. An overall classification accuracy of 84% was achieved for the four transition types.

It can be seen that a common trend in fall detection is to provide an automated call for help mechanism—Jeon *et al.* show an alert on a PDA that calls for help if the subject does not respond, Nyan *et al.* transmit an SMS message to another phone if a fall is detected, and Jafare *et al.* transmit data to a medical centre. This requirement appears more often in fall detection compared to other applications, likely due to the need for rapid response, because of the old age or fragility of the subjects, and the intention for the systems to monitor continuously on a subject that is otherwise unsupervised.

2.1.5 Work-related applications

Generally, the work-related activity monitoring applications fall into one of two categories: 1) monitoring the activities performed by the subject and 2) increasing the safety of the subject (through health or environment monitoring). This section describes a sample of posture monitoring systems targeting work-related applications, with a focus on workers operating in harsh environments.

Lukowicz *et al.* [81] researched the recognition of gestures for workers in a wood shop. The tasks performed during tests consisted of assembling a simple object made of two wood pieces and a piece of metal. Acceleration sensors were placed on both wrists and on the upper part of the right arm, along with a microphone on the chest and on the right wrist. The main activities performed were hammering, sawing, filing, drilling, sanding, grinding, screwing, and using a vice, for which classification was performed with an accuracy of 83.5%. Lukowicz *et al.* describe several advantages of wearable context-sensitive computing devices in the workplace, including the reduction in cognitive load (compared to accessing information on a traditional desktop computer) and the ability to automatically record actions performed by the worker. Recognition of tasks would allow automatic display of manual pages and alerting of the worker if steps are missed.

While not directly treating postures, Sung *et al.* [112] present a system for detection of shivering aimed at workers in cold climatic conditions, using the case study of Army Rangers on missions. The aim is to develop a real-time instrument able to classify cold exposure using non-invasive sensing methods and minimal processing power. The instrumentation used for testing consists of two accelerometers (on the right arm and chest), a 12-lead Electrocardiograph (EKG) set, heat flux sensors and rectal/oesophageal body temperature thermometers. Subjects were submerged between waist and chest deep in cold water (at either 10 ℃ or 15 ℃) and walked on a treadmill. Several models were tested, with the best—a Hidden Markov Model (HMM)—being "effectively 100% accurate" when providing core temperature classification (into "baseline", "cold", and "coldest" temperature regions) based on shivering activity. In addition to offline testing, a real-time evaluation was performed, giving similar results. The system is interesting as it involves real-time processing of acceleration data as well as mining over mixed physiological sensing data sets towards arriving at well-being decisions.

Kemp [67] provides details of several wearable body sensor systems for monitoring workers in dangerous environments. Of these, two systems utilised accelerometers for monitoring activity. The commercial LifeShirt system by VivoMetrics [77] includes a lightweight, machine washable chest strap with embedded sensors that monitor the subject's breathing rate, heart rate, activity level, posture (see description below), and skin temperature, while the LifeGuard system presented by Montgomery *et al.* [92] includes accelerometers and a variety of physiological sensors. In both cases, acceleration is measured at a single location on the chest (the LifeGuard system uses two 2-axis accelerometers arranged perpendicular to each other to capture the three independent axes of movement and one redundant measurement). The LifeShirt is described as providing postural information (the specifics of this information are not stated but appear to be the chest rotation relative to vertical), while the LifeGuard system provides only information on activity level.

Similarly to the LifeShirt, the commercial PSM Responder [101] system by Zephyr Technology Corporation is aimed at monitoring of workers in dangerous environments, including EOD operatives. Monitoring is performed by the BioHarness (available as a chest strap or integrated into a garment) and the captured data is transmitted to a PC where it is visualised. The specific outputs shown are the rotation of the chest from vertical, the individual axis readings, and the vector magnitude of the axis readings. The system thus does not classify posture in the sense considered by this thesis, giving only the angle of the chest.

Biswas and Quwaider [23] describe a posture classification system for monitoring the activity of soldiers in the field, which performs classification in real-time but not on-body. Primarily, the focus is on transmitting contextual information and safety alerts to other personnel so that appropriate decisions can be made with regard to rescue attempts or provision of medical aid. The specific postures identified are sitting, standing, walking, and running, and classification is performed using a PC rather than on-device. Extension to support online on-body classification of posture is suggested as a future work direction.

The work presented in this thesis is based around the provision of real-time posture classification via wearable nodes performing on-body processing of data. The primary reasons for this are 1) that on-body transformation of data into information reduces the transmissions required by the system to an external base station (as touched on by Curone *et al.* [35]) and 2) to allow autonomous operation in the event that the communication link to the base station cannot be maintained. The systems described in this section either do not perform posture classification (providing either the rotation angle of the device or the raw accelerometer data itself) or perform the classification using a PC to which all data is transmitted. Furthermore, crawling and kneeling are postures that must be classified in certain applications such as firefighter monitoring [26, 36], but are not classified by systems presented in the works reviewed.

A consideration in the design of wearable systems is that of wearability. This is particularly the case in medical applications (as described in Section 2.1.3) as the subject may be wearing the monitoring system for long durations of time. The next section thus reviews published works on wearability.

2.2 Wearability

As described in Sections 2.1.2 and 2.1.3 the wearability of an on-body system directly affects the user's satisfaction with the system and thus the likelihood of them opting to use it. Furthermore, for posture and activity monitoring applications, it is desirable to avoid impeding or restricting the wearer's movements.

When considering a system's wearability, Knight *et al.* [70] states that the level of comfort may be affected by a number of factors such as: physical dimensions of the wearable devices (for examples their size and weight), how they affect movement, and any pain caused either directly (for example rubbing against the skin or producing heat) or indirectly (for example muscle fatigue).

Gemperle *et al.* [47] conducted research to locate, understand, and define the locations on the human body where wearable objects can be placed without interfering with the movement of the wearer. The most unobtrusive locations found on the body for wearable objects were: (a) collar area, (b) rear of the upper arm, (c) forearm, (d) rear, side, and front ribcage, (e) waist and hips, (f) thigh, (g) shin, and (h) top of the foot. A location that is often suggested as suitable is the hip, as it is closer to the center of gravity and the weight of the object is therefore less perceivable [17, 41, 47, 121]. Dunne [41] states that the weight that can be easily carried differs across user groups (men versus women, adults versus children versus older adults, and so on), meaning that establishing acceptable limits requires the target user group to first be determined. Overall, while it is accepted that the weight of a wearable object should not hinder the subject's movement or balance, there is no precise measure given in the literature as to what suitable limits are.

The wearability of on-body systems can be enhanced via integration of the devices into clothing and the use of flexible electronics, both of which have been made possible by advances in miniaturisation of electronics generally and sensors specifically. Lymberis and Dittmar [82], Meng and Kim [88], and Patel *et al.* [98] all provide examples of monitoring systems targeted at health-related applications incorporating one or both of these techniques in order to increase wearability. Alternatively, sensors may be made both less intrusive and less visible by disguising them as jewellery. Asada *et al.* [7], for example, present a photoplethysmographic (PPG) sensor designed to be worn as a ring, while Degen *et al.* [37] present a system for fall detection built into a wrist-watch form factor.

In this thesis the sensor positions (lower arms, upper arms, chest, hip, ankle, thighs and calves) match those considered by Gemperle *et al.* [47] to be unobtrusive. The weight of the sensors is distributed across different parts of the body, while the weight of the on-body nodes is located around the waist in a pouch closer to the center of gravity. This should allow the system to be unobtrusive to the wearer.

2.3 BSN platforms for posture monitoring

Due to the need to use small and lightweight on-body components as shown in Sections 2.1.2 and 2.1.3, BSN systems have historically been built around computationally constrained hardware platforms with low power consumption ([39, 62, 49]). While the increasing popularity of smartphones and tablets has driven the production of lower cost and more capable platforms, these tend to consume more power and thus have a shorter battery life.

In surveying the literature in the posture and activity classification area, is clear that there has been no convergence on a particular hardware platform to support either data acquisition or end-to-end posture classification systems. Of a sample of 21 papers from the posture monitoring literature, as described in Table 2.2, 13 used an off-the-shelf monitoring platform, while 8 used custom node designs. In both cases, hardware platforms based on a range of technologies have been used. In addition, a variety of methods were used to transmit the data from the on-body nodes or to store it, as shown in Table 2.3.

Farella *et al.* [43, 44] and Young *et al.* [128] both developed a bespoke on-body node as the basis of their systems (the WiMoCa node and the Orient-2 system, respectively). The WiMoCa is a wireless sensor node containing an ATmega8 microcontroller, a TR1001 868 MHz radio chip, and an LIS3L02DQ accelerometer, while the Orient-2 is based around a dsPIC 30F3014 microcontroller, a CC1100 868MHz radio chip, and incorporates an MMA7260Q three-axis accelerometer, two HMC1052 two-axis magnetometers, and three ADXRS300 gyroscopes.

In the case of off-the-shelf hardware, the lack of commonality in platform choice between different researchers is likely to be due to a combination of reasons, with the main two being: 1) availability of particular platforms at the time the research was conducted and 2) the researchers' prior experience with particular hardware platforms or associated software (such as the OS or programming languages supported). Bespoke systems, by their nature, are varied and, in addition to the reasons of hardware availability and familiarity of the researcher with specific technologies, reflect a need to optimise performance for a given application. These factors may help to explain the wide variety of platforms (described in Table 2.2) using a number of different wireless communication protocols (as listed in Table 2.3). Practically, the criteria for platform selection specified by the researchers are often similar (most commonly around the devices being small and light and having a low power consumption) and much of the work presented could feasibly be implemented on a common platform. The devices selected (as listed in Table 2.2) are usually based around a 8- or 16-bit microcontroller with a small amount of RAM and frequently an integrated ADC, along with a radio (integrated with the microcontroller or as a separate chip).

Notably, it can be seen in Table 2.2 that in 18 of the 21 works reviewed, classification was performed on a PC or laptop rather than using an on-body device (two of the remaining three [39, 65] performed classification partially on the device and partially on a PC, while Maurer *et al.* [87] performed classification on the device but did not give any results). As already established in Section 1.1, autonomous operation is an essential feature for posture classification systems deployed in applications i) requiring a high degree of mobility for the wearer of the system or ii) where the postural information is used as input to another subsystem deployed in the on-body system. Deployment of the classifier on an on-body node was noted as a goal in several works, such as Curone *et al.* [35] and Zhang *et al.* [129]. However, the the fact that none of these algorithms have been deployed and evaluated on-node means that their real-life performance is not known. Deployment in this way is a crucial step in determining not only their performance generally with regard to the figures quoted as the state of the art, but also their suitability for deployment in applications requiring real-time on-body classification (such as the ones presented in Section 2.1.4 and 2.1.5).

The lack of standardisation on a given platform or communication method means that work cannot easily be shared and thus development effort is likely to be heavily duplicated between projects. Given

Table 2.2: Off-the-shelf and bespoke systems used in the literature.

	On body platform	uC	Sensors	Processing unit	References
OTS	Xsens MTx-28A53G25		Accels, gyros, mags	Laptop	[46]
	TelosB	MSP430	3D accel, 2D gyro	Laptop or PC	[62, 49]
	Witilt v2.5	PIC 16LF688	3D accel (MMA7260Q)	PC	[69]
	MicaZ (MPR2400 + MTS310)	ATmega128L	2D accel	Stargate+PC	[39]
	SunSPOT	ARM920T	2D accel (MMA7455L)	PC	[4]
	eWatch	LPC2106	2D accel (ADXL202)	eWatch	[87]
	Mica2Dot	Atmega 128	2D accel (MTS510)	PC	[103]
	Philips NWS Activity Monitor			PC	[79]
	Alive Heart Monitor		3D accel	PC	[20]
	HTC touch phone (HD T8282)		3D accel	PC	[129]
	Freescale D3965MMA7660FC in	MC68HC908JW32	3D accel (MMA7660FC)	PC	[122]
	Embla A10 (logger) ,		2D accel (ADXL202)	PC	[42]
Bespoke	TA unit	MSP430F149	2x2D accel (MXR7210GL)	TA unit+PC	[65]
	Wearable board	C8051F330	3D accel (MMA7260Q)	PC	[126]
	MITes	NRF24E1 (8051)	2D accels (ADXL202/210)	PC	[95]
	SensorButton	MSP430F1611	3D accel (LIS3L03AQ), microphone, light sensor	PC	[78]
	BodyANT	ATmega88V	3D accel (Bosch SMB380)	PC	[72]
	WiMoCA	ATmega8	3D accel (LIS3L02DQ)	PC	[44]
	Orient-2	Microchip dsPIC 30F3014	3D accel (MMA7260Q), mag., gyro	PC	[128]
		ADUC7027	3D accel (ADXL330)	PC	[35]

Table 2.3: Methods used for data communication / storage in literature.

Communication / Storage	References
802.15.4	[4, 27, 39, 49, 62, 65]
Bluetooth	[20, 35, 69, 87]
Bespoke 2.4 GHz protocol	[72, 78, 95, 126, 122]
Bespoke 868 MHz protocol	[44, 128]
Local storage	[42, 79, 129]
Wired comm.	[46]
Not stated	[103]

no "standard" accepted platform for this type of work and to allow for 1) faster prototype iterations and 2) the deployment of complex algorithms, a relatively powerful hardware platform was chosen for this work in comparison to the platforms described in this section. The chosen platform is described in detail in Section A.1. Furthermore, the classification system described in this thesis is capable of real-time on-body posture classification—of the works reviewed here, only one is presented as having this capability (described by Maurer *et al.* [87], though the posture set is smaller than that considered in this thesis and no results are given for online classification).

2.4 Classification methods for posture monitoring: data processing and system design space

Within the posture and activity classification literature, a variety of algorithms have been applied to the task of processing and classifying sensor data, commonly acceleration and/or gyroscopic data. Most approaches in the literature make use of some sort of machine learning to support classification, and as such require training data in addition to data used for validation/testing. This section describes some of the methods used, along with other relevant details regarding training and testing classifiers.

This section identifies the common steps taken in designing, implementing and evaluating machine learning based classifiers. The first step in this process is data gathering for training and testing of the classifier. This data is pre-processed (for example to adjust for sensor calibration and filter out noise in the samples) and then any required data features (such as variance or Signal Vector Magnitude (SVM), as described in Section 2.4.3) are extracted. Following classification, additional steps may be taken to attempt to increase the accuracy of the classifier (as described later in this section).

The literature revealed that the *number of subjects used when gathering experimental data* to support classifier training and testing is highly variable. Xue *et al.* [123] for example used data from 44 subjects, while Takeuchi *et al.* [114] used data from only 3 subjects—using one for training the classifier and two for testing it—and Huynh *et al.* [56] used data collected from one subject (performing the set of activities four times). It is commonly accepted that the training of a classifier that will generalise well to unseen subjects requires data gathered from a variety of subjects. However, the wide variety of subject numbers found in the literature and the general lack of justification for employing those numbers of subjects seems to imply that subject selection is based largely on availability and convenience rather than scientific rigour, particularly in cases where very low numbers of subjects are used. It is possible that the root cause of this is that the researchers are focused on designing and implementing a prototype posture monitoring system rather than performing rigorous validation of the classifier.

It is clearly important that rigour be applied to the data gathering process to ensure meaningful validation of the classifier. This thesis therefore brings well-evidenced analysis of the optimal number of subjects to be used for training (see Chapter 6) as well as introducing a robust methodology for testing and evaluation of classifiers (see Chapter 4).

The *processing of acceleration data prior to classification* is not commonly considered in the literature beyond extracted data features. In some cases, however, the data is filtered to remove noise or in an attempt to separate movement and gravitational signals. Khan *et al.* [69], for example, apply a moving average, Kang *et al.* [64] apply a low-pass filter, Karantonis *et al.* [65] apply a median filter to remove noise and a low-pass filter to separate gravitational from movement-related acceleration, Mathie *et al.* [84] applied a high-pass filter to remove the gravitational component of the acceleration and then a median filter to remove noise, and Sharma *et al.* [106] apply a moving average to remove noise and a high-pass filter to separate gravitational and movement acceleration. Other types of pre-processing are performed based on the data gathered and the needs of the analysis performed. Parkka *et al.* [97] perform calibration and resampling of the data, while Sharma *et al.* [106] perform a combined unit conversion and calibration step. Other than the examples above, sensor calibration is rarely mentioned in the literature. Calibration was considered by the author here and reported in Section A.1.3 on page 147. An analysis of the effect of uncalibrated sensor data on decision tree accuracy was performed. It was found that the classification accuracy was not affected by the lack of sensor calibration (with linear calibration and C4.5 decision tree classifiers as used here).

The next stage of data processing is *feature extraction*. The use of data features provides classifiers with more information than can be observed from the raw samples alone, particularly with regard to the evolution of the signal over time. Features are fundamental to the work in this thesis and their use in existing work is investigated in Section 2.4.3. Researchers have investigated a large number of features in an attempt to provide sufficient information for the classifier to distinguish the required postures (for example Frank *et al.* [46]). Lombriser *et al.* [78] and Atallah *et al.* [8], started with a large set of feature candidates and selected a small number of features based on an analysis of their outputs.

Once the data has been pre-processed and features have been extracted, *classification* is performed. There are a large number of classification methods used in the literature, often based around supervised machine learning algorithms. Methods used include Support Vector Machines [4, 53, 56, 71, 113, 123, 124], Bayesian techniques [8, 14, 46, 78, 79], Artificial Neural Networks [4, 69, 97, 126], Decision trees [14, 21, 31, 78, 79, 87, 97, 109], Fuzzy rule based methods [4], Hidden Markov Models [22, 39, 50, 56, 59, 90, 95, 114, 103] and K-means clustering [19, 58, 87]. Decision trees are relatively popular as they are computationally simple compared to many of the other methods, and provide, by and large, high classification accuracy. This makes them ideal for deployment on resource constrained on-body nodes. Their use in the posture classification literature is described in more detail in Section 2.4.2.

Following classification, *post-processing* is occasionally performed to increase the accuracy of the classification results. Parkka *et al.* [97], for example, applied a median filter to the output of the classifier to eliminate short duration postures.

2.4.1 Testing and evaluation

There are two important factors in the *testing and evaluation* of a classifier: 1) whether the testing is performed offline, based on data traces, or as part of a deployed classification system and 2) in the offline case whether the data is *truncated* to remove transitions, or is fully representative of natural movement. In the literature surveyed, the majority of the classifier testing was performed offline. There are some examples of realistic deployments for testing. Karantonis *et al.* [65] provided the results of online testing of their classifier with six subjects performing 12 tasked activities, and reported an overall accuracy of 90.8%. Dong *et al.* [39] presented results for a system deployed on a single subject in real-time, finding a classification accuracy of around 90% during various physical exercise type activities. Quwaider *et al.* [103] deployed their system on a single subject performing several activities such as sitting, standing, and lying, and reported an accuracy above 90%. Online testing of this type is important in confirming the performance of the classifier, particularly when performed in realistic (non-laboratory) settings.

In comparison the results above, the author's work resulted in classifiers that were tested and evaluated both offline and deployed, using real-time classification in realistic scenarios. An average classification accuracy of 96.3% was obtained during offline testing using 17 subjects. Real-time classification accuracy was 97% (as an average for five subjects performing tasked activities while being monitored).

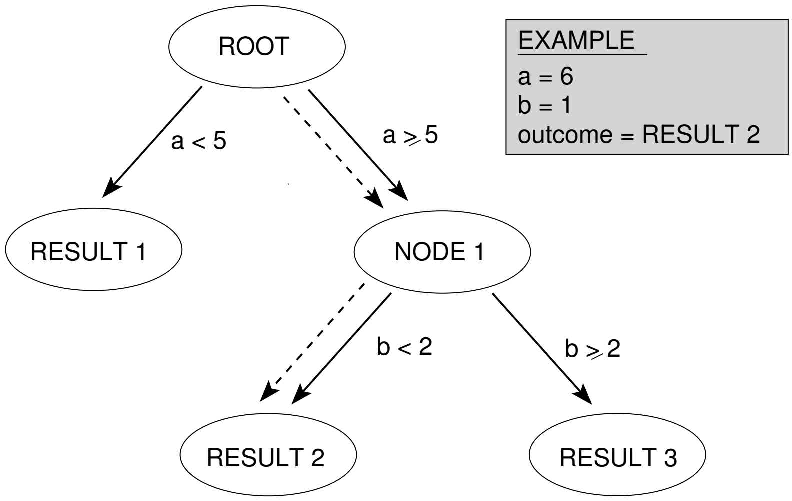

Figure 2.1: Simple decision tree example. *The dashed arrows indicate the nodes that would be visited given the example data values.*

The second factor, that of *truncation*, is rarely discussed in the literature surveyed. It is likely that in most cases only the results belonging to annotated periods of activity are used in calculating the classifier accuracy, as the technique for determine a "correct" result during a transition is not stated. In some cases, a particular point during the transition is assigned as the "change over" moment (for example, Parkka *et al.* [97] performed annotation via a tablet device during experimentation and created a change over point when selecting a new posture during each transition). The author here presents an alternative method in Chapter 5.

2.4.2 Decision Trees

Decision trees are a common classification tool that use a tree-like graph as a predictive model. A common decision tree based algorithm used in posture classification is C4.5 [14, 21, 31, 78, 79, 87, 97, 109], developed by Quinlan [102]. A simple example of a decision tree of this type is shown in Figure 2.1. An advantage of decision trees is that once trained, they are computationally simple and thus suitable for implementation on constrained embedded platforms (as evaluated by Maurer *et al.* [87] in terms of clock cycles and execution time per classification). Additionally, they do not contain loops and thus the time taken to perform a classification is bounded by the microprocessor specification and the depth of the tree. This makes decision trees a suitable choice for applications that require real-time classification (such as those described in Sections 2.1.4 on page 21 and 2.1.5 on page 22).

Based on results reported in the literature, the classification accuracy obtained for posture classification when using decision trees is similar across different research projects: for example, 84.3% [14], 86% [78, 97], and 89.3% [3]. Maurer *et al.* [87] reported classification accuracies of between 85.2% and 92.8% depending on the sensing location selected. The accuracy found by Tapia *et al.* [116] was only 56.3% when different intensities of activity were considered (for example, walking at various speeds), increasing to 80.6% when the intensities were merged into a single class per activity. Karantonis *et al.* [65] used a method similar to binary trees to distinguish between periods of activity and rest, recognise the postural orientation of the wearer (sitting, standing, walking, and lying on front, back, and side), and provide an estimation of metabolic energy expenditure. Recognition of postural orientation was carried out with 94.1% accuracy. The similarity in classification accuracy across the reported results may reflect a property of decision trees as applied to realistic data. As a rule, researchers will attempt to produce a system that matches or exceeds the performance of the work reported in the existing literature. For example, if the state of the art consists of classifiers capable of achieving a given accuracy, then researchers will attempt to create a classifier that either 1) achieves a higher accuracy or 2) achieves the same accuracy but expands upon the capabilities of existing classifiers (for example, classifying a wider range of postures, running on a more constrained hardware platform, or targeted at use in an application that imposes additional requirements).

To give an example beyond decision trees, He *et al.* [53] reported a classification accuracy of 92.3% using a Support Vector Machine based approach classifying data gathered from an ADXL330 accelerometer. However, in a publication based on the same work two years later, Xue *et al.* [123] reported an accuracy of only 86.8% using the same type of classifier and the same sensing device. A contributing factor for the decreased accuracy in this case appears to be that a larger number of postures are considered—ten in the later paper compared to four in the earlier one. Despite the decreased accuracy, the work is considered to be an advance because it increases the capability of the system while still providing an accuracy similar to that found elsewhere in the literature.

2.4.3 Posture classification system design space

The design space for posture classification systems is complex, encompassing a variety of system parameters that can impact the accuracy of the system (such as the positioning of sensors, the sampling rate, number of sensors and feature extraction methods used). However, similar to the selection of hardware, there is little commonality between systems reported in the literature with regard to these parameters. In addition to causing duplication of investigative work and lack of unification over the classifier design space, this may also make it difficult to meaningfully compare performance across different work, obscuring the reason for particular systems performing better or worse than others. To demonstrate the

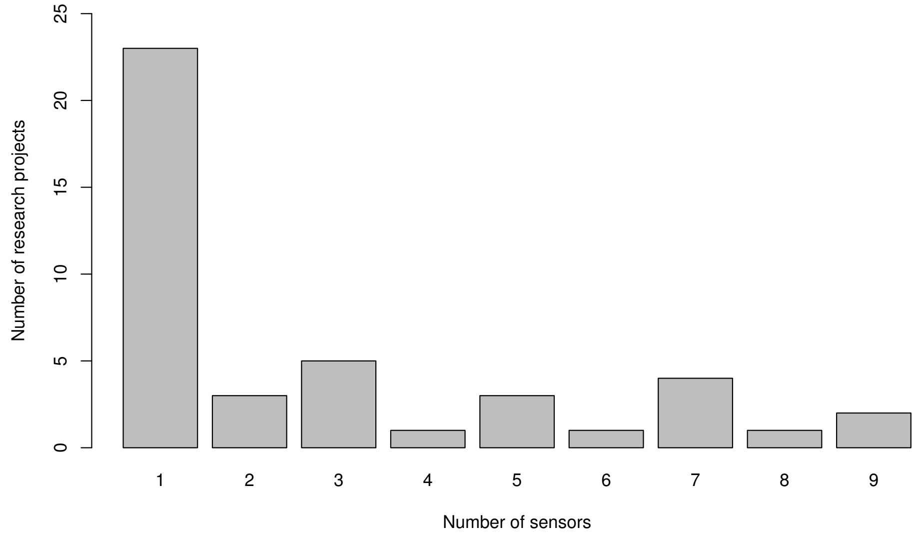

Figure 2.2: Number of on-body sensors used in posture/activity classification research.

range of options and design parameter settings, 43 papers were selected from the posture classification literature. This section analyses these papers in terms of five major parameters: the number of sensors used, the positioning of the sensors, the sampling rate used, the features extracted from the gathered data, and the postures targeted for classification.

Number and positioning of sensors

Figures 2.2 and 2.3 show the distribution of number of sensors and sensor positions respectively over the set of papers considered. It can be seen that there is little consistency in sensor number and placement between different research projects and that:

- The most common number of sensors used was 1. Common locations include the hip [46, 114, 84, 79, 65, 20], wrist [126, 90, 74, 59, 109, 108] and chest [68, 15].
- None of the papers surveyed used more then 9 sensing positions.
- The most common sensor positions are the hip and wrist. This is likely to be due in part to the method of placing the sensor on the subject, as wearability is increased by designing the node to be mounted on a belt or building it into a watch style housing.

Table 2.4 provides additional detail on the sensor positions used in the reviewed works, where it can be seen (in addition to the points already discussed) that there is little consistency between selection of the left or right side of the body for sensor placement.

Table 2.4: Sensor positions selected in the posture/activity monitoring literature. *R: right; L: left; T: trouser pocket; S: shirt pocket; J: jacket pocket; N: necklace. * only one location used at one time.*

	Ankle	Calf	Thigh	Hip	Wrist	Forearm	Upper arm	Chest	Pocket	Bag	Ear	Total sensors
[20, 46, 65, 79, 84, 114]				✔								1
[62]	✔ (×2)				✔ (×2)			✔				5
[97]					✔			✔				2
[52, 68]									✔ (T or J)			1
[39]		✔ (×2)	✔ (×2)	✔			✔ (×2)					7
[59, 74, 90, 108, 109, 126]					✔							1
[4]							✔					1
[56]			✔ (R)	✔ (R)	✔ (R)							3
[14]	✔ (R)		✔ (L)	✔ (R)	✔ (R)		✔ (L)					5
[42]				✔	✔							2
[22]	✔ (×2)		✔ (×2)	✔	✔ (×2)		✔ (×2)					9
[44, 43]		✔	✔					✔				3
[103]	✔ (R)		✔ (R)				✔					3
[116]	✔		✔ (R)	✔	✔		✔					5
[87]				✔	✔ (L)			✔ (N)	✔ (T and S)	✔		6
[95]				✔ (L)	✔ (R)			✔				3
[49]	✔ (×2)		✔ (×2)			✔ (×2)	✔ (×2)					8
[19]			✔(R)	✔		✔(R)						3
[72]			✔		✔			✔				3
[71]	✔ (R)		✔ (L)									2
[21]				✔					✔ (T)			1*
[8]	✔		✔	✔	✔		✔	✔			✔	7
[113]			✔ (R)		✔ (×2)			✔				4
[68]								✔				1
[63]	✔											1
[93]		✔	✔					✔				3

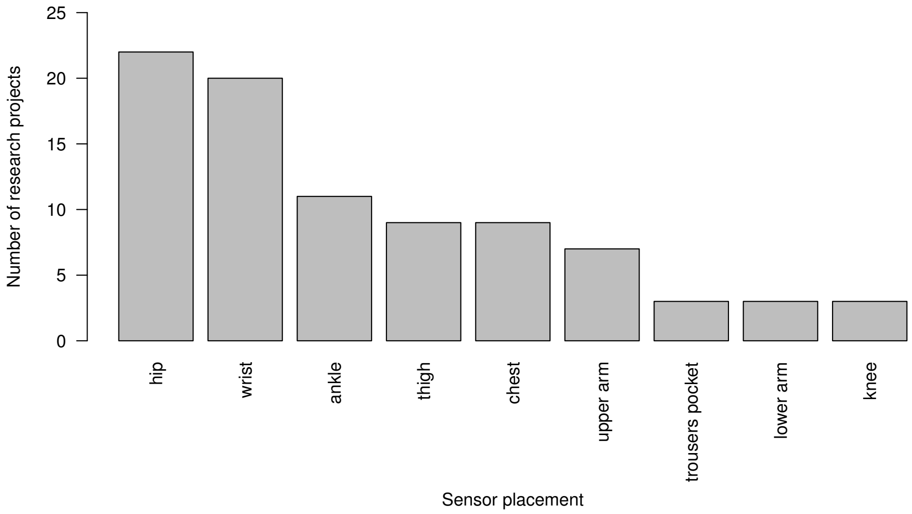

Figure 2.3: Sensor positioning in posture/activity classification research.

Sampling rate

The sensor sampling rates used in the papers surveyed varied from 10 Hz to 100 Hz, as shown in Figure 2.4. 100 Hz appears to be the most common sampling rate. No justification is generally given for the choice of sampling rate, though it may be related to the capabilities of the hardware platform used to collect the data. Despite 100 Hz being commonly used it has been shown that posture classification can be performed at much lower rates. Karantonis *et al.* [65], for example, stated that almost all measured body movements involved frequency components below 20 Hz and that even while walking 99% of the energy is contained below 15 Hz. Antonsson and Mann [5] concluded that 98% of the power for gait analysis is contained below 10Hz, and Bouten *et al.* [28] state that "[when] walking at natural velocity the bulk of acceleration power in the upper body ranges from 0.8–5 Hz". Using the results of Bouten *et al.* as an example, it is possible to conclude that 10 Hz is the lower bound (Nyquist rate) to capture the frequency components of walking.

In the work here, a sampling rate of 10 Hz was used in the implementation of the deployable real-time systems and Section 6.7 on page 116 demonstrates that this is sufficient to allow accurate classification of the set of eight postures considered here.

Extracted data features

As with the other parameters considered, there is little commonality in extracted data features used in the work surveyed. In some cases (particularly where only offline classification was performed), researchers

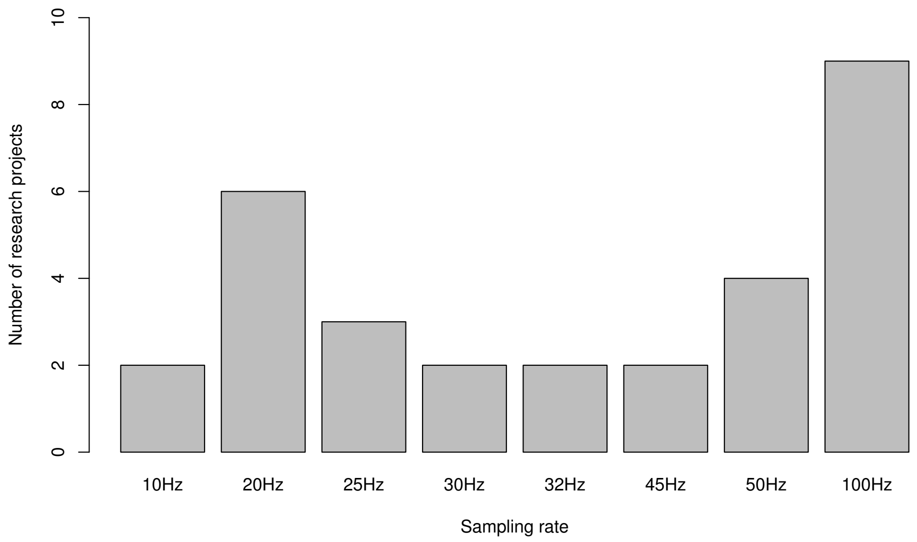

Figure 2.4: Sampling rates used for accelerometer data gathering in posture/activity classification research.

have used a large number of features simultaneously in an attempt to provide sufficient information for the classifier to distinguish the required postures. Frank *et al.* [46], for example, selected 19 features from a larger set. Other researchers have started with a large set of feature candidates and selected a small number of features based on an analysis of their outputs for the different postures. Lombriser *et al.* [78], for example, evaluated 8 features and selected 3 (mean, variance, and energy), while Atallah *et al.* [8] analysed a set of 13 features and determined that entropy, covariance, and energy provided the best results.

Over the set of papers surveyed, the most popular features were: mean (17 papers), variance (15 papers), energy (10 papers), Root Mean Square (RMS) (6 papers), and Signal Magnitude Area (SMA) (4 papers). Note that these features and others have been considered by the author here and are fully described in Section 6.4 on page 105, along with an analysis of classification accuracy delivered.

Targeted postures

Figure 2.5 shows the postures and activities classified in the work surveyed. The five most common postures are walking, standing, sitting, running, and lying, followed by more complex activities that add additional movements to (or combine) these five (for example, vacuuming will involve standing and walking, while cycling is similar to sitting but with additional leg movement). The choice of postures and activities to classify is heavily dependent on the application researched—whether it is focused on daily activities or sports for example.

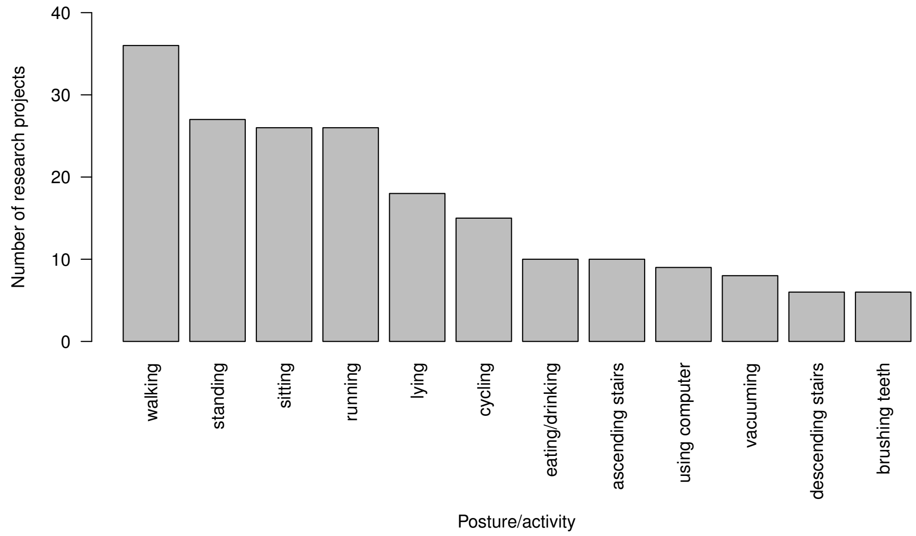

Figure 2.5: Postures and activities classified in the literature.

In the work here, the postures selected include four of the common postures studied by others (walking, standing, sitting, and lying), and add kneeling and crawling to these. The two additional postures appear to be rare in posture classification research and yet are required for applications such as firefighter monitoring [26, 36].

To summarise the discussion in this section, the design space for posture classification systems has a large number of options available (in the number of sensors used, the positioning of the sensors, the sampling rate used, the features extracted from the gathered data, and the postures targeted for classification) and little standardisation has occurred to date with regard to the best selections for any given purpose. Chapter 6 provides an in-depth analysis of the design space and gives advice with regard to selection of the optimal system parameters.

2.4.4 Main limitations of existing posture classification research

A workshop [80] highlighted the main areas in which existing posture classification literature was lacking. Primarily, the issues identified were related to inadequate or incomplete reporting, a lack of reasoning and justification for the work, and the use of unrealistic evaluation methods. For example, Amft [2] states that "the particular kind of work (e.g. user study, algorithm development, etc.) and deployed algorithm class (e.g. activity classification, repetitive or single-instance recognition, activity spotting) is typically not sufficiently specified upfront" and that "aspects of how evaluations are performed, are often left unspecified, unconsidered, or are just omitted in reports". Of the problem areas identified by the workshop participants, the following are investigated in this thesis:

1. Moving from laboratory-based evaluation towards addressing realistic challenges. This includes analysis of the proposed solution on realistic data, assessment of system robustness to realistic conditions, and establishing a link between what was studied and the real-world applications it could be applied to [24]. More fundamentally, a suitably defined motivating scenario or reasoning for the work presented is required [100]. The theme throughout this thesis is the requirements and constraints set by a class of applications including, as the driving case study, monitoring of EOD operatives during missions. Chapters 4 and 5 include realistic evaluation of a prototype posture classification system incorporating the algorithms proposed in this thesis.

2. Clear and detailed reporting of methodology (including annotation methods [29]) and evaluation methods (including analysis/performance metrics [2]) [34]. Chapter 4 describes in detail the data gathering and evaluation methodology followed for the work presented in this thesis.

3. Justification of system design considerations, including sensor placement, classification algorithms, and calculated data features [99, 100]. Each of these items is discussed in this thesis: system design is detailed in Chapter 3, sensor placement and calculated data features (among other system and data gathering parameters) are investigated in Chapter 6, and the selected classification algorithm is justified in Chapter 4. Furthermore, Chapter 5 gives an in-depth discussion of an extension to the classification algorithm to handle postural transitions in a meaningful way.

The concerns summarised above are related to the way in which results are reported—justifying why design choices were made, evaluating classifier performance on realistic data, and clearly stating how the experimentation was performed and the results were analysed. This information is important in allowing other researchers to reproduce the work or adapt elements of it and apply them to a different usage scenario. Furthermore, as pointed out by Amft [2], the results presented are often the “best case” results and do not take into account, nor do the authors explicitly report, realistic limitations. The areas of concern summarised here are thus a major driver for the reporting in this thesis, demonstrated in Chapters 4 (data gathering and evaluation methodology), 5 (transitions handling, an inherently real-world challenge), and 6 (system design parameters).

2.5 Handling transitions between postures

The handling of transitions is an aspect of posture classification that is not commonly considered during the classifier design stage. The focus of classifier research is often purely on creating a system that can perform classification of the selected postures. Data gathered to support the offline design and testing of the classifier is thus truncated to include only the postures of interest. This means that the accuracy is not as high as anticipated when the classifier is deployed for real-life monitoring, because the classifier is presented with data samples that it was not trained to classify.

There has been some effort within the literature to: 1) analyse specific types of transition to directly perform classification of them, and 2) develop fall detection systems based around transitions from a given posture to lying down.

A transition that is often targeted for detection or classification is that of sitting-to-standing (and the inverse, standing-to-sitting). Atallah *et al.* [8], Barralon *et al.* [15], and Jiang *et al.* [62] all looked into detecting such transitions. Godfrey *et al.* [48] investigated detection of sitting-to-standing and standing-to-sitting transitions to aid in classifying standing and sitting when using only a single sensor placed on the chest. Aloqlah *et al.* [1] looked into transitions between standing, sitting, and lying using data gathered from a three axis accelerometer mounted in a headband. A discrete wavelet transform is used in combination with a fuzzy logic inference system to detect the transitions and infer the current posture.

Fleury *et al.* [45] investigated transitions that occur in daily life such as sitting-to-standing and standing-to-lying down. A MMA7260Q accelerometer and a HMC1053 magnetometer are integrated into a data acquisition board, which is placed under the subject's left armpit. Classification is performed by first segmenting the signal using thresholds and then applying a wavelet analysis. Thirteen subjects performed a prescribed regime that involved sitting on a chair, moving around and lying on a bed, recorded by five webcams for verification. An accuracy of 70% correct classifications was achieved over the 13 subjects. Jafare *et al.* [60] also proposed a methodology to detect transitions between sitting, standing, and lying. Two sets of experimental data were analysed: 1) two young subjects imitated 68 types of falls, and 2) four young subjects and seven elderly subjects performed a prescribed regime of walking, sitting down and lying down. An overall accuracy of 84% was achieved for the four transitions. Khan *et al.* [68] investigated transitions between sitting, standing, walking, and lying, along with several postures, achieving an average classification accuracy of 97.9%.

Li *et al.* [76] present a fall detection system using two sensor nodes placed on the chest and thigh. Data was collected from three male subjects undertaking the following activities for 5 seconds each: daily activities (walking on stairs, walking, sitting, jumping, lying down, running, running on stairs), fall-like motions (quickly sitting down upright and reclined), falls on a flat surface (falling forward, backward, right, and left), and falls on an inclined surface (falling on stairs). When a transition to lying is detected, the acceleration and angular velocity are analysed to determine if the transition was intentional. If the transition was not intended then it is classed as a fall.

Detection of transitions can be useful in some applications such as detecting falls or as additional information to aid in classifying postures. However, the work in this thesis required a different approach to that found in the literature surveyed. Here, the goal is not to detect the transitions as such, but to diminish their negative impact on the accuracy of a classifier when used in a real-life deployment. Chapter 5 describes the method used to achieve this goal.

2.6 Summary

Posture classification using a BSN-based system is the topic of a wide variety of research projects, targeted at applications ranging from providing long-term remote healthcare to increasing work safety. In these applications, the benefit of such systems lies in either 1) replacing bulky existing equipment with smaller lighter on-body sensing nodes or 2) in providing postural information where none was previously available. The systems presented in the literature generally focus on two areas of BSN-based posture classification: evaluation of posture classifier performance (usually offline), and the use of small and light on-body devices to increase wearability. Real-time on-node classification of posture is required for several applications (such as monitoring of firefighters) but is generally not performed, despite being noted as a goal in some cases.

Despite the wide range of applications and the number of research projects targeting them, the design space for posture classification systems has not been extensively analysed and there is little commonality in the hardware platforms used. For example, a survey of 21 papers showed that 13 of them used off-the-shelf devices (12 different devices in total) and the remaining 8 each developed their own hardware platforms. Of the projects, 17 used wireless communication: 802.15.4 (6 projects), Bluetooth (4 projects), generic 2.4 GHz (5 projects), and 868 MHz (2 projects) radios. The system parameters selected in the literature also showed little commonality in terms of the number of sensors used, the positioning of the sensors, the sampling rate used, the features extracted from the gathered data, and the postures targeted for classification. Chapter 6 presents an investigation into the design space for posture classification systems and provides guidelines for developing a posture classification system targeted at real-life application deployment. Reporting of data gathering methodologies and robust classifier evaluation methods are considered to be lacking in the posture classification literature. The methodology used here is described in Chapter 4.

Handling of transitions within the literature surveyed is focused on detecting specific types of transition either as an end in itself, to aid in classifying other postures, or as a step towards detecting falls. Detection of transitions involving sitting and standing is particularly common within the literature surveyed. The approach taken within this thesis, however, is to target the classifier only at the specific postures required and to implement a means of reducing the negative impact of transitions on the classifier accuracy. Chapter 5 provides details of the technique used.

The work in this thesis builds upon works from the literature with regard to the system design and classifier selection. Specifically, the C4.5 decision trees used here (see Chapter 4) were also used for classification by Bao and Intille [14] and Tapia [116], the positioning of accelerometers on the body (as shown in Chapter 6) are similar to those used by Guenterberg *et al.* [49] (and match those later used by Xu *et al.* [122]), and the set of features extracted from the raw sensor data (presented in Chapter 6) was based on the work of Ermes [42], Bao and Intille [14] and Mathie *et al.* [85].

Chapter 3

Posture classification platform

3.1 Introduction

It was shown in Sections 2.3 on page 25 and 2.4 on page 28 that the literature provides no standard way of designing and building posture classification systems, though there are some commonalities in broad terms with regard to the processing stages implemented. This chapter presents the concept and design of a general end-to-end platform for real-time posture classification. The platform presented is named Class-act, since it is a classifier of activity. The platform is targeted at two usage scenarios:

- A self-contained system providing postural information to an external system (for further processing/modelling or visualisation).
- A configurable investigative instrument for posture-related investigations.

The Class-act platform architecture specifies two roles that individual BSN sensor nodes can perform: 1) Primary Nodes responsible for classification and relaying configuration commands to the Secondary Nodes and 2) Secondary Nodes that are responsible for gathering data and passing it to the Primary Node. The architecture allows for a single Primary Node along with any number of Secondary Nodes as required by the application. The hardware used is not specified by the platform design, allowing flexibility in specific implementations (in the use of less-wired or completely wireless communication, the number of sensors per node, and so on). Section 3.2 describes the platform design in detail.

The work in this chapter, in combination with Chapter 4, forms one of the three contributions brought by this thesis (to quote from Section 1.5):

- The design of a platform that allows real-time on-body classification of static and dynamic postures—a capability not present in existing work. The specific posture set consists of six static postures (sitting, standing, kneeling, and laying on back, front and one side) and two dynamic postures (walking and crawling), of which kneeling and crawling are not commonly considered in the literature. Classification is performed on a small, light embedded device using a simple easy-to-implement algorithm. The classification algorithm used is a C4.5 decision tree, with a temporal feature (windowed variance) to aid in distinguishing dynamic and static postures.

The chapter is structured as follow: Section 3.2 presents the Class-act platform design and architecture. Section 3.3 describes two application examples demonstrating the additional requirements that such applications can impose. Section 3.4 presents a prototype system implementation example. Finally, Section 3.5 summarises the work presented in this chapter.

3.2 Design/architecture

This section presents the design and architecture of an end-to-end posture classification system, along with the hardware requirements for such a platform.

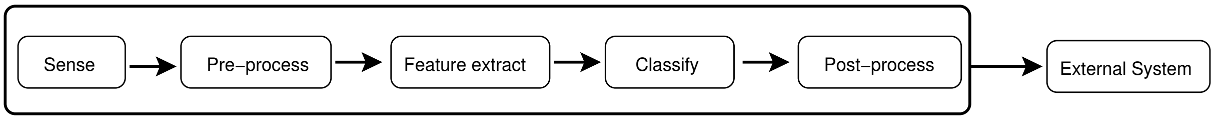

Figure 3.1: Generic data processing chain for posture classification systems.

Figure 3.1 shows the general data flow specified by the Class-act platform design, derived from systems presented in the literature (as described in Section 2.4 on page 28), with the following stages:

Sense Acceleration/gyroscope data is sampled from the attached sensors.

Pre-processing The sampled values are processed to prepare them for use by the classifier. This may involve steps such as filtering or adjusting for calibration.

Feature extraction Data features (such as variance) are extracted to aid in classification. Chapter 6 presents an investigation of the effect of a number of features on classification accuracy.

Classification The sampled values and extracted features are used to determine the current posture of the monitored subject. Chapter 4 describes the classifier used in the work here.

Post-processing The classifier output is manipulated to achieve goals such as reducing the number of posture changes identified. For example, Chapter 5 presents transition smoothing filters—a post-processing step to solve the problem of rapid classifier output changes during postural transitions.

The postural information generated via this processing chain is provided to an external system. The external system may be one of several different devices such as a remote PC used for observation of the subject or another on-body system that performs further processing/modelling using posture as an input. To simplify the discussion here, these will all be referred to as "external system" unless the distinction is important to the discussion.

In order to provide the described processing stages, the Class-act platform architecture specifies two roles for BSN nodes:

Secondary Nodes A Class-Act system[1] contains any number of Secondary Nodes. The actual number of these nodes in an implemented system will be based on the application requirements (such as a need to keep system components at different locations on the body physically separate) and hardware constraints (such as the maximum number of sensors a given node can support). The Secondary Nodes are responsible for gathering data, performing pre-processing and feature extraction, and passing the data to the Primary Node.

Primary Node A Class-act system contains one Primary Node. The Primary Node enables the system to meet the need for on-body classification of posture. The node is responsible for gathering data and performing pre-processing and feature extraction (as with the Sensing Nodes), but also 1) aggregates data from all nodes, 2) performs classification, 3) transmits postural information and/or sensed data to an external system, and 4) relays configuration commands from an external system to the Secondary Nodes (for the investigative system usage scenario).

Figures 3.2 and 3.3 demonstrate the data flow for each of the two usage scenarios specified (self-contained posture classifier and investigative instrument). The data flow through the Class-act system in each scenario consists of stages representing data gathering, pre-processing of the data (calibration, filtering, and so on), feature extraction, posture classification, post-processing (such as transition smoothing filters), transmission of the data, and remote configuration capability. It can be seen that the processing chain shown in Figure 3.1 is suitable in both cases. In fact, the only difference in node capabilities between the two scenarios relates to the external system—the investigative instrument is capable of being reconfigured during use. Examples of configuration options that may be implemented are:

- Selection between a set of classifiers.
- Selection of the data feature to extract.
- Selection of the transmission mode to use: 1) transmission of all sensed data and postural information, 2) transmission of postural information only, or 3) transmission of posture changes only (i.e. an event-driven transmission mode).

[1] "Class-act system" is used as shorthand to refer to any system implemented based on the Class-act platform design.

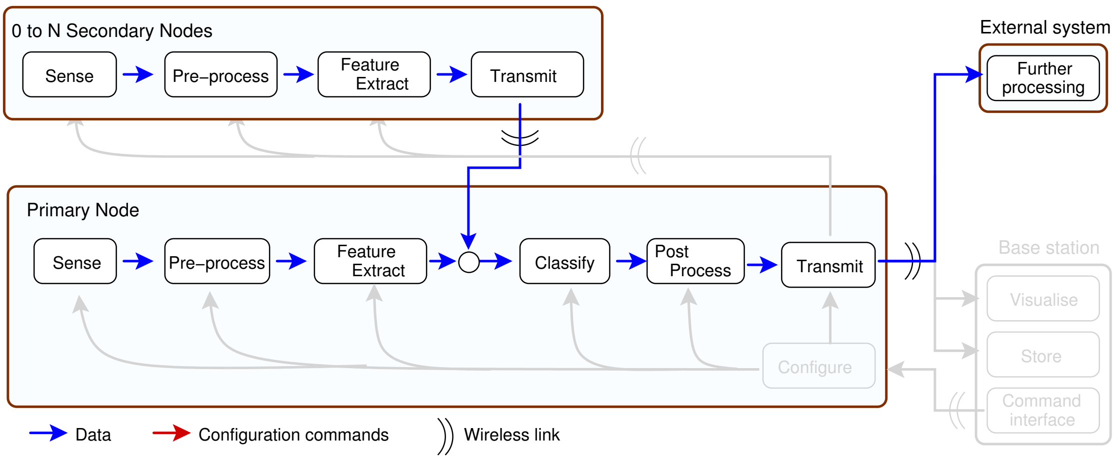

Figure 3.2: Self-contained system usage scenario data flow.

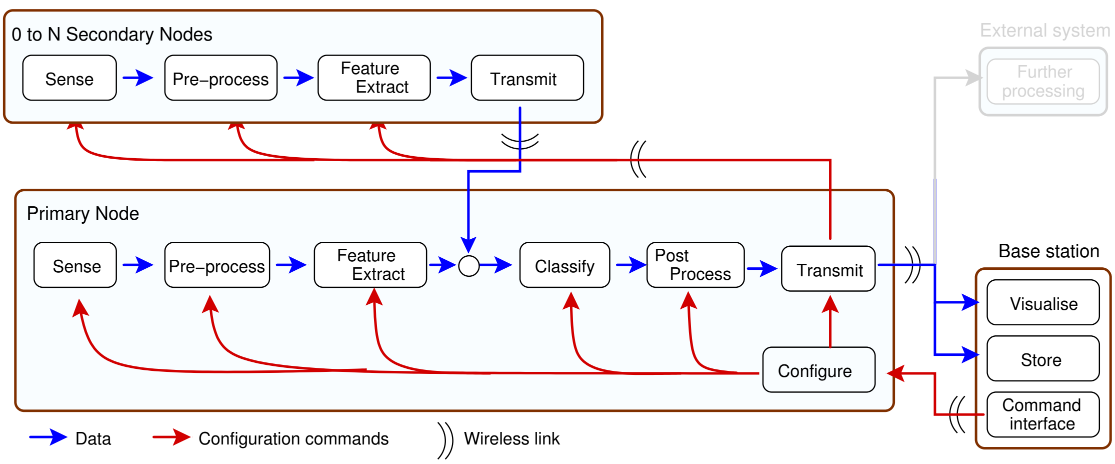

Figure 3.3: Investigative instrument system usage scenario data flow.

The self-contained system, on the other hand, will be pre-configured with a particular configuration and will only transmit posture changes—this avoids transmission of redundant information and, will therefore extend the battery life. Section 5.4 on page 91 demonstrates the transmission reduction obtained in this mode for a prototype system implemented based on the Class-act platform.

The Class-act platform design does not specify the hardware needed to implement Primary and Secondary nodes. However, there are several generic requirements that can be derived from the literature and from practical system implementation considerations:

- The BSN hardware must include sensors that can provide data relevant to posture classification. Usually this will be accelerometers and/or gyroscopes.
- The BSN hardware must be capable of sampling at a sufficient rate to allow accurate classification, particularly where time-dependent data features are extracted.
- The BSN hardware must not restrict the subject to a certain area, leading to two sub-requirements:
  - When classified posture is provided to an external system or base station (not located on the subject's body), communication of postural information must be performed wirelessly.
  - The system must be battery powered (or self-powered in some alternative way).
- The BSN hardware must be light and unobtrusive so that the subject's comfort and mobility/natural movement are not affected by use of the system.
- In the case of the self-contained system usage scenario, the system must be capable of real-time operation (as defined in Section 1.1 on page 2).
- The BSN hardware must be capable of supporting the pre-processing, classification, and post-processing to be performed (while maintaining real-time operation in the self-contained system usage scenario).

3.3 Class-act platform application examples

This section provides concrete examples of how a system based on the Class-act platform architecture can be targeted at specific applications. The first is targeted at use in EOD operative monitoring (see Section 1.2 on page 4), where postural information is supplied to a second system that performs heat stress prediction and helmet CO_2 concentration modelling. The second is targeted at use as an investigative laboratory instrument.

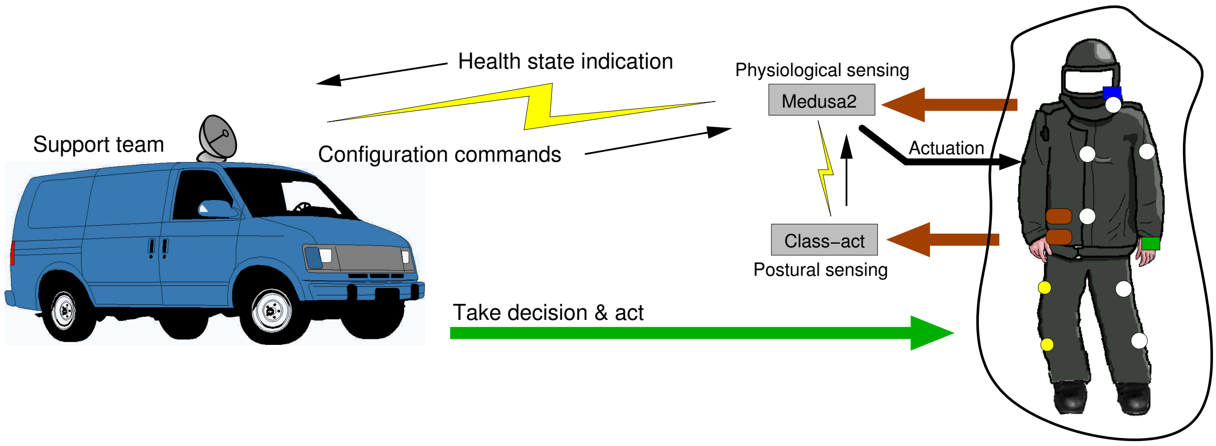

Figure 3.4: Self-contained system application example: EOD operative monitoring. *Sensor types shown on the operative for demonstration purposes: white: skin temperature; yellow: accelerometer; blue: helmet* CO_2*; green: pulse oximeter (pulse rate and blood oxygenation).*

3.3.1 Self-contained system

Figure 3.4 gives an overview of an application example for the self-contained system usage scenario—that of EOD operative monitoring. As described in Section 1.2 on page 4, the EOD operative monitoring example is only one possible application of the work here, with the broader class including applications such as monitoring of firefighters and infantry. The expectation is that the EOD application can be generalised to the other applications within this class. In this example, the Class-act system provides postural information to the Medusa2 system [66], which performs further modelling and prediction with regard to the health status of the operative. The Medusa2 system was developed as a monitoring system to enable increased safety of EOD operatives through: 1) monitoring of physiological parameters, 2) inference of health state information from the gathered data, 3) autonomous actuation of the in-suit cooling system, and 4) provision of appropriate data, information and alerts to both a remote observer and the operative. Two of the algorithms implemented within the Medusa2 system (specifically, real-time prediction of 1) the risk of UHS occurring in the operative and 2) helmet CO_2 concentration) require posture as an input due to the large influence that posture has on the evolution of the state of the system.

The EOD application brings several requirements beyond the generic ones described in Section 3.2:

- The eight postures specified for classification in the work here map to the postures required in the EOD application—the system must therefore be capable of classifying these.

- Any instrumentation on the upper and lower body must be physically separate. This is to aid the operative in donning and removing the EOD suit and to prevent damage to the system at those

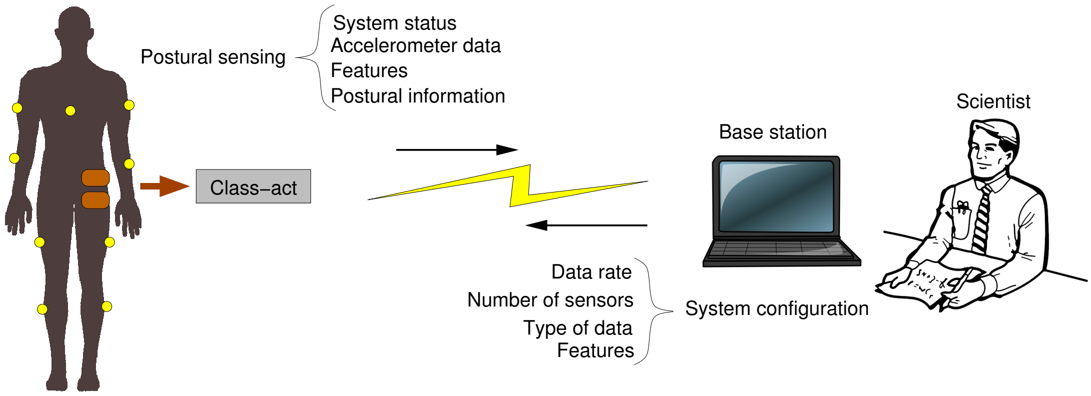

Figure 3.5: Investigative instrument application example.

times.

- The wireless communication method used by the Class-act system must match that used by the Medusa2 system.

- Co-location of sensors between the Medusa2 and Class-act systems would be preferable in order to reduce wiring for data and power.

3.3.2 Investigative instrument

Figure 3.5 gives an overview of an example implementation for the investigative instrument. This maps directly to one of the usage scenarios for the platform and so does not introduce new requirements to the extent that the self-contained system example does. The only additional requirement, for the sake of convenience in investigations, is:

- The hardware platform should allow the number of attached sensors to be varied as required between deployments.

3.4 Prototype implementation example

A prototype system was implemented meeting the requirements of the two example implementations described in Section 3.3. Due to the use of the Class-act platform design, both application examples were supported via a single prototype implementation. Full integration with the EOD suit was not a goal for the prototype system, it serves as a proof-of-concept for the EOD application. The implemented prototype has been deployed in the work here for evaluation of the algorithms presented in Chapters 4 and 5. Full hardware details are given in Appendix A.

In response to the requirements presented in Section 3.2 and 3.3, several system design choices were made with regard to: 1) the number of on-body nodes, 2) the communication methods used (sensor to node, node to node, and node to other components in the system), 3) the types of data/information transmitted from the on-body nodes, and 4) the number and location of the on-body sensors.

3.4.1 Number of on-body nodes required

The EOD application requires physical separation of upper and lower body sensing, while the investigative instrument application does not specify a requirement with regard to the number of nodes. Based on this, the decision was made to use two on-body nodes (one for the upper body and one for the lower body) in the prototype system for consistency across both applications. During initial testing of the implemented prototype, it was found experimentally that a single node could not reliably gather data simultaneously from more than eight of the sensor boards used. The decision to use two on-body nodes when building the prototype system therefore means that a total of up to 16 sensors can be supported. Note that Section 6.6 demonstrates that accurate classification can be performed using only two sensors (on the thigh and calf) and therefore only one node is needed for a final implementation of the system for the EOD application.

3.4.2 Communications

A generic requirement of the Class-act platform is that of wireless communication from the on-body nodes to the external system/base station. This applies for the investigative instrument as the base station is located away from the subject and to the EOD application as the communication method must match that of the Medusa2 system (in this case, Bluetooth). Bluetooth was selected as it meets the needs of both applications. For simplicity, Bluetooth was also used for node-to-node communication. Neither of the applications are expected to involve communication distances greater than that allowed by Bluetooth. Based on the scenario descriptions given, there are two transmission modes that must be supported by the on-body system: 1) transmission of postural changes only (for the EOD application), and 2) all three communication modes with online selection (for the investigative instrument). The prototype system supports each of these modes.

A wired connection was selected for sensor-to-node communication since: 1) wired links are simpler and less error-prone than wireless links, and 2) power could be supplied to the sensors alongside the data connections, reducing the size, weight, and complexity of the sensors compared to a self-powered wireless

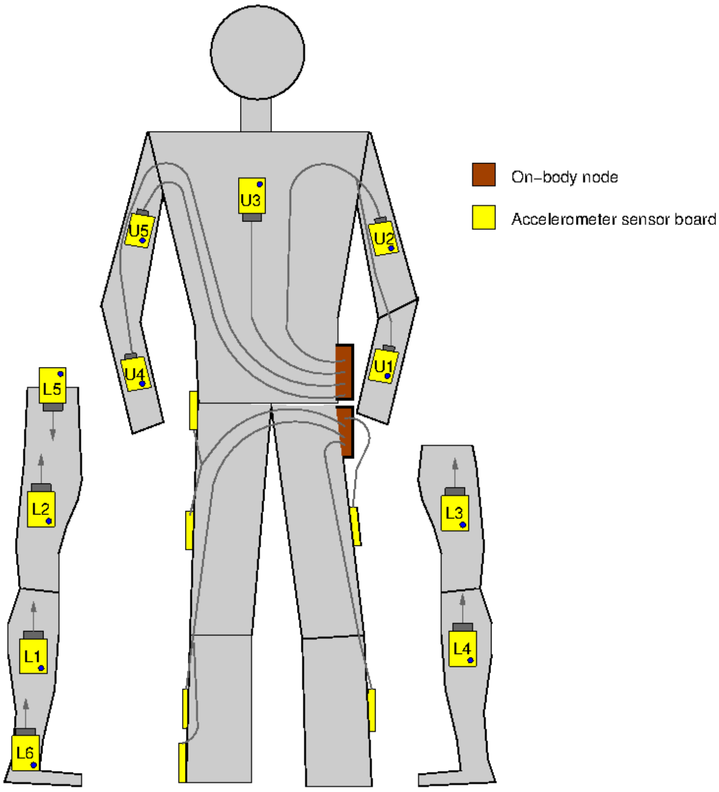

Figure 3.6: Positioning of sensors and nodes on the body for the prototype system.

solution.

3.4.3 Sensor positioning

The position of sensors for the prototype system is a superset of the Medusa2 locations to simplify tight integration of the two systems. If the two systems were to be merged into a single combined monitoring system (sharing the same hardware nodes), co-location of the sensors would reduce the amount of wiring needed between the sensors and nodes. The temperature sensors for Medusa2 are located at the subject's neck, upper arms, chest, abdomen, thighs, and calves. The final locations selected correspond to the distinct body segments: upper arms, lower arms, chest, thighs, and calves. These locations match those used by Xu *et al.* [122] and are similar also to those used by Guenterberg *et al.* [49]. An investigation was conducted (described in detail in Section 6.6 on page 111) towards determining the optimal set of sensor placements to provide sufficient data for accurate posture classification while also minimising the number of worn body sensor. Figure 3.6 shows the sensor locations and connections to the on-body nodes given sensor placements on the subject's chest, upper arms, forearms, calves, thighs, hip, and ankle. Each of these eleven locations is considered as a potential mounting position for a triaxial acceleration sensor (see

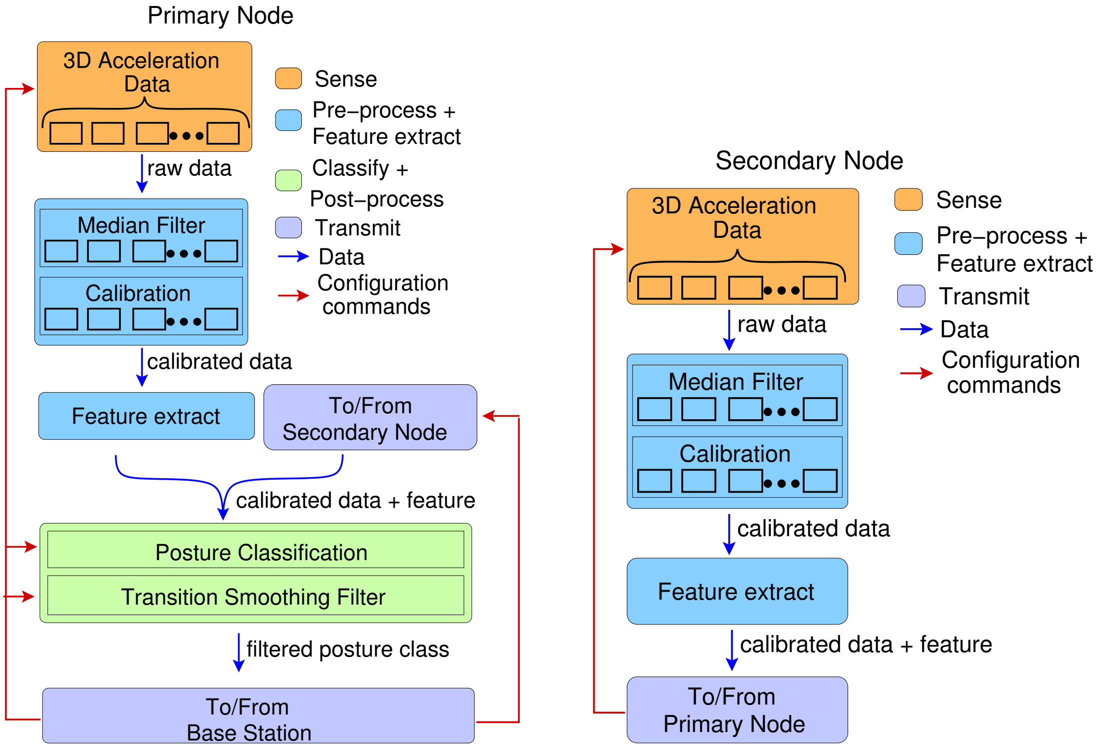

Figure 3.7: Data and information flow for the Primary Node and Secondary Node.

Section 6.6 for discussion of the effect of choosing specific location subsets).

Due to the directional nature of acceleration measurement, consistency of orientation of the sensors is important for accurate classification. In order to ensure this consistency, reference diagrams were produced to show the location and orientation of each sensor (the boards were not packaged for the prototype, so their orientation was clear visually). While every effort was made to match the diagrams closely, it is natural that some small errors in orientation would occur from one trial to the next, particularly across subjects with varying body builds. The trials conducted therefore established an informal test of the effect of small inconsistencies in mounting, which was found to have little noticeable impact on classification accuracy—the variation in accuracy results was generally small, with a standard deviation of 3.7% when WVar was used. This forms an upper bound for the effect of mounting inconsistency (assuming the experimenter is following the diagrams correctly).

3.4.4 On-body node software

Figure 3.7 shows the data flow and processing steps for the Primary and Secondary Node within the prototype system. The stages are as follows:

Sense At the Sense stage, data are gathered from the attached acceleration sensors. The classification

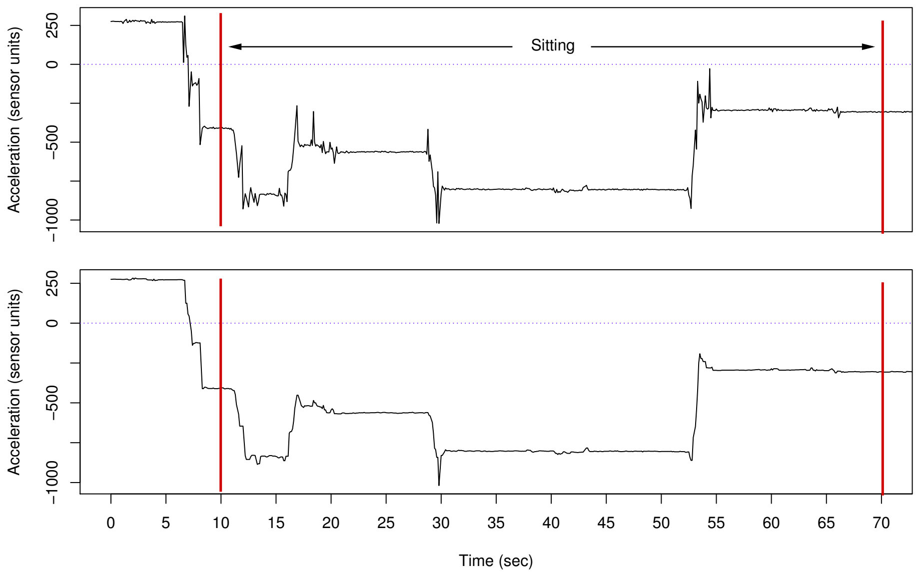

Figure 3.8: Median filter applied to x-axis of an accelerometer placed on a subject's calf while sitting. *Top: original accelerometer data. Bottom: filtered data.*

accuracy obtained with varying numbers of sensors is described in Section 6.6 on page 111.

Pre-process The Pre-processing stage consists of two data manipulation steps: median filtering and calibration. First, a median filter with a window size of three samples is applied to remove spurious data "spikes". The median filter was also used in this way by Karantonis *et al.* [65]. The median of an array of values, $\widetilde{x}$, is calculated as

$$\widetilde{x} = \begin{cases} Y_{(w+1)/2} & \text{if } w \text{ is odd} \\ \frac{1}{2}\left(Y_{w/2} + Y_{1+w/2}\right) & \text{otherwise} \end{cases}$$

where Y is a sorted array of values and w is the number of values in the array. Given $w = 3$ (a fixed window of three values), $\widetilde{x}$ is thus always obtained from Y_2. Figure 3.8 shows an example of the median filter applied to sample data gathered from the x-axis of an accelerometer placed on a subject's calf while sitting. It can be seen that the data are smoothed to an extent, removing "spikes" that could lead to misclassification of the subject's posture.

The second step is to adjust the accelerometer sensor data in order to compensate for sensor calibration errors. The process used for calibration is discussed in detail in Section A.1.3.

Feature extract The Feature extraction stage consists of features such as Windowed Variance (WVar) being extracted from the calibrated raw data. An analysis of the classification accuracy benefits gained from feature extraction is given in Section 6.4 on page 105.

Classify Prior to classification being performed, the data from the Primary and Secondary Node are appended together (or concatenated) to form a single data vector containing all of the body acceleration data and data features. The data is then provided to a classification mechanism. The method used here is a C4.5 decision tree trained using experimental data, as described in detail in Chapter 4.

Post-processing The Post-processing stage consists of applying a transition smoothing filter to the classified postural information to improve the overall accuracy and output stability. The particular filter used may be selected by the user and the filters implemented here are described in Chapter 5.

Transmit Transmission from the Secondary Node to the Primary Node is performed wirelessly and includes all of the gathered acceleration data from the attached sensors. Transmission from the Primary Node to the base station has several modes as described in Section 3.2. In all cases, the postural information itself is transmitted (continuously or only when the posture changes) and the acceleration data may be transmitted depending on the application.

3.4.5 Base station software (visualisation and system configuration)

As part of the prototype system implementation, a visualiser was developed to support the investigative instrument application. The EOD application is already provided with a visualiser developed as part of the implementation of the Medusa2 system [66]. The visualiser developed for the investigative instrument application is described in this section and provides several options with regard to system configuration.

A screenshot of the visualiser developed is shown in Figure 3.9. The visualiser is split into three areas. The left-hand side is dedicated to sensors and communications, the central area shows the current classification result, and the right-hand side shows options related to data acquisition and processing. The visualiser is implemented in Python using the wxPython Graphical User Interface (GUI) libraries, providing portability between operating systems. As the set of postures that the system can classify is pre-defined, the current posture is displayed using one of a set of images (one for each classifiable posture).

The specific functions supported by the visualiser are:

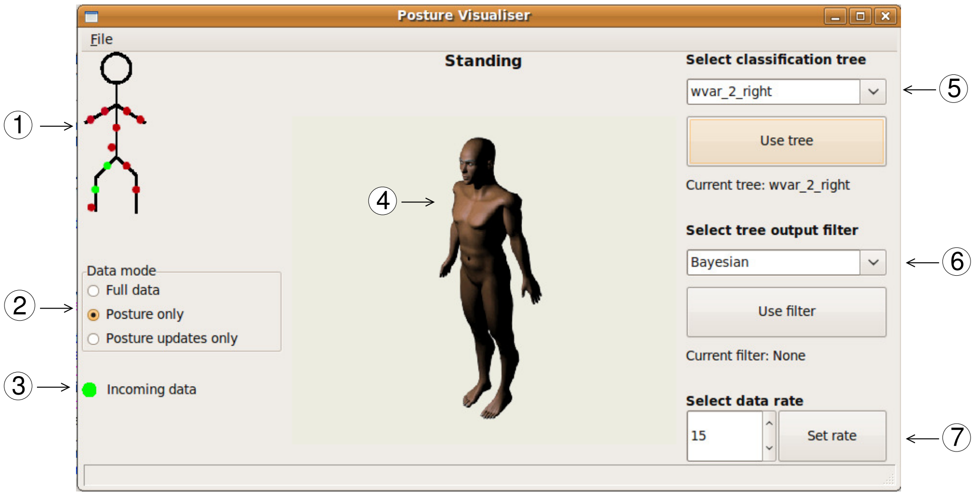

Figure 3.9: Interactive visualisation and configuration software running at the base station.

1. display of the active sensors for the specific classification tree selected (a stick man with coloured sensor markers—green for active, red for inactive),

2. configuration of the data transmission mode (as detailed in Section 3.2),

3. indication of whether data is currently being received (green when data is received, red if a defined period—one second as implemented here—has passed with no received data),

4. display of the current posture of the subject using a 3D human graphic,

5. configuration of the posture classification tree that the Primary Node should use,

6. configuration of the transition smoothing filter to apply to the classifier output, and,

7. configuration of the sampling rate in Hz to be used by the sensors.

3.5 Summary

The design and architecture of an end-to-end platform enabling on-body posture classification was presented. The platform was devised to meet the requirements of two usage scenarios described in Section 3.2: 1) self-contained system and 2) investigative instrument scenarios. These scenarios impose a number of requirements that are generic to all implementations of a Class-act system, such as on-body classification and use of battery power. A data flow was devised for each usage scenario to demonstrate the inherent similarities in system design. The specific hardware used is not specified by the platform design, allowing flexibility in specific implementations (in the use of less-wired or completely wireless communication, the number of sensors per node, and so on).

The needs of two example applications that are suitable candidates for Class-act systems were described: EOD operative monitoring and investigative instrument. These application examples introduce additional requirements beyond those generically specified for the platform.

A prototype system was implemented based on the platform design. This prototype system was used for online evaluation of the algorithms described in this thesis (see Sections 4.12 on page 76 and 5.5 on page 94). Additional detail on the prototype hardware system is given in Appendix A.

The Class-act platform and example instrument implementation described form the basis of the contribution in this chapter and Chapter 4—a wearable real-time instrument performing on-body classification of posture. To the author's knowledge, as described in Section 2.3 on page 25, no such system has previously been demonstrated in the literature.

Chapter 4

Posture classification algorithm and data gathering process

The previous chapter presented the design of an end-to-end on-body posture classification platform to address two usage scenarios (self-contained and investigative system), along with two application examples and an example implementation.

This chapter continues the work described in Chapter 3 through: 1) a posture classification algorithm suitable for deployment in a wearable (resource constrained) system such as one implemented based on the Class-act platform, 2) a method for gathering empirical data for training and testing of posture classification algorithms, 3) a demonstration of the fitness for purpose of the algorithm (used in conjunction with a suitable data feature) for the EOD application considered in this work and the wider class of related applications generally, and 4) an evaluation of the algorithm's accuracy in classifying the defined set of eight postures considered in this work. Section 2.4.4 showed that one of the main limitations of the current literature is the lack of a clear and detailed description of the design and evaluation methodology used when reporting on posture classification systems. This chapter addresses this gap with regard to the design and evaluation methodology used in this work.

The chapter is structured as follows: Section 4.1 provides an overview of the classifier testing and evaluation process. Section 4.2 describes the classification algorithm and classifier training algorithm adopted by the author. Sections 4.3 to Section 4.10 describe the experimental regimes, data gathering tools, experimental subjects, data annotation method, classifier training method, and testing method used in the work here. Section 4.11 presents findings related to the suitability of the chosen classification algorithm for the work here, specifically with regard to classifying the set of required postures and the effect of wearing an EOD suit on classification accuracy. Section 4.12 provides a real-time, real-life functional evaluation of the classifier. Finally, Section 4.13 summarises the work in this chapter.

4.1 Classifier testing and evaluation process

Figure 4.1 on the next page shows the relationship between the data gathering process and classifier testing and evaluation, as reported on in this chapter. The Class-act platform requirement of on-body classification (on a potentially resource constrained node) guided the selection of an appropriate algorithm for classifying the target postures (Section 4.2). The EOD case study application guided the design of the regimes (Section 4.4) and selection of the subjects (Section 4.6) for the experimental trials.

The next section describes and justifies the algorithm chosen for posture classification. It also introduces two central aspects of the design of the classifier and the data gathering process that affect the suitability of the classifier for the application.

4.2 Classification algorithm—C4.5 decision trees

Decision trees are a natural choice for the work in this thesis given the requirement for real-time classification on an embedded system. Once trained, decision trees are computationally simple and thus easy to accommodate on constrained embedded platforms (as demonstrated by Maurer *et al.* [87]). A further advantage of decision trees is that there are no loops within the tree. Thus, the time taken to perform a classification has a natural limit based on the depth of the tree. This aids in real-time classification of posture as the time required for classification has an upper limit that can be determined, and particularly the lack of loops means that there can never be a situation wherein a classification attempt does not complete. Prior examples of the use of decision trees in posture classification applications were discussed in Section 2.4.2 on page 31.

Of several other algorithms considered by the author, Hidden Markov Model based classifiers are used often in existing works [22, 39, 50, 56, 59, 90, 95, 114, 103]. These are generally relatively computationally complex, however, and the literature does not indicate a clear benefit in terms of classification accuracy [56, 95, 113].

The algorithm chosen for classifying the set of eight postures considered in this work was the C4.5 algorithm [30]. Generically, the C4.5 algorithm creates a decision tree by finding, at each node, the attribute (and threshold for that attribute) that allows the data samples to be most effectively divided into subsets containing particular classes. The effectiveness of a given attribute in achieving this is determined via the difference in entropy (or "information gain") resulting from choosing one attribute instead of another. Quinlan [102] provides an in-depth description of the method of creating decision trees via the C4.5 algorithm. The process of selecting attributes at each node based on the information

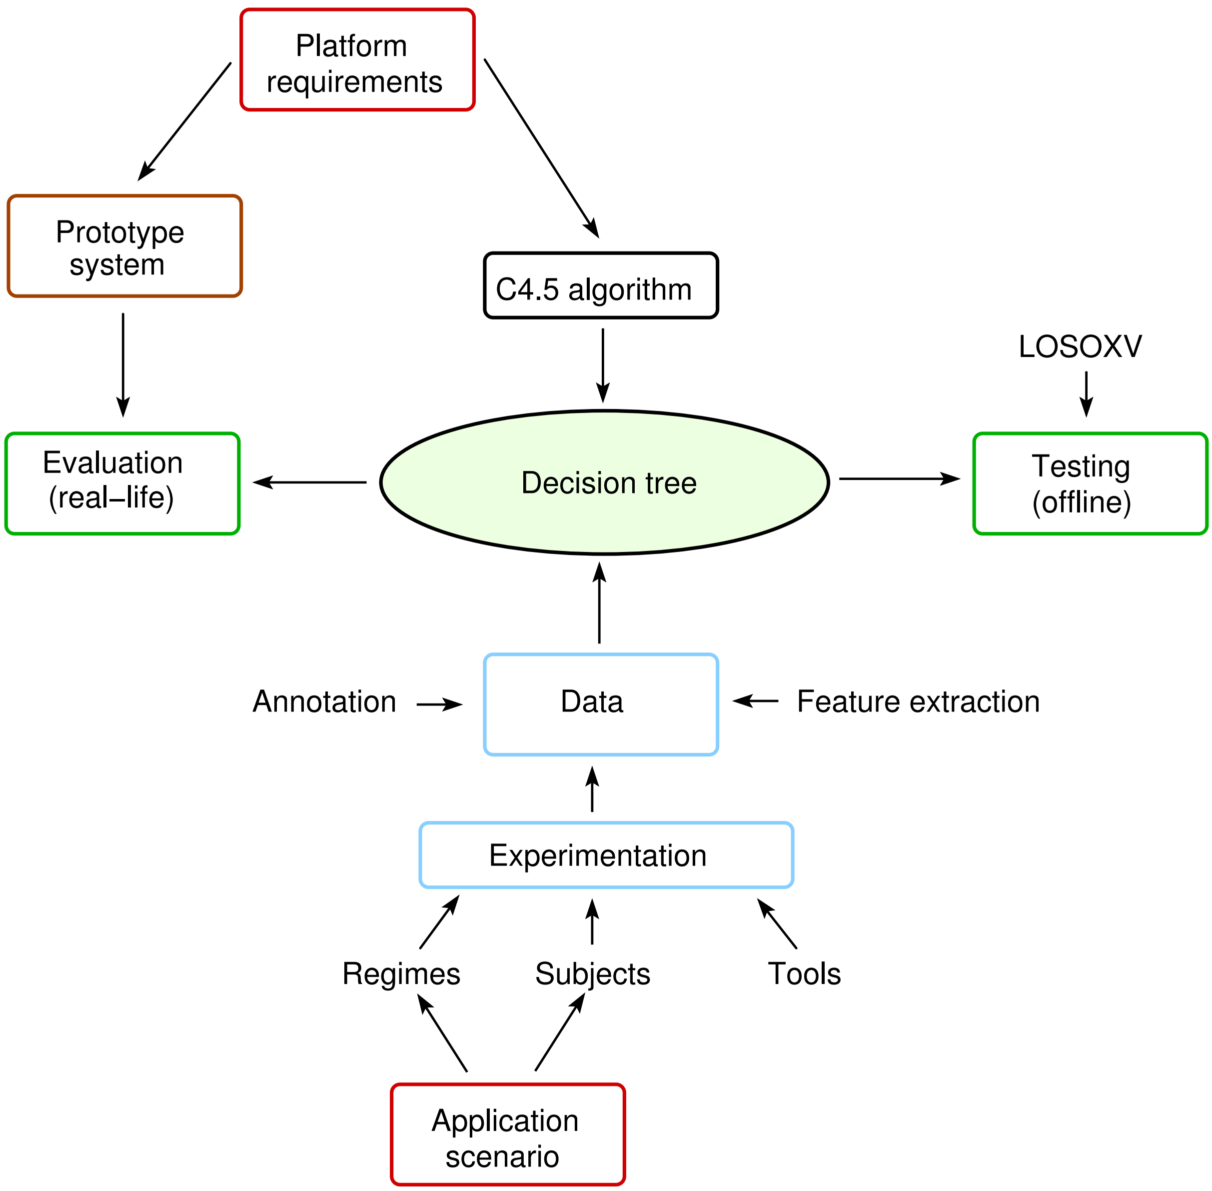

Figure 4.1: Process for data gathering and posture classifier testing and evaluation. *The links represent the way in which the different aspects of the work support each other.*

gain means that important attributes (that is, sensor locations) appear closer to the root of the tree, and redundant sensor locations are likely to be excluded from the tree. This can aid in system development by highlighting sensors that provide useful data towards classifying the required postures.

The Waikato Environment for Knowledge Analysis (WEKA) toolkit [120] was used to generate C4.5 decision tree classifiers (via the Java implementation of the C4.5 algorithm as used in WEKA, named J48). Weka is a free, easy to use, cross-platform tool implemented in Java, providing a comprehensive range of machine learning algorithms. Testing was performed using the LOSOXV method in almost all cases (as described in Section 4.3). The testing procedure was implemented outside of WEKA. LOSOXV was selected in preference to 10-fold cross-validation since the aim was to determine the expected classification accuracy of the classifier on unseen subjects. There is considerable variation in how different human subjects move. Furthermore, there will be slight variations in how sensors are fitted from one subject to the next. These two factors mean that LOSOXV forms a more stringent test of a posture classifier than ordinary 10-fold cross-validation.

In the process of developing the classifier, two main questions required investigation with regard to the suitability of the classifier for the work here:

1. Is it possible to classify the full set of required postures with a high accuracy?

2. Is it possible to train on subjects wearing light clothing and still provide a high classification accuracy when deploying the Class-act system on a subject wearing heavy protective clothing such as an EOD suit?

Classification of the full set of postures required by real-world applications such as EOD operative monitoring requires the ability to classify not only static postures (such as sitting or standing) but also dynamic postures (in this case walking and crawling). Information in addition to raw acceleration data is required for classifying dynamic postures in order to capture some of the history of the subject's movements. The solution applied here is the use of extracted data features to capture the history of the data and form part of the posture classifier's input set. Section 4.11.1 on page 74 demonstrates that use of features when classifying static and dynamic postures results in similar accuracy to using only raw data when classifying only static postures. The effects of eight different data features on classification accuracy were investigated in this work as described in Section (6.4).

The ability to train the classifier based on subjects wearing light clothing is important in the case study application since, for cost and convenience reasons, it would be preferable to train the classifier without needing to obtain application-specific clothing such as an EOD suit. Section 4.11.2 on page 75

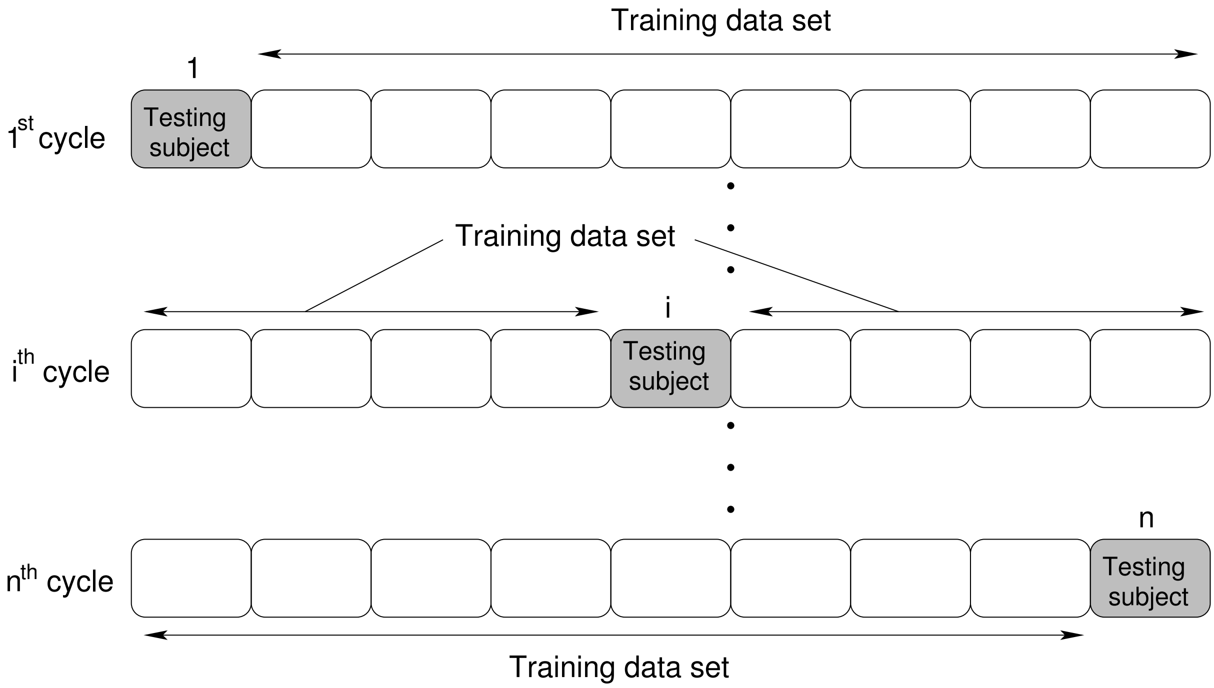

Figure 4.2: LOSOXV training and testing data selection process.

provides the results of this analysis, showing that classification accuracy is not hindered by the type of clothing the subjects wear.

The following sections provide a description of the data gathering process used in this work. This is composed of the following phases: experimental planning, data collection, pre-processing, training, testing and evaluation phases. The resulting data sets were used in the analysis in this chapter, as well as in Chapters 5 and 6.

4.3 Data gathering and classifier evaluation process overview

When developing machine learning based classification algorithms, a generic process is usually followed consisting of data gathering, training, testing, and evaluation. Details of this process, however, are lacking in the literature, as discussed in Section 2.4.4. A description of this process is given starting in this section and following on to Section 4.10, focusing on the types of testing and evaluation used in the work here. The steps are as follows: experimental planning, data collection, pre-processing, training, testing and evaluation.

1. Prior to experimental data gathering:

    (a) An experimental protocol is designed that will allow the experimenter to gather data representative of that which the classifier is expected to encounter during real-life system use.

   (b) Data gathering instrumentation is selected from the available instruments based on its ability to supply the required data for classifier training. For example, the system should support use of the number and type of sensors that will be used in system deployments using the trained classifier.

   (c) A group of subjects are selected to provide adequate coverage of the range of body types, ages, and so on that is expected for monitored subjects within the application.

2. Data is gathered via experimentation and accurately annotated (manually or automatically), providing the means to apply supervised learning techniques and to evaluate the accuracy of the classification algorithm.

3. Following the data gathering experimentation, the data is pre-processed into a form that is suitable for training decision trees (as described in Section 4.9).

4. After pre-processing the data, classifier training, testing and evaluation are performed via one of the following methods:

   (a) LOSOXV is performed in order to assess the overall accuracy of the classification algorithm based on the data gathered. For LOSOXV, the following process is applied (summarised in Figure 4.2 on the previous page).

      i. A subject is selected as the testing subject (the subject “left out”).

      ii. The data from the remaining subjects is combined to form the training data subset and a classifier is trained using this.

      iii. The classifier is tested using the data from the “left out” subject.

      iv. This process is repeated for each of the subjects in the set and the classification accuracy results from each iteration are summarised.

   (b) Real-time evaluation

      i. The best classifier obtained using the steps above is deployed as part of the Class-act system and the classification accuracy is evaluated in real-life.

The remainder of this chapter describes each of the steps given above in detail with regard to how they were applied in training, testing and evaluating the classifier. The next section describes the data gathering regimes for the work here (part of the experimental planning phase).

4.4 Data gathering regimes

Predefined *regimes* are used in the work here to ensure that the data is gathered and archived in a controlled manner. A *regime* is a description of a sequence of postures and their durations, along with the duration of the transition periods between the postures. An appropriate regime specification aids in ensuring that: 1) the data gathered is consistent between subjects with regard to instructions given to the subjects and the experimental conditions, 2) all the required postures are fully represented in the training and testing sets, and 3) correct annotation of the data is applied. Eight postures are targetted in this work. Six of the eight postures (walking, sitting, standing and lying on back, on front and on one side) are commonly targeted in the literature (as demonstrated in Section 2.4.3), while kneeling and crawling are rarely encountered in the literature but are required for applications such as monitoring of EOD operatives or firefighters. An important consideration while planning the experimentation was that data be gathered while the subjects are performing the required postures and also other activities (such as kneeling while also unpacking objects from a rucksack). This is expected to be important regardless of the specific application considered as it represents the need to train and evaluate classifiers using data gathered in realistic conditions and to prevent overfitting of machine learning based classifiers.

The regimes used here were based on existing research in the area of EOD operative safety [118] and on feedback from an EOD suit manufacturer regarding the types of activities that would be performed during EOD missions. Three increasingly complex regimes were developed by the author, progressively as the research advanced, focusing on: 1) the eight postures alone (R1), 2) the eight postures combined with natural movement (R2), and 3) mission-like activity (R3). The three regimes were as follows:

R1 Regime R1 was posture focused, requiring the subject to sit, stand, walk, kneel, crawl, lie on one side, lie on their front, and lie on their back. Each posture was maintained for one minute, with the subject performing light arm movement tasks combined with variations from the set positions (such as, for example, leaning slightly back, forth, or sideways whilst standing). The postures are exemplified in Figure 4.3 on the following page.

R2 Regime R2 was posture *and* natural movement focused, and expanded on R1 by including natural movements (such as lifting weights whilst standing, or moving objects from a rucksack whilst kneeling) as shown in Figure 4.4 on the next page. The aim with this regime was to provide the decision tree training process with data that more accurately represented movements performed by people in real-life situations (i.e. free movement).

R3 Regime R3 was mission activity focused, matching experimentation presented in existing EOD-

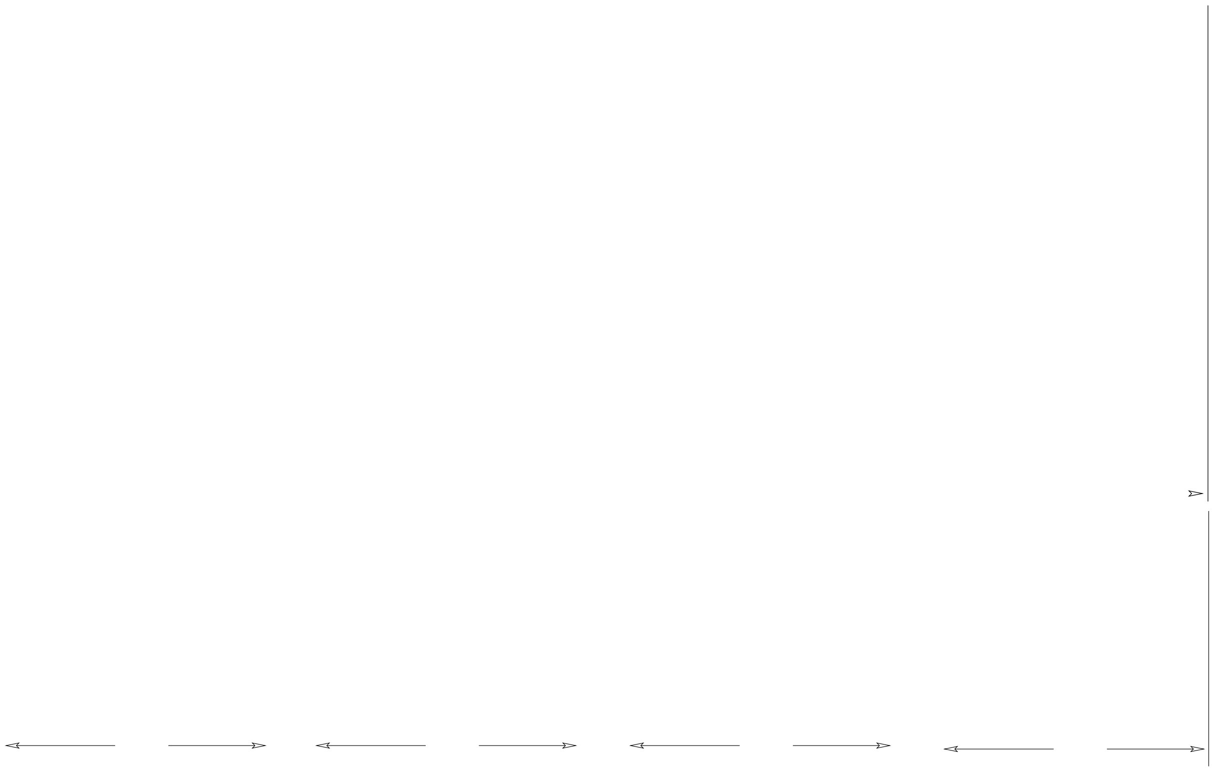

Figure 4.3: Overview of Regime 1 posture timing.

These images have been removed

Figure 4.4: Overview of Regime 2 posture timing.

These images have been removed

Figure 4.5: Overview of Regime 3 posture timing.

Figure 4.6: Sensor placement on the outside of the EOD suit.

related physiological research [117]. The aims of this regime were to reflect the activities that are most likely to be performed during EOD missions (a subset of the eight described previously) and also to reflect the expected relative durations of the activities (whereas R1 and R2 aimed to provide equal coverage of all eight postures). The activities performed were: walking (3 minutes); kneeling while moving weights into and out of a rucksack or reading (2 minutes); crawling (2 minutes); arm exercise while standing (4 minutes); sitting (3 minutes). These activities are shown in Figure 4.5.

During the experimentation, efforts were made to duplicate the environment of EOD missions and acquire data from subjects wearing the EOD suit. Variations of regimes R1 and R3 were thus performed with the subjects wearing an EOD suit. In the EOD suit trials the sensors were placed on the outside of the suit (rather than directly on the subject) as shown in Figure 4.6 on the previous page. However, analysis of the data showed that use of the EOD suit did not have a significant effect on the results obtained (Section 4.11.2 on page 75 demonstrates the effect of the suit on classification accuracy).

The investigation of transitions in Chapter 5 required variations of R1 and R2:

RT30 A combination of R1 and R2 where each posture was maintained for 30 seconds.

RT40 A combination of R1 and R2 where each posture was maintained for 40 seconds.

4.5 Data gathering tools

Data gathering was performed using three BSN systems based on two distinct hardware platforms: two systems, named DG1 and DG2 (for Data Gathering instrument), based on Gumstix Verdex devices and similar to the Class-act platform described in Chapter 3 but lacking the ability to perform on-body classification (which was not needed for data gathering purposes), and one system using an off-the-shelf hardware platform named Sensing Health with Intelligence, Modularity, Mobility and Experimental Reusability (SHIMMER) [107]. The hardware components of the DG1 and DG2 systems are described in Section A.1 on page 145 while the SHIMMER platform is described in Section A.3 on page 154. Acceleration readings were taken at a rate of 10 Hz for DG1 and DG2 and at 100 Hz for the SHIMMER-based system. DG1 and DG2 were synchronised with the base station using Network Time Protocol (NTP) [89], while the SHIMMER-based system used a different method of synchronisation described in Section A.3.

4.5.1 Sensing location configurations

During the course of experimentation, several combinations of body sensor positions were used. The SHIMMER based tool used 7 sensing positions, while DG1 and DG2 used 9 or 11 sensing positions as required by the analysis performed. These locations reflect those in the literature used by Xu *et al.* [122] and Guenterberg *et al.* [49]. The sensor locations for each system (illustrated in Figure 4.7 on the next page) were:

SHIMMER-based system The SHIMMER-based system used seven sensors due to a limitation on the number of devices that can be part of a single Bluetooth network. The sensors were placed on the ankle, lower leg, upper leg, hip, lower arm, and upper arm on the right side of the body, plus the chest.

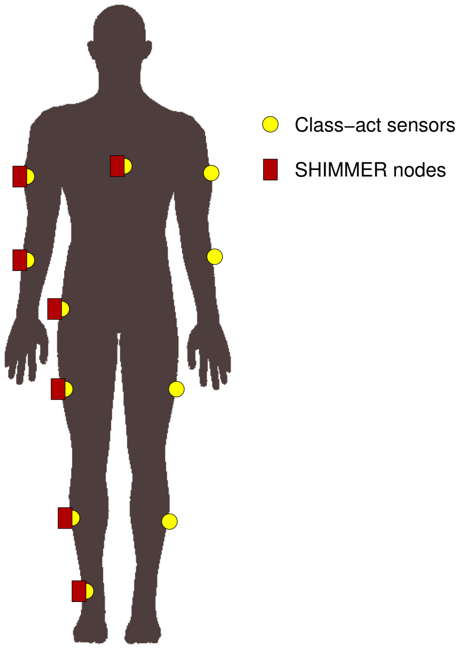

Figure 4.7: Sensor positioning for DG1/DG2 and SHIMMER systems.

DG1/DG2—9 sensors The most commonly used number of sensors was 9, selected to correspond to each distinct body segment except for the head (each segment is assumed to be rigid and thus multiple sensors at several locations on a single segment would not add additional information). The locations were: thighs, calves, upper arms, forearms, and chest.

DG1/DG2—11 sensors When analysing the classifier accuracy with differing numbers of sensor locations, two additional sensors were added to evaluate locations as used in the literature [50, 65, 95, 116]. The locations were the same as for the 9 sensor configuration, with the addition of two sensors placed on the hip and the ankle.

Based on the data gathered by these systems, Section 6.6 on page 111 investigates the number of sensors required to allow accurate posture classification, concluding that two is sufficient for the postures considered here.

4.6 Subject selection

Over the set of experimental trials, data was collected from 7 females and 15 males with a range of ages, heights and weights as shown in Table 4.1 on the next page. From the application point of view the range of experimental subjects selected were expected to provide sufficient coverage in terms of sex, age, height

Table 4.1: Experimental subject characteristics.

Subject	Age	Sex	Height (m)	Weight (kg)
S1	20	Male	1.79	76
S2	20	Male	1.82	73
S3	22	Male	1.78	62
S4	22	Male	1.80	73
S5	22	Male	1.83	89
S6	22	Male	1.84	82
S7	22	Male	1.85	70
S8	22	Male	1.87	70
S9	23	Female	1.60	49
S10	23	Female	1.63	53
S11	23	Female	1.64	60
S12	23	Male	1.70	74
S13	23	Male	1.72	72
S14	23	Male	1.75	72
S15	24	Female	1.67	62
S16	25	Male	1.80	75
S17	26	Female	1.59	54
S18	27	Male	1.86	72
S19	28	Male	1.80	77
S20	29	Female	1.68	64
S21	31	Male	1.64	56
S22	36	Female	1.60	60
Min	20	-	1.59	49
Max	36	-	1.87	89

and weight. Particularly, it was assumed that the operatives monitored by a deployed system are likely to be males between 20 and 30 years old. A large group of subjects were used as this is important during training to help reduce the possibility of over-training (associating a posture with a restricted range of readings that do not generalise well) and thus increase the accuracy of the system when classifying data from unseen subjects.

4.7 Summary of data sets gathered

All of the classifier training and testing in this work were based on experimentally gathered data. Table 4.2 on the following page summarises the experimentation performed to support this work.

4.8 Data annotation

The data gathered for classification algorithm training and testing must be annotated with the correct posture for each sample. An example of data annotation is shown in Figure 4.8 on page 71. In this example, the subject begins by sitting for 30 seconds, followed by transition to lying on their side. An example of annotated acceleration data over all eight postures for an accelerometer placed on a subject's right calf is shown in Figure 4.9 on page 71.

Over the course of the work presented in this thesis, three methods of annotation were used:

1. Manual annotation was performed based on notes taken during the experimentation. In this case the start and end times of each posture were noted by the experimenter.

2. Manual annotation was performed based on video footage filmed during the experimentation. To ensure data annotation was performed correctly, a reference point was established in the data—the subject made a particular movement at the start of the test that showed up clearly within the sensor data and on the video. In the experimentation here, the subject held both arms in the air for several seconds.

3. Automatic annotation was performed by a script that prompted the subject to change posture at defined intervals. A visual and audio countdown was provided to allow the subject to change posture at the correct time as defined by the regime.

Table 4.2: Summary of experimentation performed.

	Data acquisition system	With EOD suit?	No. of sensors	No. of subjects	Annotation method	Regimes performed					Hours of data	Subjects
						R1	R2	R3	RT30	RT40		
D1	DG1	-	9	4	Manual	✔	✔	-	-	-	1.9	S3, S5, S13, S22
D2	DG1	-	9	1	Manual	✔	✔	✔	-	-	0.5	S2
D3	DG1	-	11	5	Manual	✔	✔	✔	-	-	2.6	S6, S8, S10, S14, S19
D4	DG1	✔	9	4	Manual	✔	-	✔	-	-	1.8	S4, S7, S12, S13
D5	DG2	-	9	7	Automatic	✔	✔	-	-	-	2.5	S4, S13, S15, S16, S18, S20, S21
D6	DG2	-	9	7	Automatic	-	-	-	✔	-	4.2	S4, S12, S13, S15, S16, S20, S21
D7	DG2	-	9	5	Automatic	-	-	-	-	✔	2.9	S4, S13, S15, S20, S21
D8	DG2	-	4	6	Manual	✔	-	✔	-	-	3.0	S4, S11, S12, S16, S18, S21
D9	SHIMMER	-	7	7	Manual	✔	✔	-	-	-	2.2	S1, S4, S9, S11, S13, S15, S17
D10	SHIMMER	-	7	1	Manual	✔	-	-	-	-	0.2	S21

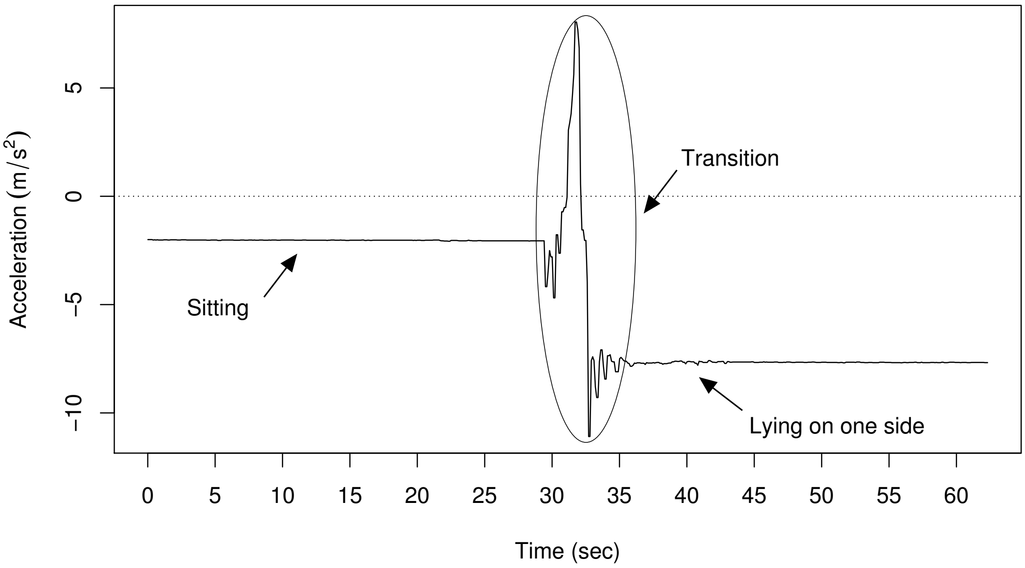

Figure 4.8: Data gathering annotation example for the x axis of the right calf sensor.

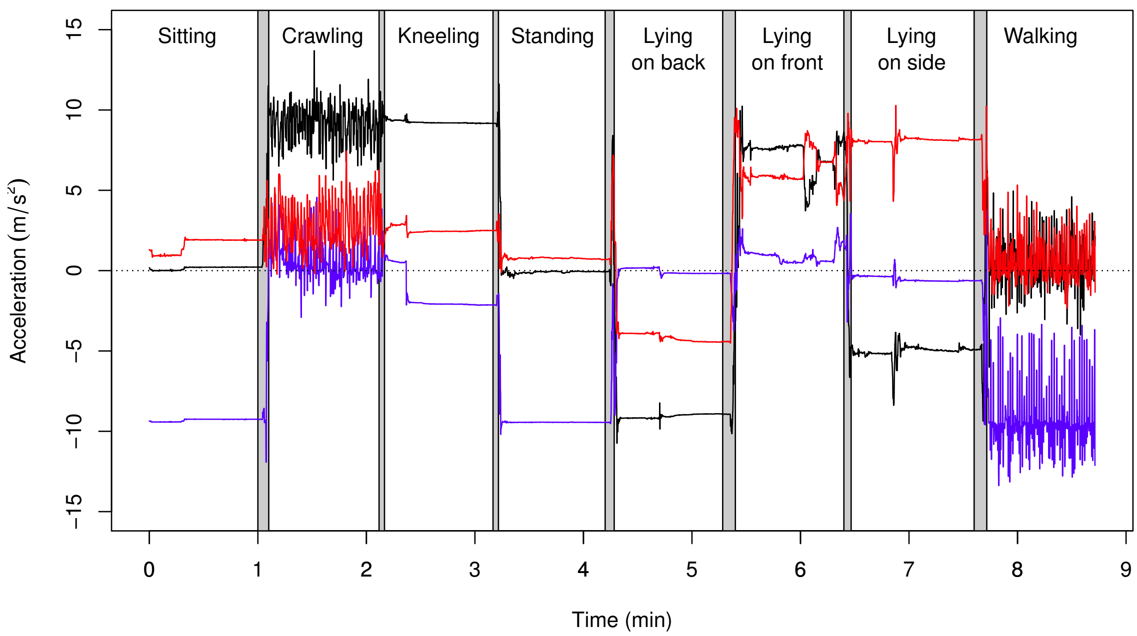

Figure 4.9: Accelerometer signal from the right calf sensor over all eight postures. *Black, red, and blue lines show the x, y, and z axis data respectively, while shaded areas indicate transitions.*

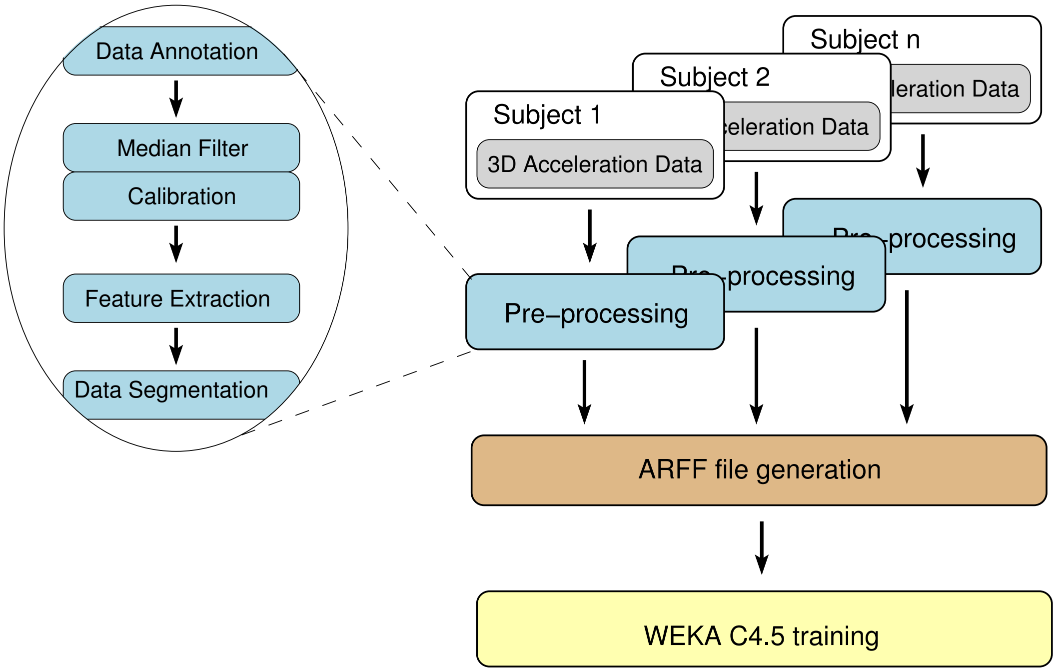

Figure 4.10: Classifier training process overview.

4.9 Data pre-processing and classifier training

Decision trees, as a machine learning based approach to data classification, require training. However, before a decision tree can be trained on the gathered data, the data must be pre-processed into the required form. Figure 4.10 shows the data pre-processing and classifier training. Starting with the raw 3D acceleration data logged during experimentation, several pre-processing stages are applied:

Data annotation Data annotation is performed via one of the methods described in Section 4.8.

Median filter The data is filtered using a median filter with a window size of three samples. This eliminates single-sample spurious readings as demonstrated in Section 3.4.4 on page 52.

Calibration Calibration coefficients are applied to the data, as described in Section A.1.3 on page 147. While the decision was made to apply calibration correction to the data, it should be noted that the accuracy of a decision tree will not generally be affected by a linear data transform applied to both the training and testing data. However, calibration allows for sensors to be swapped between body locations without causing classification errors (not applying calibration correction means that two sensors will produce different values for the same body location in the same situation).

Feature extraction A selected data feature is extracted and appended to the raw data. For the analysis here, WVar was used (calculated over a 30 sample sliding window) as it was found to provide the highest classification accuracy (more in-depth analysis is provided in Section 6.4).

Data segmentation The data is segmented as needed for the particular investigation and unneeded segments are discarded. For example, in some of the investigations, transitions are not considered (transitions are discussed in detail in Chapter 5), while in other cases only specific postures are considered.

Once the data has been processed and is suitable for decision tree training, an ARFF format file is created. This is the format required by WEKA, consisting of a header section describing the input variables and a data section which has each of the samples to be considered. This file is then supplied to WEKA, which trains the decision tree.

4.10 Decision tree testing methods

Decision tree testing is performed with experimentally gathered data via two methods: 1) LOSOXV and 2) training on one group of data sets and testing on a different group.

- LOSOXV is used to determine the effect of different parameters on accuracy (for example, selecting the data feature that provides the highest accuracy) and results in the training of a tree, and the calculation of metrics such as classification accuracy, for each iteration (each subject “left out”). The mean classification accuracy across these iterations is referred to in this work as the overall accuracy. Section 4.3 provides a description of the LOSOXV testing process.

- Training and testing on two different groups of data sets is used as an alternative to LOSOXV in specific investigations, such as training on data from subjects wearing light clothing and testing on data from subjects wearing an EOD suit (Section 4.11.2). With this method a tree is trained using all of the training data and is then applied to each of the unseen testing data sets. This results in an accuracy result for each testing data set.

4.11 Classifier suitability evaluation

As described in Section 4.2 on page 58, there were two main criteria for determining the suitability of the chosen classification algorithm for the work here: 1) the ability to accurately classify the required set of eight postures and 2) the ability to accurately classify the posture of EOD suit wearers based on

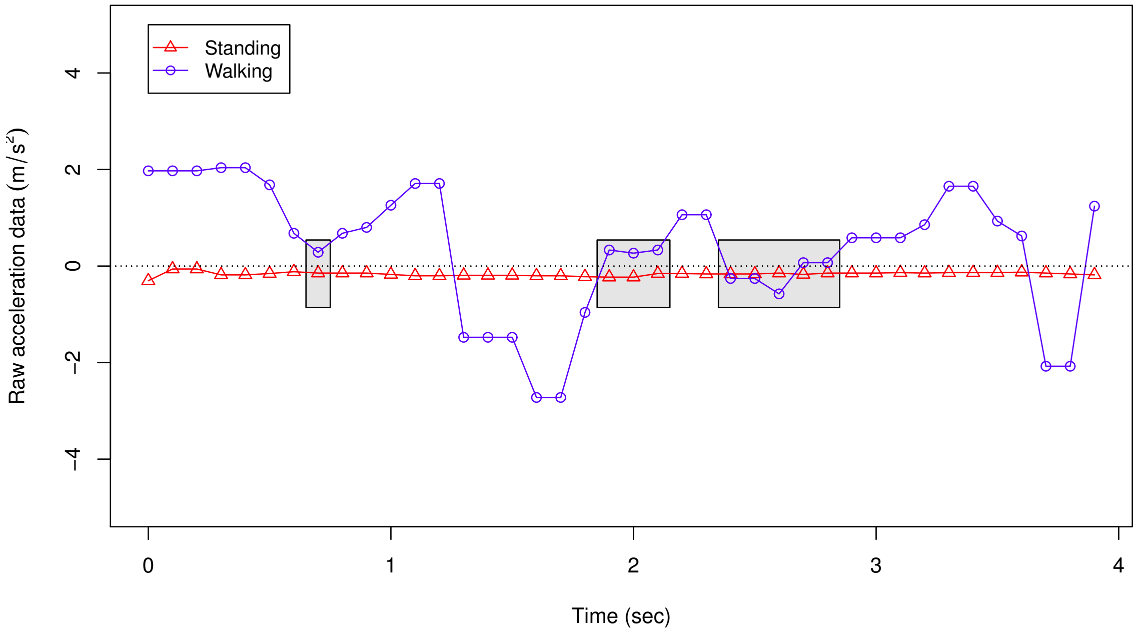

Figure 4.11: Demonstration of raw sensor data for walking and standing taken from the x-axis of the right calf sensor. *The shaded areas indicate samples in the walking data that closely match those of the standing data.*

training data from subjects wearing light clothing. This section analyses C4.5 trained decision trees to answer these questions based on the data sets described in Section 4.7 on page 69.

4.11.1 Ability to classify the full set of postures

Of the eight postures considered in this work, six are classed as static postures (sitting, standing, kneeling, and three variations of lying down) and two are classed as dynamic postures (walking and crawling).

It is expected that static postures can generally be classified using only raw acceleration data as they each involve different orientations of the body segments. Dynamic activities such as walking, however, require knowledge of the history of the data as they involve movement that cannot be evaluated based on single time step data samples. As an example of this, walking produces acceleration values that *at some time instances* cannot be distinguished from standing. During walking, the motion of the leg is similar to a pendulum and thus one can expect the acceleration to be roughly sinusoidal along the axis closest to the direction of motion. Conversely, standing produces near to zero acceleration in the forward axis. This effect is demonstrated in Figure 4.11 where it can be seen that acceleration readings generated by walking overlap at some points in time with readings generated by standing. Information about the history of the data can be provided by extracting data features (such as variance) that provide a significantly different output trace for static and dynamic postures. For the example case of standing and walking, calculating WVar for standing results in an output of around zero (as there is little variance in the values), while calculating it for walking results in output values that are usually between around $2\ \left(\mathrm{ms}^{-2}\right)^2$ and $3\ \left(\mathrm{ms}^{-2}\right)^2$ (and always above zero).

This section demonstrates the effect of the different posture types on classifier accuracy and the benefit of using an extracted data feature. In doing this, it shows that the use of an appropriate feature allows the classifier to classify the set of eight postures considered in the work here. A LOSOXV-based evaluation was performed over data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects).

First, the data was truncated to include only the static postures: sitting, standing, kneeling, and the three variations of lying down. The overall classification accuracy obtained was 98.0% when using only sensor data as input to the classifier. This accuracy is considered here to be sufficiently high to confirm that static postures can be correctly classified using only raw data. Next, the same test was performed with all of the postures included in the data sets, giving a reduced accuracy of 90.5%. Finally, a test was performed using all postures and with WVar as an extracted data feature added to the sensor data inputs (the analysis given in Section 6.4 shows that the use of WVar provides the best classification accuracy overall). This resulted in a classification accuracy of 96.3%. While not as high as the accuracy for classifying only the static postures, it is sufficiently high to conclude that all eight postures can be accurately classified if an extracted data feature is used.

4.11.2 Simplifying data gathering for the EOD scenario

When devising a method for gathering data to train a classifier, it is helpful to minimise the complexity of the exercises and the equipment required. Particularly, for the application case study here, it is desirable to require minimal use of an EOD suit. Wearing such a suit is very strenuous for the subject and in the context of this research such use required consent from volunteers and additional supervision. Furthermore, depending on the resources available to the experimenter, access to EOD suits is not necessarily guaranteed.

Given the issues described, the ideal scenario is to collect training data from subjects wearing their normal clothing. Therefore, it is important to understand the impact of clothing type on the accuracy of the classifier when deployed on EOD suit wearers. A significant reduction in accuracy would render the above method unsuitable for system training. The analysis in this section thus tests the suitability of gathering data from subjects wearing normal clothing to train a classifier for deployment in EOD missions.

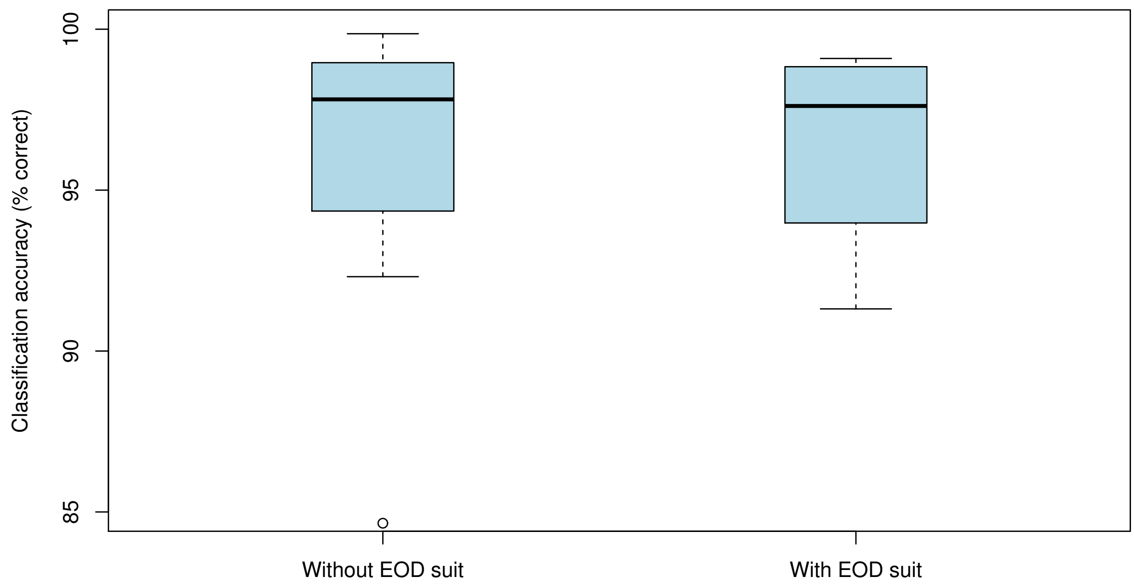

Figure 4.12: Impact of wearing an EOD suit on classification accuracy.

A decision tree was trained using data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects) with the subjects wearing light clothing (for example, a shirt and jeans). This is the same group of data sets as used for the analysis in Section 4.11.1 on page 74. For testing purposes, eight sets of data (from data set D4) were used. These were gathered from four subjects with nine sensors placed over an EOD suit. The analysis here is based on the data gathered by two of those sensors (left thigh and calf) as Section 6.6 establishes that two sensors on the thigh and calf are sufficient for accurate classification.

Figure 4.12 shows the results of this analysis. The testing results were compared with the results given in Section 4.11.1 for subjects not wearing an EOD suit and it was found that testing on subjects wearing an EOD caused little effect on the classification accuracy—the difference in the mean accuracy between the two cases was 0.1%. These results indicate that a classifier trained on subjects wearing light clothing will be suitable for deployment on subjects wearing an EOD suit.

4.12 Real-life functional evaluation

This section presents an evaluation of the Class-act system implementation (described in Section 3.4 on page 49) performing real-time on-body classification. The system was deployed on five subjects performing tasked activities with classification being performed in real-time on the Primary Node. The five subjects were two females and three males with a range of ages, heights and weights as shown in Table 4.3.

These images have been removed

Figure 4.13: Tasked activities for real-life real-time prototype system evaluation. *From top-left: crawling, sitting, kneeling, walking, standing, lying on front, lying on back, lying on side.*

Table 4.3: Experimental subject characteristics for real life trial.

Subject	Age	Sex	Height (m)	Weight (kg)
S1	27	Female	1.63	67
S2	26	Male	1.72	79
S3	44	Female	1.54	56
S4	25	Male	1.8	73
S5	25	Male	1.7	74
Min	25	-	1.54	56
Max	44	-	1.72	79

The subjects performed a series of activities as shown in Figure 4.13 on the preceding page, based on regime 2 as described in Section 4.4. The specific activities performed were: crawling under a table, sitting, kneeling while moving items out of and back into a box, standing and drawing on a whiteboard, lying on their front and using a laptop, lying on their side and moving items out of a box, lying on their back and writing on a piece of paper above them. The subjects walked between each activity station.

The Primary Node was loaded with a decision tree trained using data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects). The parameters for training were: 9 sensor locations, WVar as the extracted feature, and a window size of 30 samples (details of the optimal parameters are given in Chapter 6). A video camera was used to record the experimentation and the manual video based annotation method was used as described in Section 4.8. The total experiment time over the five subjects was 1.2 hours. Though Section 6.6 on page 111 shows that two sensors are sufficient for accurate classification, this test was performed using nine sensors (the full sensor load for the Class-Act system implementation used) in order to capture the highest accuracy configuration available. As classification was performed on-node and only the final postural information was reported, it is not possible to give the results for these tests based on the use of only two sensors. However, as demonstrated in Section 6.6, the classification accuracy for two sensors and nine sensors was similar (95.5% average for the two versus 96.3% for all nine).

The system was evaluated in terms of the classification accuracy and the information yield. The classification accuracy for each test was calculated by comparing the actual posture performed by the subject (using the annotations) with the classified posture determined by the Primary Node. This was calculated over truncated data (with transition periods removed). Yield was calculated by comparing the number of classified postures over the course of the test with the expected number based on the test duration. The primary aim of this experimentation was to show that the classification accuracy during deployment in a realistic environment with a subject performing tasked activities was consistent with the

Table 4.4: Summary of real-time evaluation results for five subjects.

	S1	S2	S3	S4	S5
Test duration (minutes)	13.6	15.7	14.3	17.7	9.6
Classification accuracy (%)	98.5	98.1	94.2	97.1	97.3
Accuracy reduction w/transitions (%)	2.6	4.3	3.9	3.8	4.3
Information yield (%)	99.8	100	99.2	99	99.9

accuracy found during offline testing and that the yield was acceptable.

Offline testing using LOSOXV over data sets D1, D2, D3, and D5 with the parameters used here provided an accuracy of 96.3%. Based on the results presented in Table 4.4, it can be seen that the classification accuracy was within the range found during offline testing, meaning that the training process used is also applicable outside the laboratory environment. When transitions are considered (with both the initial and final posture counted as a correct classification during a transition) the accuracies are reduced by between 2.6% and 4.3%. Chapter 5 discusses transitions in depth and described the method used to reduce their impact on the classification accuracy and stability of the output of the system. The information yield for the system in the "information" transmission mode was above 99%, which is an acceptable yield for the applications considered in this work. The literature implies that a 99% yield is generally acceptable [38].

4.13 Summary

Given a requirement for real-time on-body classification on a resource constrained system, a natural choice for the posture classification algorithm was decision trees. Decision trees are computationally light and suitable for deployment on constrained platforms such as the ones that would be required by Class-act. The algorithm chosen for training of the decision trees was the C4.5 algorithm. The WEKA toolkit was used to generate C4.5 decision tree classifiers. A complete method for data gathering and training and testing of classifiers were described, including annotation and data preprocessing.

The methods of gathering data for tree training, testing and various analyses performed in this thesis were standardised around a series of regimes based on existing research in the area of EOD operative safety and on feedback from an EOD suit manufacturer. These regimes focused on: 1) the eight defined postures alone, 2) the eight postures combined with natural movement, and 3) mission-like activity. Data was collected from 7 females and 15 males with a range of ages, heights and weights. From the application point of view the range of experimental subjects selected were expected to provide sufficient coverage in terms of sex, age, height and weight. A total of 22.2 hours of data was collected for use in training and testing the decision trees. Decision tree testing was performed with experimentally gathered data via two methods: LOSOXV and direct comparison of a trained tree's output against the correct posture when tested on unseen data.

The suitability of the classifier required investigation in two main regards. Classification of the full set of postures considered in the work here requires the ability to classify not only static postures (such as sitting or standing) but also dynamic postures (specifically walking and crawling). Furthermore, these postures must be classified while the subject is performing other activities during the posture (such as unpacking items form a rucksack while kneeling). The ability to train the classifier based on subjects wearing light clothing is important for the EOD application since, for cost and convenience reasons, it would be preferable to train the classifier without needing to obtain an EOD suit. The testing results showed that all eight postures can be accurately classified if an extracted data feature such as WVar is used (classification accuracy of 96.3%), and that a classifier trained on a subject wearing light clothing will produce accurate classification when deployed on a subject wearing an EOD suit.

The end-to-end Class-act system implementation was also evaluated when deployed and classifying posture in real-life and real-time for five subjects performing tasked activities. The classification accuracy was shown to be consistent with the results found during offline testing , and the information yield was sufficient for the applications considered here (above 99%).

Analysing the results of the real-life evaluation showed that transitions reduced the overall classification accuracy by between 2.6% and 4.3%. The next chapter investigates transitions and describes a method of 1) reducing their impact on classification accuracy and 2) maintaining a more stable posture output.

Chapter 5

Transition smoothing filters

Natural human movement and activities are composed of well-defined *postures* along with *transitions* between those postures. In relation to transitions, the goals for a system such as the one developed in this thesis are 1) to provide high classification accuracy of natural human movement in the field (i.e. for sequences of postures with transitions between), and 2) minimise the number of transmissions required by an event-based system (transmitting only posture updates).

It has been shown in the previous chapter and elsewhere in the literature that well-defined postures may be accurately classified (96.3% accuracy for truncated data was demonstrated in Chapter 4). However, the classification of posture *and* transitions poses two additional problems (the first of which was demonstrated in the previous chapter) in the context of the goals described above :

1. Classifiers trained on truncated, well-defined posture data will see a degradation in accuracy due to transitions when deployed in realistic scenarios. The results here show that the classification accuracy of tree-based classifiers, when evaluated on untruncated data, degrades at a rate of 2% for each transition/minute encountered in the evaluation protocol.

2. Rapid changes in the classifier output will occur during transitions (also observed by Parkka *et al.* [97]). The primary benefits of an event-based system are a lower bandwidth requirement and reduced power consumption, both resulting from reduced transmissions. Rapid posture changes will prompt a higher number of transmissions and thus counteract these benefits.

This chapter presents a method of reducing the negative impact of transitions with regard to the effects described above. Three transition smoothing filters were designed, implemented and evaluated, corresponding to the Post-process stage of the processing chain described in Section 3.2 on page 44. They act on the classifier output and limit the accuracy loss to a predictable rate of 1% per transition/minute when deploying the classifiers in realistic scenarios. When applied to an event-driven postural classification system, the filters further reduce the number of events transmitted in realistic scenarios by 75%. When compared to a continuous posture-reporting system, the number of posture updates transmitted is reduced by 99.6% (a 270-fold reduction).

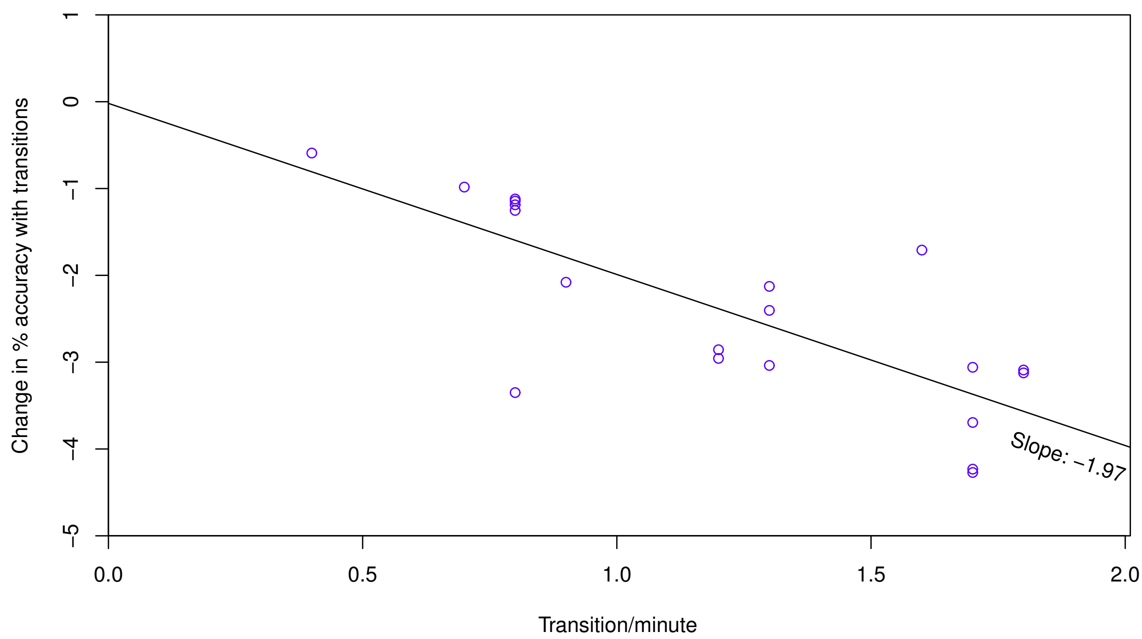

Figure 5.1: Effect of transition frequency on classification accuracy.

The contribution to knowledge in this chapter is brought by an investigation into posture transitions. Three posture filters that remove such transient posture changes are designed, implemented and tested on experimental data. The best performing filter, Exponentially Weighted Voting (EWV), is shown to reduce posture change events by 75.2% and increase accuracy by 1% (over unfiltered results). The reduction in change events is in addition to the transmission reduction for switching from a continuous reporting system to an event-based system (98.5%, a 66-fold reduction).

The structure of the chapter is as follows: Section 5.1 describes the issues related to handling transitions in a posture classification system and the goals of the methods developed by the author. Section 5.2 proposes three transition smoothing filters. Section 5.3 presents the implementation of the filters within the prototype system and describes the testing method used to select the best filter, followed in Section 5.4 by the test results. Section 5.5 presents a realistic evaluation of the benefits of the selected transition smoothing filter when implemented into the real-time, event-based prototype classification system. Finally Section 5.6 gives a summary of the work in this chapter.

5.1 Handling transitions in posture monitoring systems

In the literature, a common approach to training, testing and evaluation when using supervised learning is to truncate the datasets and consider only periods of stable posture or activity [14, 46, 53, 65, 123].

The focus for the majority of work is on classifying a set of well-defined, stable postures; the effect of transitions on classification accuracy in a realistic deployment is not considered or evaluated. Reported systems are therefore likely to deliver poorer classification accuracy when deployed in realistic scenarios, particularly if the regimes monitored present a high number of transitions per minute. Figure 5.1 on the facing page demonstrates the effect on classification accuracy of the number of transitions per minute for a tree trained on truncated data (following the methodology presented in Chapter 4). It can be seen that classification accuracy decreases as the frequency of transitions increases. The accuracy of the classifier can thus only be predicted using knowledge of the expected frequency of transitions.

To avoid such performance degradation and variability in realistic scenarios, one solution could be to provide both well-defined postures *and* transition data to the machine learning algorithm. Each posture would be annotated correspondingly while the change between one posture and another would be annotated as a separate "transition" class. However, this approach is flawed as transitions are themselves composed of a variety of short-lived postures, many of which are from the very set of postures which need to be classified.

In essence, not only is the learning algorithm being instructed that data samples in particular ranges will sometimes belong to a specific posture class and other times will belong to the "transition" class, but also that almost any sample could belong to the "transition" class. Depending on the algorithm in use this can have different effects, including ignoring the less commonly occurring class or constructing a classifier that uses very small (such as single measurement unit) differences in the readings to make the decision. In the case of the decision trees used in this work, a classification accuracy of 62.4% was found when testing this method of annotation using data sets D5, D6, and D7 (9.6 hours of data from 8 subjects). This is clearly unacceptable when compared to the state of the art.

Some researchers have looked at the classification of specific transitions related to their application, for example Godfrey *et al.* [48] looked into sitting-to-standing and standing-to-sitting transitions in order to aid in classifying standing and sitting using a single sensor placed on the chest. Similarly, Li *et al.* [76] and Jafare *et al.* [60] considered transitions such as standing to lying, which are important in classifying falls. Methods considered in relation to classification of transitions include segmentation of the signal using thresholds [60] and discrete wavelet transform [1, 48]. However, it is desirable to handle transitions in a way that is generally applicable, whereas these types of method require, for example, knowledge of specific types of transitions that may occur.

In applications such as the EOD operative monitoring case study, it is more important to provide a consistent, stable view of the operative's posture (that is, minimise the posture fluctuations during

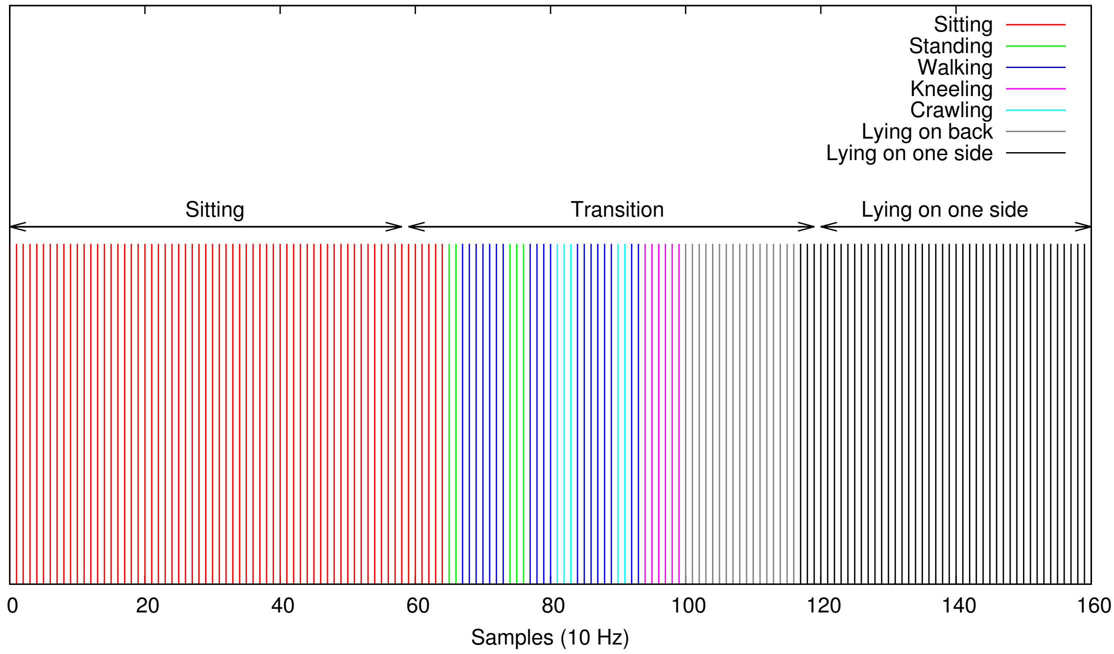

Figure 5.2: Classification output sequence when a subject is transitioning from sitting to lying on one side. *Labels indicate the correct annotations.*

transitions) than to specifically determine that they are in any given type of transition. Fluctuation in the classification output will cause fluctuations in both the automated processing/modelling performed by an external subsystem and also the visualisation shown to an observer. This requirement can thus be seen to arise generally where postural information is to be visualised or used as input to a control system. The need for stability means that the previously described issues related to detecting transitions in real-time between the postures can be side-stepped entirely and a different approach used with the goal of providing a posture output that does not fluctuate during transitions. The desired output from the system here during transition periods is therefore either one of 1) the posture immediately prior to the transition (the initial posture) or 2) the posture immediately following the transition (the final posture).

Figure 5.2 shows a classified sequence of data where the subject transitioned from sitting to lying on one side. During this transition the classified posture changed a total of 11 times. An unstable visual output of this type leads to more transmissions from an event-based system, as well as presenting an external modelling subsystem or visualiser with several different posture inputs in a short space of time. In this case, the ideal output would be for the classification to switch from sitting and lying on one side at some point (and only once) during the transition.

The justification for the work proposed in this chapter is therefore as follows:

1. Over a period of time around a transition, the output from a classifier trained on truncated data will consist of: 1) two stable postures, pre and post transition, and 2) short-lived postures from the defined set during the transition. This work proposes the use of a *transition smoothing filter* that outputs a posture estimate on the basis that the postures of interest will be of relatively long duration and that fast changes are not a desirable output. Ideally, this would result in the output during the transition consisting of only the postures observed before and after the transition occurred. This filter would be applied to all of the classifier output resulting also in the removal of spurious incorrect classifications during stable postures.

2. In many cases it is desirable to implement an event-based system instead of a continuous monitoring one, particularly when: 1) the system's battery life is impacted by the number of transmissions made or 2) the application implies long periods where the classifier output values will be constant. In the former case a continuously reporting system will quickly drain the batteries, while in the latter almost all of the postural information transmitted will be redundant. In the case of a battery-powered on-body posture monitoring system, both of these conditions are true. Thus, an event-based system would be preferable, where only updates to the subject's posture are transmitted instead of every classified posture. The method of handling transitions (such as the filter described here) influences the number of posture change events that will be generated by such an event-based system.

In summary, the approach proposed here is to define the initial and final postures around a transition as being "correct" for the duration of the transition, and apply a filter to the classifier output that will satisfy this criteria and achieve a stable view of posture (maintaining high classification accuracy and reducing the number of the posture changes that are detected by the system). Three different filter candidates were designed, deployed and evaluated using experimentally gathered data. The following sections describe them, the experimental data used, and the results of the evaluation.

5.2 Candidate transition smoothing filters

Three options were considered for the transition smoothing filters: a voting filter, a weighted voting filter, and a Bayes filter. Conceptually, these filters take a time-series of classified postures as input and attempt to "smooth" them (that is, produce a more stable series of output postures) based on the assumption that posture tends to be static over time.

Voting filter

The voting filter uses a sliding window where the last n classification results are summarised to find the posture that appears most frequently. This is based on the assumption that stable postures are likely to appear more frequently in the recent classifications than transient/spurious postures (such as observed during transitions). Given a set of past unfiltered posture estimates $d(t), d(t-1), ..., d(t-n+1)$ the posture chosen c^* at time t is given by,

$$c^*_{voting}(t) = \arg\max_{c \in C} \sum_{i=0}^{n-1} [c = d(t-i)]$$

where the term in square brackets yields 1 if true and 0 otherwise (following Iverson's bracket notation) and the set C denotes the possible postures.

Exponentially weighted voting

EWV is inspired by the Exponentially Weighted Moving Average (EWMA) filter. Although the Voting filter is simple and robust, it makes the assumption that all votes are equal. The EWV filter improves on this by attributing greater weight to recent posture estimate inputs. This is based on the assumption that more recent posture estimates are likely to be a better indicator of actual posture than less recent ones. This also means that the filter will be faster to respond to an actual posture change than the Voting filter would. In operation, each posture class c is associated with a separate EWMA-based filter w_c. Given the current unfiltered posture estimate $d(t)$ and the prior filter output $w_c(t-1)$, the new output for each posture is calculated as,

$$w_c(t) = w_c(t-1) + \alpha([c = d(t)] - w_c(t-1))$$

for all $c \in C$. A constant α controls the relative weight of newer values over old. Once $w_c(t)$ has been calculated for each c, the posture with the largest filter output is chosen.

Bayes filter

A Bayes filter is a general algorithm for filtering on the basis of a Dynamic Bayesian Network model [119]. The Bayesian net model for this filter is shown in Figure 5.3 on the facing page and consists of a time-based dynamic net where the postural state x evolves over time and also affects sensor readings z.

The model contains two causal links: First, the posture x causes accelerometer sensor readings z. Second, posture x_{t-1} at time $t-1$ influences the posture x_t at time t. In principle, the intentions of

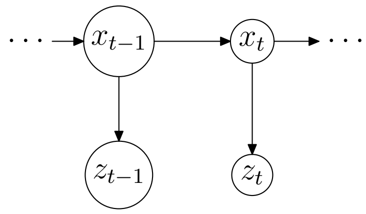

Figure 5.3: Dynamic Bayesian Network for postural state x and corresponding sensor reading z.

the wearer form a "control" causal link, however it is assumed that this is unobservable and thus is not included in the model. There may be some point to modelling intention since intermediary postures are gone through when going, say, from kneeling to walking. Therefore, a uniform set of intentions yields a non-uniform distribution between subsequent postures. It is not clear, though, what the distribution of intentions might be.

In the approach here, a further link exists between the sensor values and the unfiltered estimated posture. We collapse the two-stage link between actual posture and estimated posture into a single causal link. The estimated posture at time t is thus denoted z_t in this description. This necessarily ignores some information that would be available by considering individual accelerometer readings. The key difference between a Bayes filter approach and the HMM approaches used elsewhere [22, 55] is that in the Bayes filter, the state (which is hidden in an HMM) corresponds to a known attribute, such as the wearer's posture. In our approach, we start with an existing decision tree-based classifier that infers posture from acceleration sensor readings and that has known classification accuracy.

The filter requires us to identify the set of conditional probabilities associated with changing or keeping posture $P(x_t|x_{t-1})$ and those associated with the sensor identifying a posture, given an actual posture $P(z_t|x_t)$. These are referred to here as the transition model and sensor model, respectively. One way to obtain these conditional probabilities is to derive them from experience. In this case, it is important that the environment and behaviour of the subject is as natural as possible. Also, extensive trials are required to produce a good estimate of the true conditional probability distributions. An alternative approach is to use existing knowledge to estimate the transition and sensor model distributions. For example, it can reasonably be assumed that posture does not tend to change rapidly. Furthermore, the accuracy of the estimated posture (and thus the associated conditional probability distributions) can be derived from the precision and recall of the classifier. In this work, we fix the conditional probability of the posture

Table 5.1: Experimental data used in training the classification tree used in the transition smoothing filters evaluation. *This is a subset of data sets gathered listed in Section 4.7 on page 69.*

	No. of subjects	Regimes performed R1	R2	R3	Hours available of data
D1	4	✔	✔	-	1.2
D2	1	✔	✔	✔	0.5
D3	5	✔	✔	✔	2.6

staying the same according to,

$$P(x_t = u|x_{t-1} = v) = \begin{cases} p & \text{if } u = v \\ (1-p)\,/\,(N-1) & \text{otherwise} \end{cases}$$

for all combinations of postures u, v and where N is the number of postures. The sensor model is set according to,

$$P(z_t = u|x_t = v) = \begin{cases} q & \text{if } u = v \\ (1-q)\,/\,(N-1) & \text{otherwise} \end{cases}$$

for all combinations of postures u, v. Thus the entire set of conditional probabilities is defined by two constants p and q.

5.3 Transition smoothing filter implementation and testing method

In terms of the classification system software architecture, the processing stage responsible for handling transitions is located as shown in Figure 5.4 on the next page (labelled "Transition smoothing filter"). The filtering follows the "Posture classification" stage so that it can operate on the classifier output. All three filters described in the previous section were implemented using Python, matching the other system software components. In the implementation used in the evaluation here (as described in Section 3.4 on page 49) all three filters are made available to the classification software on the Primary Node and the visualiser provides the user with the ability to switch between them.

For testing and evaluation of the filters, a decision tree was trained using data gathered from 10 subjects as summarised in Table 5.1. The subjects wore nine sensors (located on the calves, thighs, upper arms, lower arms and chest) and WVar was used as the extracted data feature (see Section 6.4

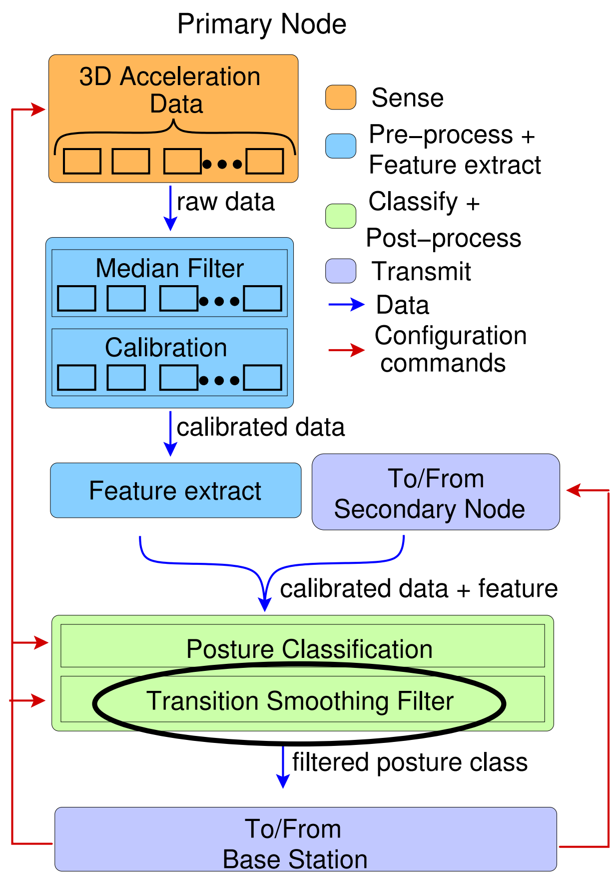

Figure 5.4: Location of transition smoothing filter in the system data flow.

Table 5.2: Summary of experimental data used for testing the transition smoothing filters. *This is a subset of data sets gathered listed in Section 4.7 on page 69.*

	No. of subjects	Regimes performed				Total time (hours)
		R1	R2	RT30	RT40	
D5	7	✔	✔	-	-	2.5
D6	7	-	-	✔	-	4.2
D7	5	-	-	-	✔	2.9

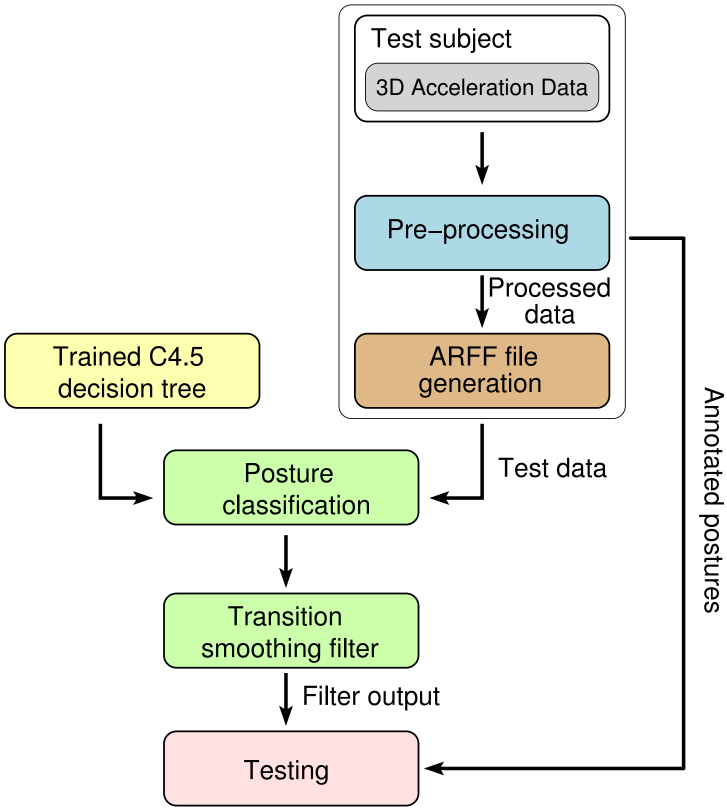

Figure 5.5: Transition evaluation method overview.

for details). While the results here are based on the full set of nine sensors used in data gathering, the optimal filter parameters were found to be the same when tested with data from only two sensors (thigh and calf). Testing data was gathered from a total of 8 individual subjects performing R1, R2 and two other regime variants (RT30 and RT40, with set transition frequencies). A summary of the testing data is given in Table 5.2 on page 89. The data was annotated using the automated method described in Section 4.8 on page 69 and the decision tree training process detailed in Section 4.9 on page 72 was followed. Figure 5.5 shows the transition testing method used here. The data from each subject was tested by passing it through the trained tree followed by a transition smoothing filter. Classification results during transitions were considered correct if they matched either the initial or final posture.

Two metrics were selected to evaluate the impact of the filters on the system classification performance: classification accuracy and number of posture change events generated. Accuracy is a common metric for this purpose and so allows comparison of different classification systems in a broad sense. As one of the goals of the filters presented here is that they provide a stable view of posture, even during transitions, the number of posture changes identified during a given set of data is used as the second metric. As the system is anticipated to send only posture updates (rather than continuously reporting), this metric corresponds to the potential improvement in battery life.

5.4 Transition smoothing filter testing results

Classifiers were trained and tested using the method described in the previous section. The output of the classifier was filtered using each of the algorithms described in Section 5.2, and the classification accuracy and number of posture change events generated were calculated for a variety of algorithm parameter values (window size for the voting filter, α for EWV, and q for the Bayes filter). The results are shown in Figures 5.6 to 5.8. The classification accuracy and number of events are calculated as an average over *all* the data samples from the 26 tests (not as averages of the accuracy and events from individual tests). Comparisons are performed against a *continuous reporting* system, and a *basic events filter* (which simply compares the current output with the previous one and transmits if different). For each filter, the parameter giving the largest accuracy gain was selected as the optimum. In all cases, the number of events generated using the selected filter parameter were near to the lowest found for that filter.

All three proposed transition smoothing filters improved the performance in terms of classification accuracy and the number of events generated. From the graphs, it can be seen that the optimal parameters were:

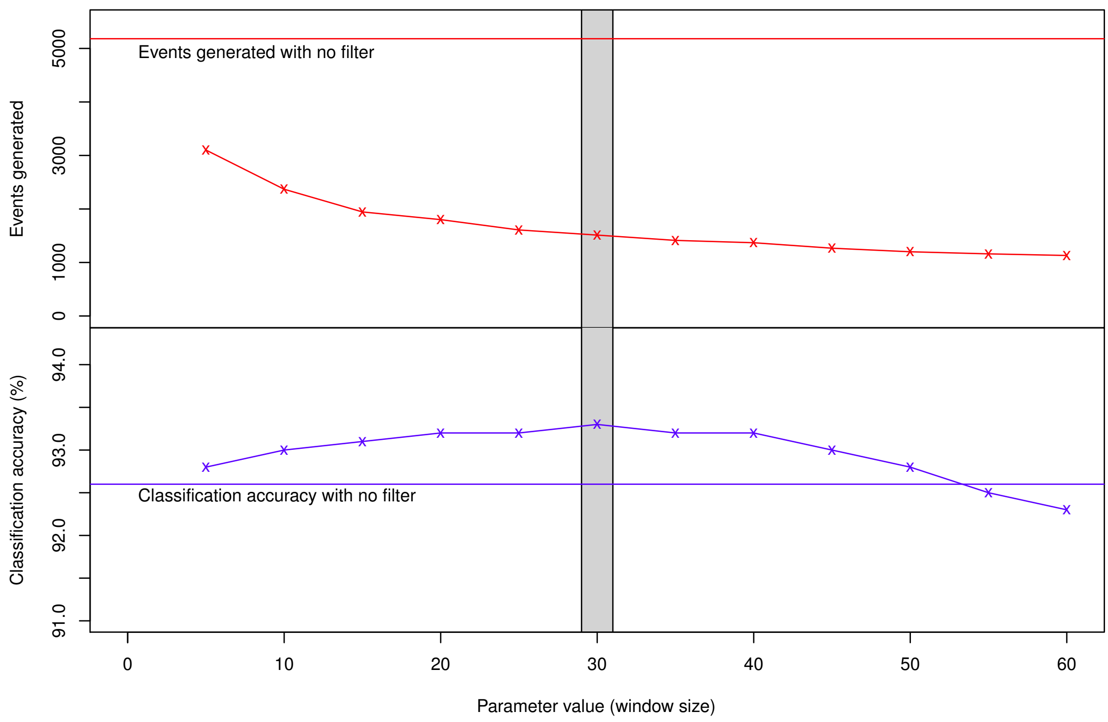

Figure 5.6: Classification accuracy and generated events for Voting filter with various window sizes.

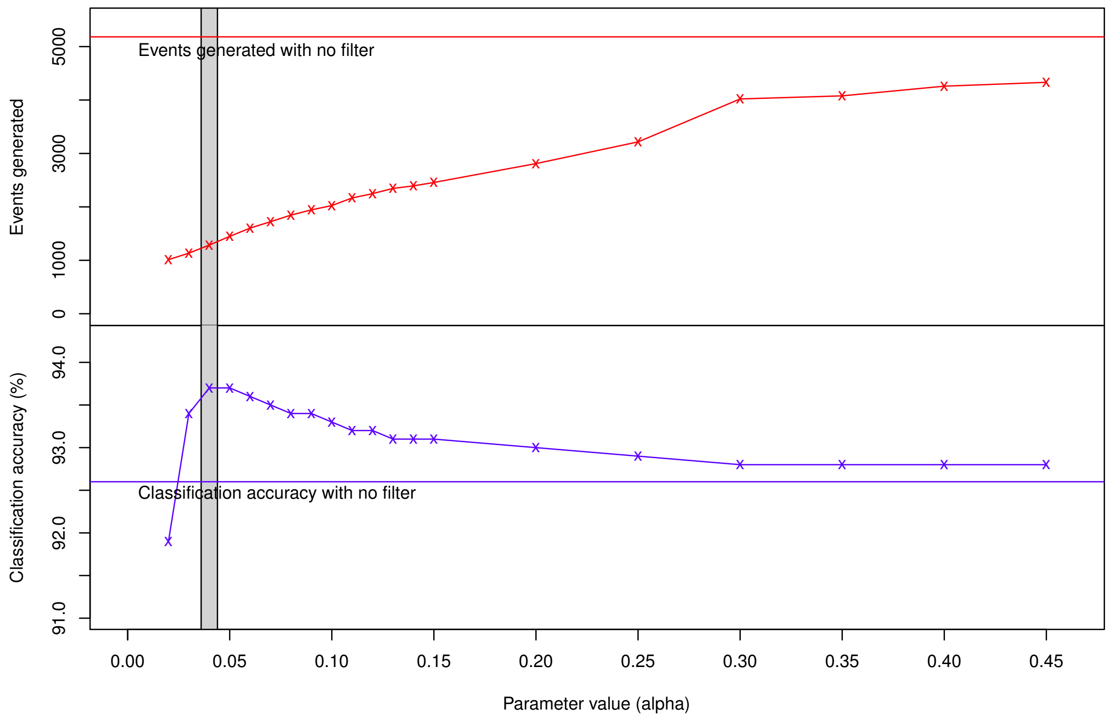

Figure 5.7: Classification accuracy and generated events for EWV filter with various values of α.

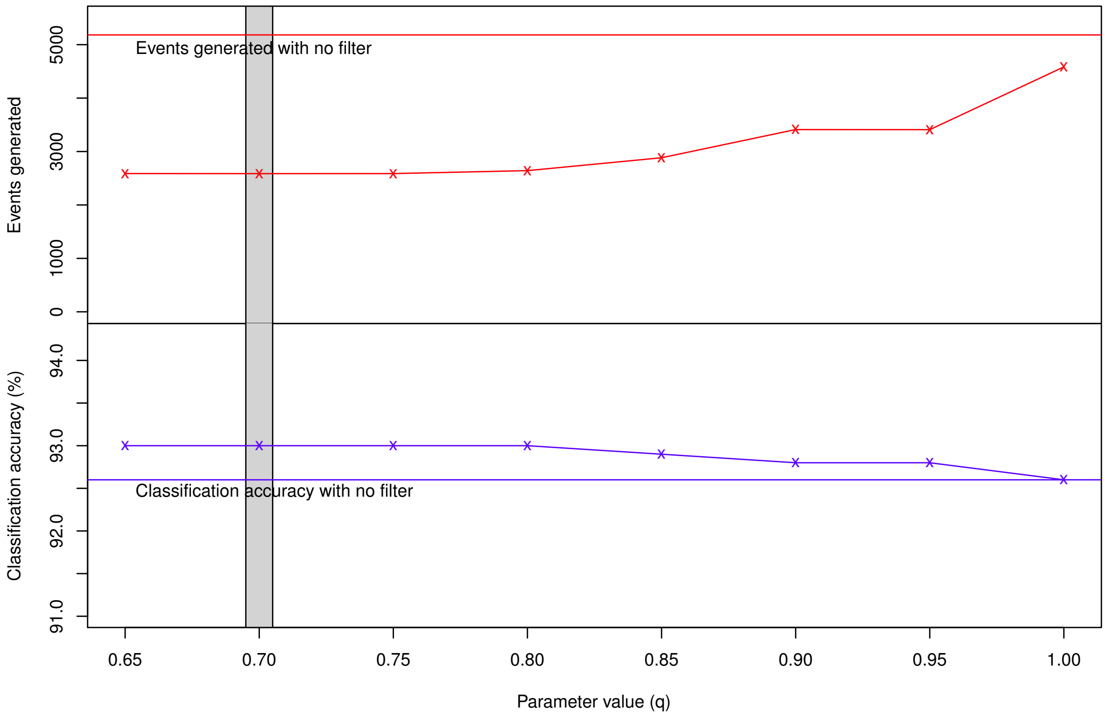

Figure 5.8: Classification accuracy and generated events for Bayes filter with various values of q ($p = 0.998$).

Table 5.3: Classification accuracy and number of events generated for optimal transition smoothing filter parameters. *Total transmissions for a continuous monitoring system were 343,140 for the data shown. RCMT=Reduction against continuous reporting transmissions.*

	Filter type and parameter value			
	Basic events filter	Vote	EWV	Bayes
		window = 30 s	$\alpha = 0.04$	$q = 0.70$
Accuracy (%)	92.6	93.3	93.7	93.0
Events	5182	1512	1285	2586
RCMT (%)	98.5	99.6	99.6	99.2

- a window size of around 30 samples (corresponding to 3 seconds) for the voting filter,
- an α of around 0.04 for the EWV filter, and,
- low values of q for the Bayes filter (values near to 1 resulted in more posture change events being generated).

A summary of the results for the optimal parameters is given in Table 5.3 on the previous page. The EWV filter had the best overall performance of the three smoothing filters, increasing the accuracy by 1.1% (compared to systems which do not cater for transitions) and reducing the number of events generated to around a quarter compared to using a basic events filter.

The increase in accuracy is dependent on the frequency of transitions. Figure 5.9 on the facing page shows the effect of the number of transitions per minute on the overall classification accuracy (for untruncated data) with no filter and when the best filter is used (EWV filter with $\alpha = 0.04$) on the output of the tree. These results are calculated using the accuracy on truncated data as the baseline for comparison. It can be seen that the classification accuracy decreases more rapidly when no filters are used (around 2% for every transition/minute) compared to when the EWV filter is used (less than 1% for every transition/minute).

Figures 5.10 and 5.11 show the effect of the EWV filter (with $\alpha = 0.04$) on the output of the classifier in terms of the postures classified during a sample sequence of data samples in which a subject transitioned from sitting to lying on one side. When no filter is used there are 11 posture changes identified, while use of the EWV filter reduces this to 3 posture changes. This is consistent with the results shown earlier, where a reduction of 75.2% was seen overall in the events generated when using this filter. It is expected that battery life will also be improved by the reduced number of transmissions.

5.5 Real-life evaluation of implemented filter

This section presents an evaluation of the EWV transition smoothing filter based on data gathered from five subjects performing tasked activities. The evaluation was performed using data gathered from the same experimental trials described in Section 4.4 on page 63. The details will not be repeated here but in summary the best performing classification tree was selected and deployed on the system hardware to classify the posture of five subjects (three males and two females) while they performed a series of tasked activities. The DG1 prototype system variant was used, with the Primary Node configured to transmit only posture information to the base station.

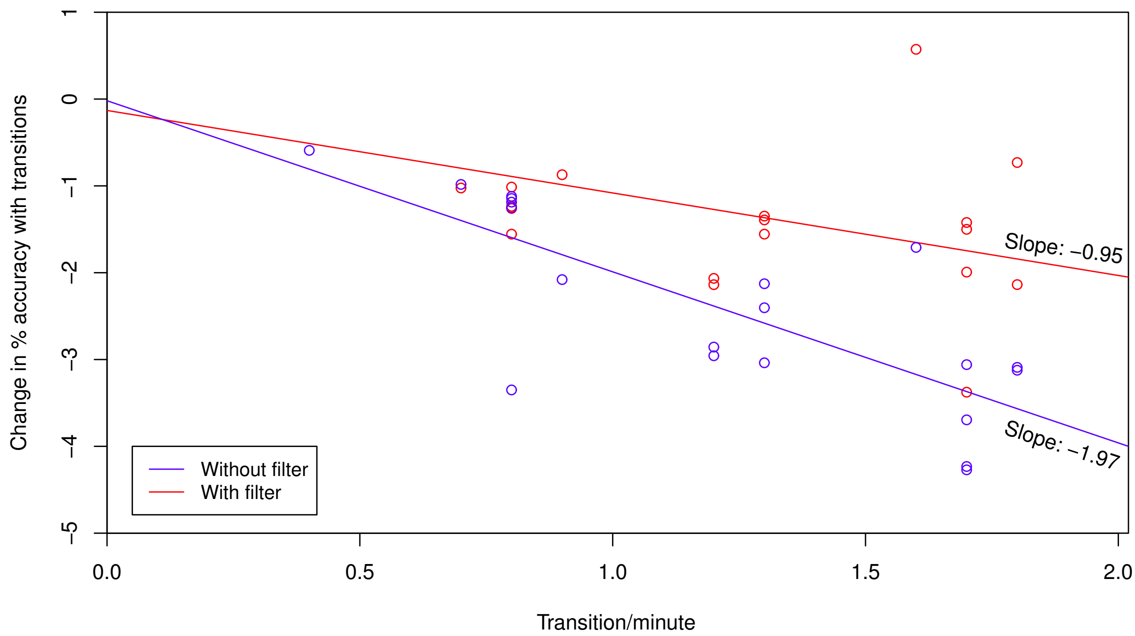

Figure 5.9: Effect of transition frequency on classification accuracy with and without transition smoothing filters. *Note that the fits predict minimal change in accuracy when there are no transitions occurring, as would be expected.*

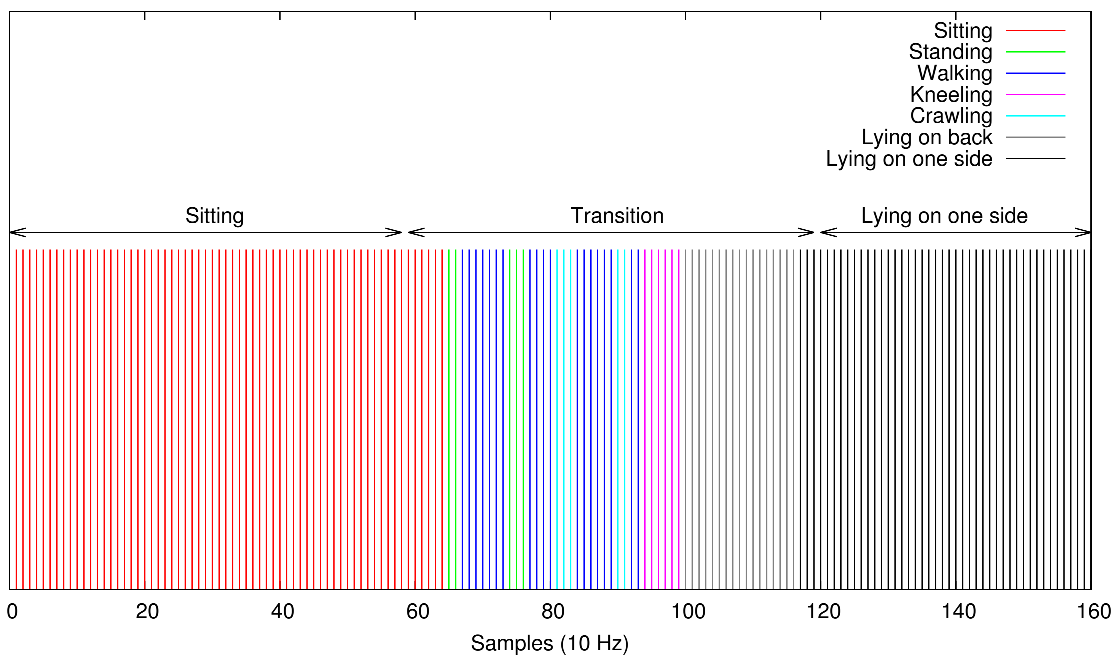

Figure 5.10: Classification output sequence when a subject is transitioning from sitting to lying on one side when no transition filters are applied.

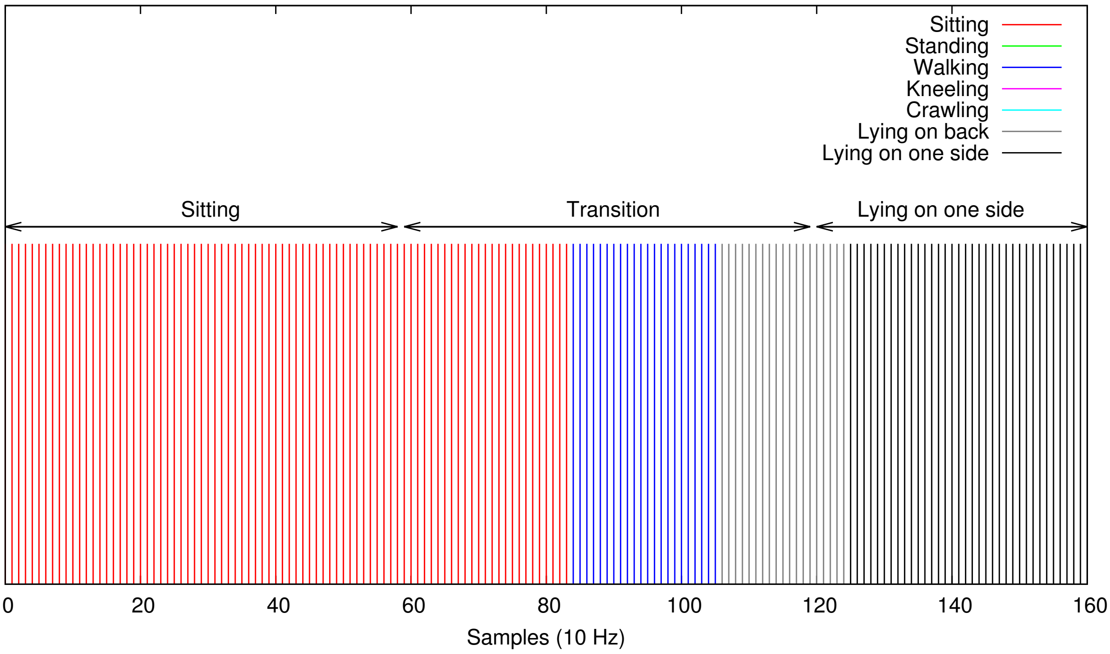

Figure 5.11: Classification output sequence when a subject is transitioning from sitting to lying on one side when EWV filter is used.

Classification accuracy was calculated over both untruncated data (as received from the node and annotated using a video based method) and truncated data (with periods annotated as *transitions* removed). The number of posture change events generated was counted using the untruncated data. The data yield was calculated by comparing the number of classified postures over the course of the test with the expected number based on the test duration and the sampling frequency of 10 Hz. The accuracy and number of posture change events was also calculated using untruncated data after filtering using the EWV filter with $\alpha = 0.04$ (giving the best overall performance as described in Section 5.4).

Table 5.4: Summary of real-time evaluation results for five subjects when the implemented transition filter is used. *A is the accuracy on untruncated data, ΔA_U is the overall accuracy loss due to transitions without filtering (comparing truncated and untruncated data), and ΔA_F is the overall accuracy loss due to transitions when the smoothing filter is used.*

	Subject	Transitions/min	A	ΔA_U	ΔA_F	Event reduction due to filter (%)
Test 1	1	1.6	98.5	-2.6	-0.6	82.9
Test 2	1	3.9	96.8	-7.4	-3.9	74.5
Test 3	2	1.7	98.1	-4.3	-1.6	76.7
Test 4	3	2.7	94.2	-3.9	-0.8	75.6
Test 5	4	1.6	96.2	-3.8	-0.5	82.3
Test 6	5	1.8	97.3	-4.3	-2.2	75.6

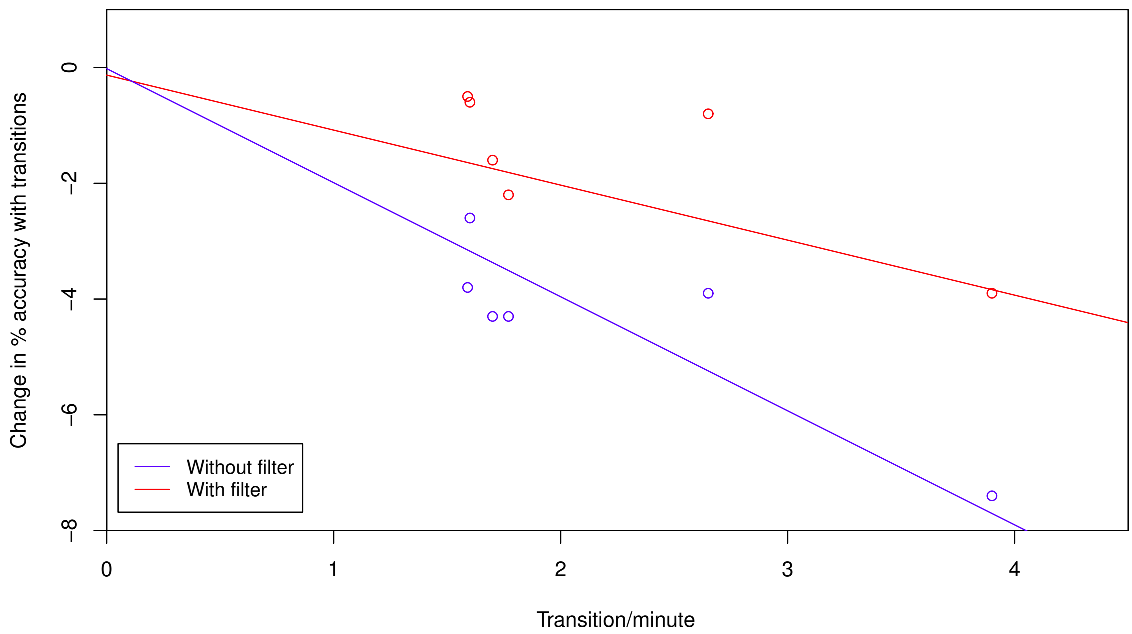

Figure 5.12: Change in classification accuracy against transitions/minute for real-time evaluation. *Lines shown are fits to the data in Figure 5.9 on page 95.*

Table 5.4 on the facing page presents the results for all five subjects. Additionally, the reduction in transmissions for an event-based system compared to a continuous monitoring system was found to be similar to the results shown in Table 5.3 on page 93—an average of 97.2% without the filter applied and 99.4% with the filter applied. Figure 5.12 shows the change in accuracy going from truncated to untruncated data based on the rate of transitions. It can be seen that the results match relatively well to the fit lines shown in Figure 5.9 on page 95, with similar RMS errors against the fit line for the values here and the values for the offline evaluation. This validates the relationship found between the reduction in accuracy and the frequency of transitions.

Based on the results, it can be concluded that:

- The relationship between transition frequency and accuracy loss is maintained in online evaluation of the classifier.
- The accuracy loss when moving from truncated to untruncated data is reduced when the smoothing filters are used.
- Use of the proposed filters significantly reduced the number of events generated.

The results found with real-time classification and the subjects performing tasked activities match with the results presented in Section 5.4 (where the data was processed offline).

5.6 Summary

Transitions between postures can be difficult to accommodate when using classification methods such as supervised learning. In the EOD case study application here and other similar applications, it is important to suppress high frequency transition reports as they are likely spurious and do not represent useful information. Based on this, transition smoothing filters were introduced as a means of achieving two main goals:

1. an improvement of the accuracy of classification for data containing both stable postures and transitions, and,

2. a reduced number of generated posture updates, providing benefits in battery life and required network bandwidth as well as providing more stable input to an external subsystem.

The approach used was based on three filters: a simple voting filter, an EWV filter, and a Bayes filter. The filter candidates were tested using experimentally gathered data.

The filter evaluation was based on two metrics: 1) classification accuracy and 2) the number of posture change event messages generated. Classification accuracy during transition periods is calculated here by assuming that the classifier should output either the initial posture or the final posture. All filters delivered an increase in classification accuracy and a significant reduction in the number of events compared to an unfiltered system. The EWV filter with $\alpha = 0.04$ provided the best overall performance of the three post-processing filters with an overall increase in classification accuracy of 1.1% and a 75.2% reduction in events generated.

The effect of transition frequency on classification accuracy was also considered. It was shown that, as the transitions frequency increased, the accuracy decreased approximately linearly. When using the best performing filter (EWV) the decrease in overall classification accuracy was less than 1% per transition/minute, compared to 2% if no filter is used.

The EWV filter was also evaluated as deployed in a realistic scenario (presented in Section 4.12 on page 76). The results confirmed the predictions derived from prior off-line analyses data in terms of: 1) change in accuracy with transition rate for both filtered and unfiltered data and 2) reduction in the number of posture change events generated when a filter is used.

In summary, the transition smoothing filters described here (and specifically EWV, the best performing filter) provide an increase in classification accuracy when integrated with a deployed system. Furthermore, they provide a significant reduction in the number of posture change events generated, which is beneficial when designing an efficient event-based system.

The next chapter provides an in-depth analysis of the system design and data gathering factors and provides guidance as to selection of an appropriate set of parameter values to maximise classifier accuracy.

Chapter 6

The design space for a C4.5 decision tree based classifier

When building a machine learning based posture classification system, there are a large number of factors to be considered such as *sensor sampling rates*, optimal (and minimal) *sensor positioning* and optimal time-based and frequency-based *features*. Such factors can have a large impact on the accuracy of the classifier. Despite this, there appears to be little published investigation into the impact of these factors. Consequently much work is duplicated by different system developers while building and evaluating posture classification systems. This further results in little commonality between systems reported in the literature. The development effort for new applications can thus be considerable.

The focus of this chapter is thus to investigate and evaluate the impact of various system related factors and training process factors on the accuracy of a machine learning-based posture/activity classifier, specifically one based on C4.5 decision trees as proposed in Chapter 4. The factors considered are: 1) extracted data feature choice, 2) data feature window size, 3) number of sensors, 4) training set size, 5) sampling rate, and 6) targeting of individual postures. The design space is explored such that the results and discussion provide guidance for posture classification system design in a range of applications with differing constraints. From these results, the optimal configuration for the case study application is extracted. This evaluation and the subsequent recommendations form the contribution of this chapter and define a decision tree classifier design space.

The chapter is structured as follow: Section 6.1 describes the classifier and data gathering factors around which the design space is formed. Section 6.2 presents the data sets used in the analysis in this chapter. Section 6.3 lists the hypotheses guiding the analysis. Sections 6.4 to 6.9 present the investigations into the various factors and the resulting conclusions about the classifier design space. Finally, Section 6.10 summarises the work presented in this chapter.

6.1 Factors affecting classification accuracy

When designing and implementing a classifier for real life applications, it is important to consider all of the factors that can impact its accuracy when deployed in the field. Despite the large number of options with regard to, for example, sensor placement and data feature extraction, the impact of these factors has not yet been subject to extensive research. Furthermore, the design and tuning of a classification system involves application-specific considerations such as the specific set of postures that must be classified. The expectation that researchers have is that system performance in a real-life deployment will be similar to that observed within the laboratory or assessment through offline testing. However from a BSN perspective, when critically evaluating the works presented in the literature, it has become clear that:

- bespoke systems and parameter settings are reported for each application treated,
- the systems reported are commonly validated only in laboratory settings,
- the performance of these systems is evaluated in well-controlled scenarios and the performance indicators vary over the research community (with precision, recall, and accuracy being most common),
- a large degree of variability exists in the system design, with factors such as sampling rate and the number and type of accelerometers used, for example, being related more to hardware availability and capability than the requirements of the application, and,
- the ability of such reported systems to perform and generalize in real-world scenarios is largely unexplored to date.

Based on these perceived gaps in the literature, this chapter presents an investigation into the following factors.

- Extracted data features. A variety of data features (such as windowed variance) can be used to capture information about the dynamic nature of some postures. When used in addition to sensor data, the classification accuracy for some postures can thus be increased. However, the impact and benefits of various data features are dependent on the set of postures that must be classified. Data features are investigated in Section 6.4.
- The data feature window size. Some features are calculated over a window of data and the size of this window will affect the classification accuracy. In particular, the larger the window size, the longer it will take for certain features (such as windowed mean) to react to a rapid change in sensed data, potentially leading to a period of incorrect classifications. Section 6.5 demonstrates this effect.

- The number and location of sensors deployed on the subject. The sensors form the basis of the system operation and thus selecting a suitable number of sensors and optimal locations for those sensors is a vital part of the system design. The results of this analysis are given in Section 6.6.

- The sensor sampling rate. The sampling rates used in the literature are often based on the capabilities of the hardware platform available, without an analysis of the need for a given rate. Sampling at a lower rate often allows the use of cheaper components and reduces the required communication bandwidth and is thus preferable as long as classification accuracy is not compromised. The classification accuracies obtained here at several sampling rates are shown in Section 6.7.

- The number of subjects forming the training set for the classifier. The goal of training a machine learning based classifier is to allow it to form rules that generalise well to unseen data. While it is not necessarily possible to train a classifier that will generalise to *all* possible future subjects, the greater the number of subjects that are included the more likely it is that the generated rules will be suitable for a large spread of subjects.The effect of training set size is investigated in Section 6.8.

6.2 Data sets used for design space investigation

Table 6.1 on the next page lists the data sets used in this chapter, showing the relevant data sets taken from Table 4.2 on page 70. The data was collected using all three data gathering systems (described in Section 4.5 on page 66. A description of the data collection process is given in Section 4.3 on page 61.

The analysis in this chapter is based on the use of LOSOXV across data from a number of subjects. The exception to this is the investigation into the number of subjects to use in the training set (Section 6.8); for this, the test method used is described separately. The particular data sets used in each analysis are listed in the appropriate section, referring back to Table 6.1 on the next page.

6.3 Summary of hypotheses

In order to guide the analysis presented in this chapter, several hypotheses were formulated in relation to the various identified factors affecting classification accuracy. These hypotheses describe the expected results for a number of tests, with the expectation based, in part, on results found in related work from the literature. The hypotheses are listed below, and the following sections treat them individually.

H1 To classify accurately dynamic postures, time domain features, such as windowed variance, are essential.

Table 6.1: Summary of experimental data used in the analysis in this chapter.

	Data acquisition system	No. of sensors	No. of subjects	Annotation method	Regimes performed R1	R2	R3	RT40	Hours of data	Used in sections
D1	DG1	9	5	Manual	✓	✓	-	-	1.9	6.4, 6.6, 6.8
D2	DG1	9	1	Manual	✓	✓	✓	-	0.5	6.4, 6.6, 6.8
D3	DG1	11	5	Manual	✓	✓	✓	-	2.6	6.4, 6.6, 6.8
D5	DG2	9	7	Automatic	✓	✓	-	-	2.5	6.4, 6.5, 6.6, 6.8, 6.9
D7	DG2	9	5	Automatic	-	-	-	✓	2.9	6.5
D8	DG2	4	6	Manual	✓	-	✓	-	3.0	6.6
D9	SHIMMER	7	7	Manual	✓	✓	-	-	2.2	6.7
D10	SHIMMER	7	1	Manual	✓	-	-	-	0.2	6.7

H2 When using the windowed variance feature, classification accuracy will be highest with a window size of several seconds, decreasing for smaller and larger window sizes (below 0.5 seconds and above 5 seconds).

H3 Increasing the number of sensing locations beyond two sensors (situated on the thigh and calf) produces an increase in classification accuracy but with diminishing return as the number of sensors increases.

H4 When using time domain features, increasing the sampling rate beyond 10 Hz will not provide an increase in classification accuracy.

H5 Obtaining a consistently high classification accuracy will require a certain minimum number of subjects in the training set. However, further increase in the dataset size will not show a matching increase in the classification accuracy.

H6 The posture classifier can classify specific targeted postures with a similar accuracy to the results found when classifying all postures.

6.4 Selection of an appropriate data feature

Hypothesis 1 To classify accurately dynamic postures, time domain features, such as windowed variance, are essential.

The need for feature extraction is explained in Section 4.11.1 on page 74. In summary: an algorithm that considers only data samples from one time step cannot capture the dynamics or history of the measured phenomenon. As a concrete example, standing and walking will result in similar or identical raw data values at some points in time. A solution to this is to extract temporal features from the data. It was shown (in Section 4.11.1 on page 74) that the classifier algorithm used achieved an accuracy of 98.0% when classifying only static postures, reducing to 90.5% when both static and dynamic postures were considered (using only raw data). The goal of using a feature such as the ones described here would thus be to counteract that loss of accuracy when classifying dynamic postures alongside static ones. Ideally, classification based on the raw data and extracted feature data would be able to achieve an accuracy similar to the 98.0% observed for static postures, across both static and dynamic posture sets.

The two primary types of data feature are *time domain* and *frequency domain* features. However, as Bharatula [18] points out, frequency domain features alone yield relatively poor classification results for accelerometer-based activity classification of the type treated in this thesis. Since the aim here is to perform all processing on a low-power microcontroller/microprocessor, and since frequency domain feature extraction is computationally complex, there seems little advantage in considering a large number of frequency domain features. Furthermore, the additional information provided by frequency domain features are not called for in this application—it is not required, for example, to distinguish between slow and fast walking rates. The one exception made was for *Energy* as it was the most common frequency domain feature encountered in literature and was thus included for comparative purposes.

This section presents the effect of various extracted data features on classification accuracy. The features considered, based primarily on their use in the literature, were: Windowed Mean (WM) [14, 42, 56, 97, 126], magnitude [15], SVM [35], WVar [42, 56, 97, 116, 126], RMS [9, 32, 46, 59, 87, 126], Energy [4, 39, 57, 116, 126], SMA [84, 65, 68, 126], and EWMA. Data was collected at 10 Hz with a 30 sample window used for WM, SVM, WVar, RMS, Energy, and SMA. In each case the raw data was used along with the additional calculated values for the feature as inputs to the classifier. The feature descriptions follow.

Windowed Mean WM is the mean acceleration for a particular axis over a fixed period of time (or window) and can be used to minimise the effect of movement. Unfortunately, this also causes a tendency to confuse periodic movement with similar stationary postures. WM is calculated as

$$x_{\mathrm{wm}} = \frac{1}{n}\sum_{i=1}^{n} x_i$$

where x_i is sample i from the window of values, and n is the size of the window in samples.

Root Mean Square Similarly to WM, the RMS value over a window represents a trade-off between minimising the impact of sudden movement and incorrectly identifying periodic movements. RMS is calculated as

$$x_{\mathrm{rms}} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} x_i^2}$$

Exponentially Weighted Moving Average EWMA is an average that gives higher importance to new observations while not discarding old observations. This means that it is potentially capable of responding faster to changes in the data than a windowed mean type feature such as WM or RMS. Strictly speaking, it is a weighted mean where the weights for older data points decrease exponentially. EWMA is often calculated recursively as follows:

$$s_t = s_{t-1} + \alpha\left(x_t - s_{t-1}\right)$$

where x_t is the sensed data point at time t, s_t is the weighted sum of the data points, s_{t-1} is the previous weighted sum, and α is a constant determining the degree of weighting applied.

Windowed Variance WVar is a calculation of the variance of the gathered data over a window. This feature provides an indication of the degree to which the acceleration is oscillating but is independent of the mean value. WVar is calculated as

$$x_{\text{wvar}} = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2$$

where $\bar{x}$ is the mean data value over the window.

Signal Magnitude Area SMA is the sum of the mean magnitudes of the x, y, and z values from a single accelerometer over a window. It gives an indication of the overall acceleration being applied to a given sensor. Prior to the SMA calculation, the individual signals are passed through low-pass and high-pass filters to attempt to isolate frequencies relevant to movement (the cut-off frequencies for the low-pass and high-pass filter were both 0.25 Hz). SMA is calculated as

$$x_{\text{sma}} = \frac{1}{n}\left(\sum_{i=1}^{n}|x_i| + \sum_{i=1}^{n}|y_i| + \sum_{i=1}^{n}|z_i|\right)$$

Energy Energy is calculated as the sum of the squared discrete Fast Fourier Transform (FFT) component magnitudes of the signal within a window, divided by the length of the window for normalisation. This is used by Ravi *et al.* [104], and is calculated as

$$x_{\text{energy}} = \frac{1}{n}\sum_{i=1}^{n}|F_i|^2$$

where $F_i = \sqrt{a_i^2 + b_i^2}$, where a_i is the real part of the FFT and b is the imaginary part of the FFT.

Magnitude For the purpose of feature extraction here, Magnitude is defined as the vector magnitude of the signal derived from the three axes of a triaxial accelerometer. It is calculated as

$$x_{\text{mag}} = \sqrt{x^2 + y^2 + z^2}$$

Signal Vector Magnitude SVM is the sum of the vector magnitudes of the signal from the three axes of an accelerometer within a window, divided by the length of the window. Essentially, it provides the average magnitude of the signal over the window. It is calculated as

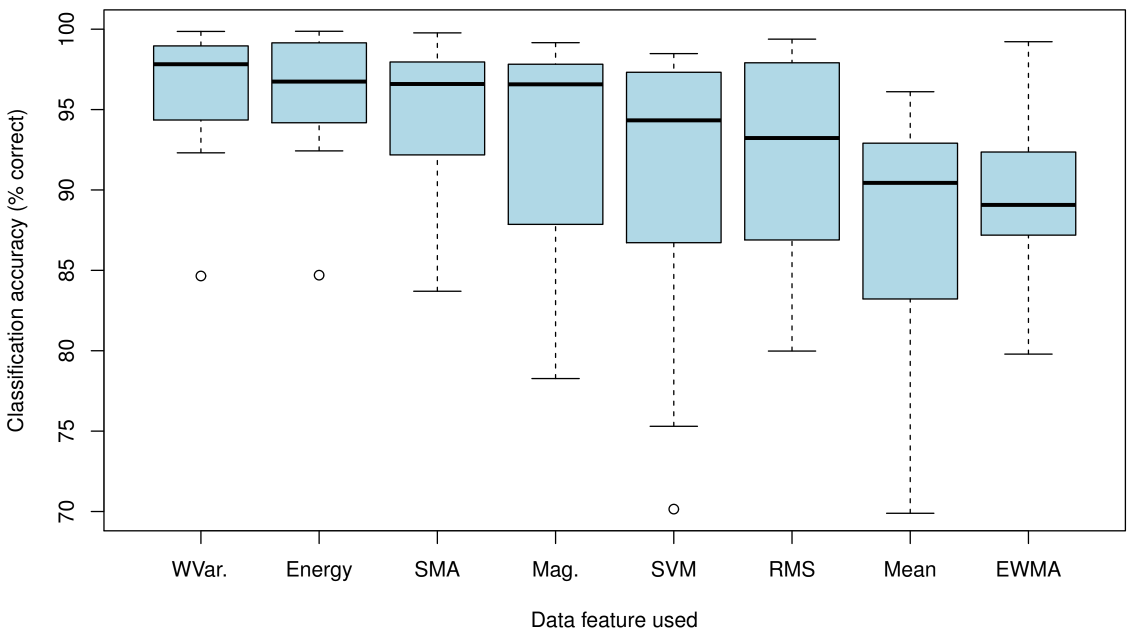

Figure 6.1: Classification accuracy when using extracted data features.

$$x_{\mathrm{svm}} = \frac{1}{n} \sum_{i=1}^{n} \sqrt{x_i^2 + y_i^2 + z_i^2}$$

LOSOXV (detailed in Section 4.3 on page 61) was used for testing purposes over data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects). Figure 6.1 shows the classification accuracy for each extracted data feature when considering the set of eight static and dynamic postures treated in this thesis. It can be seen that the best accuracy overall was obtained from WVar. The overall classification accuracy with WVar was 96.3%, which is a reduction of only 1.7% compared to the accuracy obtained when considering static postures alone and using only raw data. WVar thus appears to be the best performing feature for the set of postures defined in this work.

In addition to the above, it is useful to know the best and worst features for each posture and the accuracies achieved in each case. This information will allow a feature to be chosen in applications requiring a subset of the postures used here and provide a starting point for investigation of additional postures. This analysis was performed using the results of the testing presented previously but with the results for each posture extracted from the overall results.

Table 6.2 on the facing page presents the best and worst performing feature for each posture. While almost every posture appears to have a different best associated feature, in reality the results are very similar for each feature in some cases, as can be seen by observing the accuracy given for the worst

Table 6.2: The best and worst performing feature per posture. *PPV = Positive Predictive Value (TP/(TP+FP) where TP is true positives and FP is false positives).*

	Crawling	Kneeling	Lying on back	Lying on front	Lying on side	Sitting	Standing	Walking
Best feature	Mean	RMS	SVM	Mean	WVar	RMS	WVar	WVar
PPV	98.6	90.4	96.3	99.9	98.7	97.8	98.2	96.2
Std. dev.	1.2	15.6	9.5	0.2	3.1	4.2	1.7	6.0
Worst feature	Mag	EWMA	Mean	RMS	EWMA	EWMA	Mean	Mean
PPV	96.6	83.6	88.0	85.6	92.64	95.9	68.9	88.0
Std. dev.	6.1	28.7	27.0	32.8	16.1	10.8	26.7	16.3

feature. The main difficulty in classifying the postures appears to be for standing and walking, which are mistaken for each other by many of the features. Note particularly the 68.9% accuracy for standing when using Mean. For both of these, WVar produced the best results. Notably, the worst classification accuracy for six of the eight postures was produced by Mean and EWMA.

6.5 Selection of window size for data features

Hypothesis 2 When using the windowed variance feature, classification accuracy will be highest with a window size of several seconds, decreasing for smaller and larger window sizes (below 0.5 seconds and above 5 seconds for example).

This section provides an analysis of the effect of window size on classification accuracy. Time-domain data features are generally applied over a fixed or sliding window of data. The expectation is that larger window sizes will cause a period of incorrect classifications after each postural transition while the window becomes populated with data in the new range of sensed values. The new posture will not be classified correctly until the window includes sufficient post-transition data samples for it to be recognised. On the other hand, a small window will not allow the feature to be calculated meaningfully (as an extreme example, variance calculated over a window of one sample will always be zero). It is necessary, therefore, to quantify the bounds within which a high classification accuracy can be obtained.

This investigation was performed over data sets D5 and D7 (5.4 hours of data from 8 subjects). The WVar feature was extracted over window sizes starting at 5 samples and then increasing in steps of 5 up to 200 samples. The LOSOXV process was performed for each window size in turn over the dataset. Classification accuracy was calculated based on the truncated data by comparing the data annotation with the classifier output. Classification delay was evaluated by counting the number of samples between the

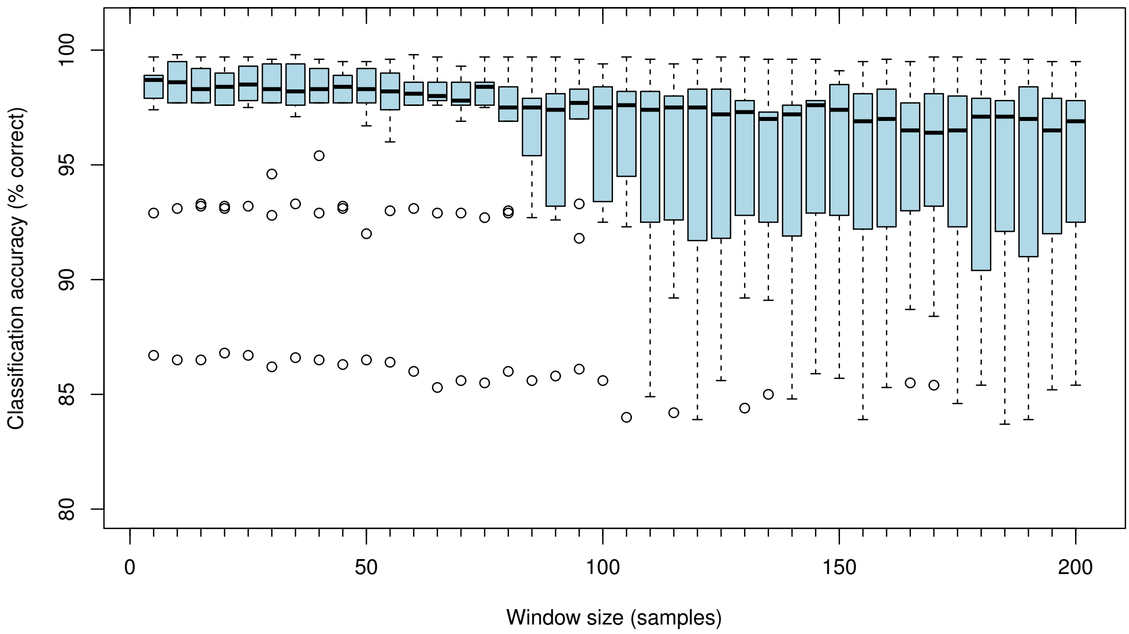

Figure 6.2: The effect of window size on classification accuracy.

end of an annotated transition to the first time that the posture was classified correctly. If the annotated posture changed before any correct classifications were encountered then the count was discarded and a new one started.

Figure 6.2 shows the effect of window size on classification accuracy. A size of between 5 and 75 samples (0.5 to 7.5 seconds) appears to result in relatively consistent accuracy results. Window sizes above 75 samples show greater variation in the results along with a decreasing median accuracy as the window size increases.

Figure 6.3 on the facing page shows the impact of window size on the time period in which incorrect classifications are output following a postural transition. While the results are variable, it can be seen that the period length tends to increase with larger window sizes, giving an average period of 0.5 seconds for a window size of 5 samples and 1.0 seconds for a window size of 200 samples. Inspecting the average period for each trial conducted, the highest for a single trial was found to be 3.37 seconds (for a window size of 200 samples). Overall, it can be seen that the window size does impact the period of incorrect classification following a posture change, though the period does not increase rapidly with increasing window sizes.

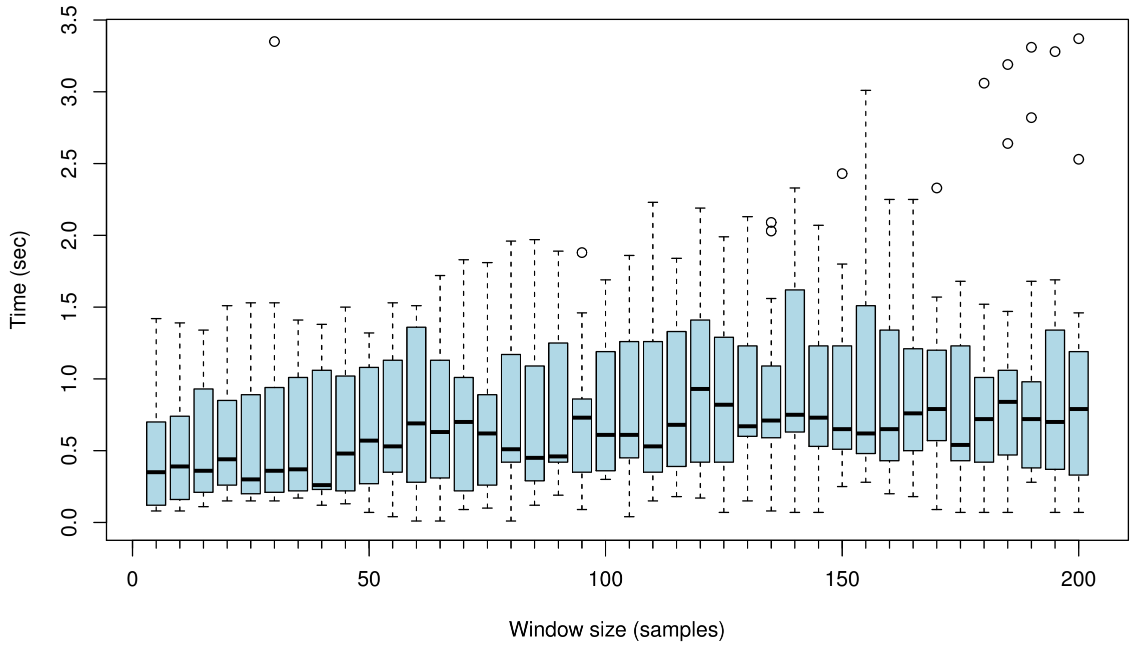

Figure 6.3: Time to first correct classification after a postural transition when using a range of window sizes.

6.6 Selection of number and positioning of sensors

Hypothesis 3 Increasing the number of sensing locations beyond two sensors (situated on the thigh and calf) produces an increase in classification accuracy but with diminishing return as the number of sensors increases.

In the literature related to classification of postures and activities, the justification for selecting particular number and placement of the sensors is not generally reported, though some commonality is evident. The hip and ankle, for example, are often used and there are several examples of posture classification using data gathered only from a single sensor. Common locations include the hip [20, 46, 65, 79, 84, 114], wrist [59, 74, 90, 109, 108, 126] and chest [15, 68]. More details of sensor locations from the literature are given in Section 2.4.3 on page 32. Reducing the number of sensors deployed in a posture classification system that uses multiple sensors wired to a wireless node can reduce the current draw from the batteries and increase the wearability of the instrument. Moreover for fully wireless solutions, the number of sensors that may be used is limited by the wireless bandwidth available on the gateway or processing node.

In the analysis here, eleven sensor positions were investigated in various combinations of between 1 and 11 sensors. The selection of the eleven locations is justified in the context of the EOD operative

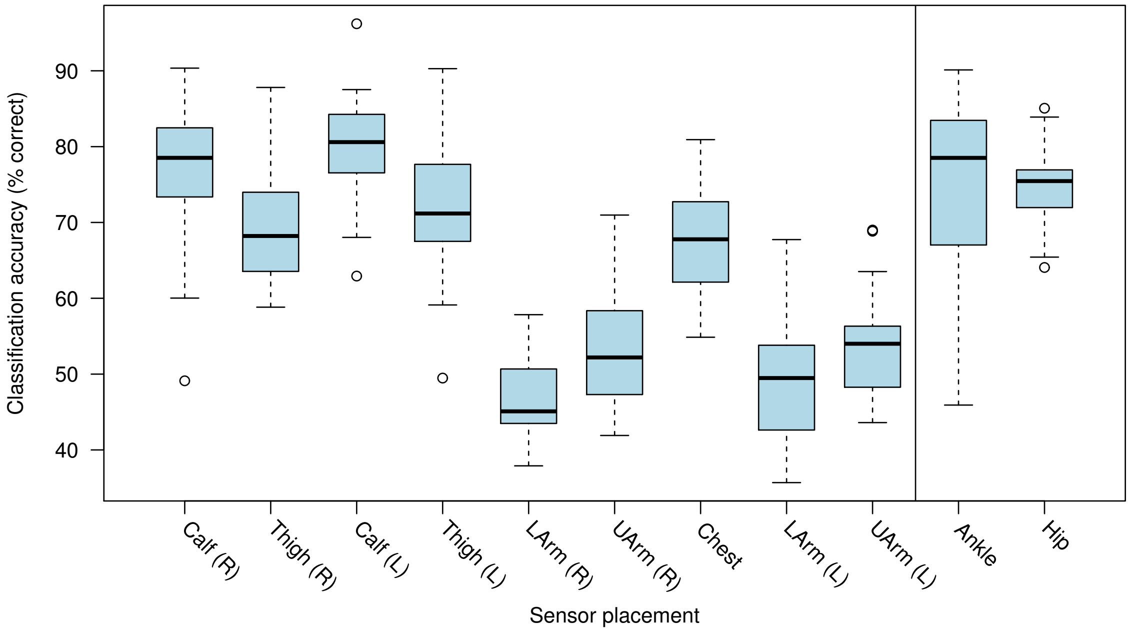

Figure 6.4: Overall accuracy for classification using one sensor location.

monitoring application by Section 3.4.3 on page 51. These locations provide coverage of the commonly investigated body segments as described in Section 2.4.3 on page 32. For each combination of sensors, the classification accuracy was tested for the set of eight postures considered here.

Data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects) were used for training for all combinations of nine or fewer sensors. Data set D3 (2.6 hours of data from 5 subjects) was used for the 11 sensor combination and data sets D3 and D8 (5.6 hours of data from 11 subjects) were used for the individual hip and ankle sensor tests. As the focus of the data gathering was on nine sensor locations, these form the basis of the majority of the discussion here. In all cases, WVar was used as the extracted data feature with a sliding window of 30 samples at 10 Hz and testing was performed via LOSOXV.

The first step of this investigation was to determine the accuracy of individual sensors when classifying the eight postures. Figure 6.4 shows the classification accuracy results when using each individual sensor location to classify all eight postures. The optimal position for a single sensor is the calf, with the arm sensors performing worst for this posture set, as expected. The chest sensor appears to perform less well than expected from the literature. The hip sensor provides adequate (but lower than calf) accuracy. Upon inspection of the generated confusion matrices, it appears that the chest location cannot be used to discriminate between standing and sitting, which is a conclusion supported by Barralon *et al.* [15].

Table 6.3: Sensors which provide an MCC greater than 0.8 when classifying individual postures.

	Crawling	Kneeling	Lying (front)	Lying (back)	Lying (side)	Sitting	Standing	Walking
Ankle	✔							✔
Lower arm (left)								
Upper arm (left)								
Lower arm (right)								
Upper arm (right)								
Chest	✔		✔	✔	✔			✔
Hip	✔		✔	✔	✔			✔
Calf (left)	✔			✔	✔			✔
Thigh (left)			✔		✔			
Calf (right)	✔			✔				✔
Thigh (right)			✔		✔			

As the results given are specific to the set of postures considered for the application here, and so may not represent the overall results for other applications, each sensor was also tested for its suitability for classifying each posture. This should provide a basis for sensor selection for alternative posture sets (particularly subsets of the postures here). For this purpose, accuracy was not used as an evaluation metric. This is because, for any given posture considered in isolation, a relatively high accuracy can be obtained by never outputting that posture from the classifier. Consider, for example, the case where the accuracy of classification for sitting is being analysed based on a dataset with equal representation of all 8 postures—a classifier which only output "other" as the result would obtain a 87.5% accuracy despite having no practical use. For this reason, the Matthews correlation coefficient (MCC) [11] was used. This coefficient is considered to be appropriate even in situations where the classes are of quite different sizes. The output is between -1 and $+1$, with $+1$ indicating perfect prediction, 0 indicating results equivalent to random prediction, and -1 indicating complete disagreement between the predicted and observed values. MCC is calculated as

$$\mathrm{MCC} = \frac{\mathrm{TP} \times \mathrm{TN} - \mathrm{FP} \times \mathrm{FN}}{\sqrt{(\mathrm{TP} + \mathrm{FP})\,(\mathrm{TP} + \mathrm{FN})\,(\mathrm{TN} + \mathrm{FP})\,(\mathrm{TN} + \mathrm{FN})}}$$

where TP is true positives, TN is true negatives, FP is false positives, and FN is false negatives. To complete the example above, the classifier outputting only "other" would produce an MCC value of 0. For the purpose of the investigation here, the LOSOXV process was performed using each sensor in turn.

Table 6.4: Use of each sensor when classifying as a percentage of total comparisons (one comparison per decision tree node). Bold text indicates results above 20%, arbitrarily chosen to indicate high values based on the maximum result found of 33.4%.

	Left lower arm	Left upper arm	Right lower arm	Right upper arm	Chest	Left calf	Left thigh	Right calf	Right thigh
All postures (aggregate)	1.5	3.4	0.6	2.7	15.0	**22.8**	14.5	18.0	**21.5**
Crawling	0.5	9.3	0.3	0.7	13.4	**29.3**	6.3	10.9	**29.4**
Kneeling	0.5	7.1	0.6	5.8	17.7	**21.2**	15.8	11.3	**20.0**
Lying on front	14.0	3.8	0.0	0.0	0.4	19.4	18.3	18.8	**25.2**
Lying on back	0.0	0.9	0.0	0.0	19.4	**20.2**	**28.7**	**24.4**	6.3
Lying on one side	1.0	2.6	1.4	0.0	8.1	15.3	**31.8**	**22.2**	17.6
Sitting	0.1	0.1	0.0	0.1	**32.8**	14.2	17.5	17.2	18.2
Standing	0.1	0.0	0.0	0.2	2.4	**33.4**	17.8	**22.0**	**24.0**
Walking	3.5	0.6	1.7	7.9	**21.8**	15.6	7.5	**23.7**	17.8

The number of true positives, false positives, and so on from each iteration of LOSOXV were summed (for each posture). A sensor was deemed to be suitable for classifying a particular posture if the MCC value for the summed results was greater than 0.8. Table 6.3 on the previous page shows the results of this investigation. It can be seen that sitting, standing and kneeling were not classified accurately using any single sensor. Using only a hip or chest sensor (as seen in the literature) resulted in accurate classification of crawling, walking, and lying on front, on back and on one side. The upper and lower arms did not allow accurate classification of any of the selected postures, reflecting the low overall accuracy shown previously for them.

The final per-sensor test looked at the importance of each sensor when classifying all eight postures. This was investigated by logging which sensor was used for each tree node visited when classifying the data described previously. The idea here is that more important sensors, such as those that appear closer to the root of the decision trees, will be used more frequently when classifying, providing a simple metric of importance. The results are shown in Table 6.4, and reflect the results given previously for individual sensors. The arm sensors are used very little, with the leg and chest sensors being responsible for the majority of comparisons made during classification.

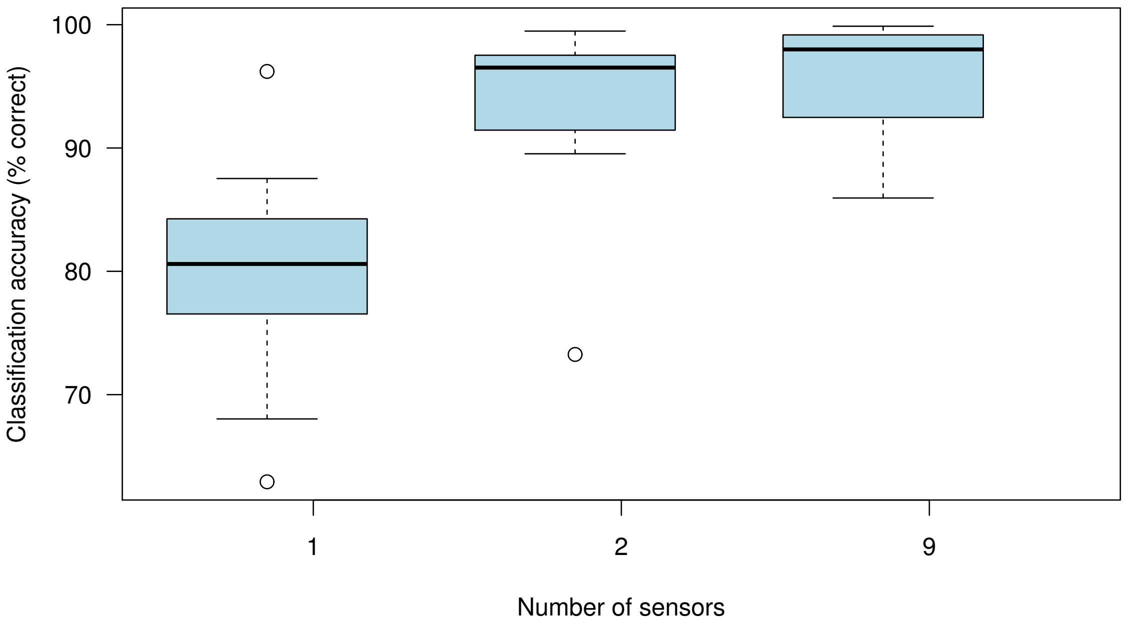

Figure 6.5: Classification accuracy when using 1 sensor (left calf), 2 sensors (left thigh and calf) and 9 sensors.

The classification accuracy found for individual sensors is significantly lower than that found when using the set of all nine sensors, so the next step of the investigation was to find the smallest set of sensors that gave a similar accuracy to using all nine. The first attempt was with two sensors, specifically the left thigh and calf as they gave the highest overall accuracies in Figure 6.4. The classification accuracy when using these two sensors was found to be 95.5%. The results are summarised in Figure 6.5. It can be seen that using these two sensors gives a similar accuracy to all nine sensors, meaning that it is possible to provide accurate classification using only two well-chosen sensors.

In conclusion, it is not sufficient to use only one sensor to recognise all eight of the postures investigated here. At least two sensors are required to recognise all postures (crawling, kneeling, sitting, standing, walking, and lying on front, on back and on one side) with a high accuracy, with two sensors on the thigh and calf giving comparable accuracy to the use of nine sensors. When using a single sensor, the best overall locations for the posture set here are the thigh and calf, though the chest will allow classification of a greater number of postures than either of these (at a lower overall accuracy) and may be preferable when targeting those specific postures.

6.7 Selection of sensor sampling rate

Hypothesis 4 When using time domain features, increasing the sampling rate beyond 10 Hz will not provide an increase in classification accuracy.

The most common sampling rates encountered in the literature are 100 Hz [32, 71, 114] and 20 Hz [22, 42, 103]. However, there are several reasons that a low sampling rate might be preferred in a given system implementation, including:

- Subjects do not normally change posture multiple times per second.
- In a continuous monitoring (rather than event based) system, a higher transmission rate can reduce battery life.

Furthermore, in the EOD Operative Monitoring application here, the Medusa2 system performs modelling at 1 Hz. Given these reasons, a low sampling rate is thus generally preferred, as long as it does not negatively impact the classification accuracy. It is expected that static postures can be accurately classified at low sampling rates as there is no movement involved and so the data will be very similar throughout a given instance of the subject performing that posture. Classification of dynamic postures (walking and crawling in this application), on the other hand, requires extraction of data features that incorporate some knowledge of the history of the data. These features are calculated over a window of data. The size of the window is set to a time interval as a trade-off between providing sufficient samples and reducing the time to provide correct classification after a transition (this effect is discussed for the WVar feature in Section 6.5).

In testing the effect of sampling rate, data was collected using the SHIMMER-based system and LOSOXV was performed over data sets D9 and D10 (2.4 hours of data from 8 subjects). WVar was extracted over window sizes adjusted to operate over the same time period in each case (a 3 second window was used, meaning that when data was gathered at 10 Hz, the window was 30 samples wide whereas when data was collected at 20 Hz, the window was 60 samples wide, for example). Resampling was performed by selecting the appropriate number of equally spaced samples from the original data that was collected at 100 Hz (for example, to test at 20 Hz, every fifth sample was used). In this analysis, only walking and crawling are considered since the static postures are expected to be classified accurately at low sampling rates.

Figures 6.6 and 6.7 on the facing page show the classification accuracy at various sampling rates for walking and crawling, respectively. It can be seen that classification accuracy appears to be relatively

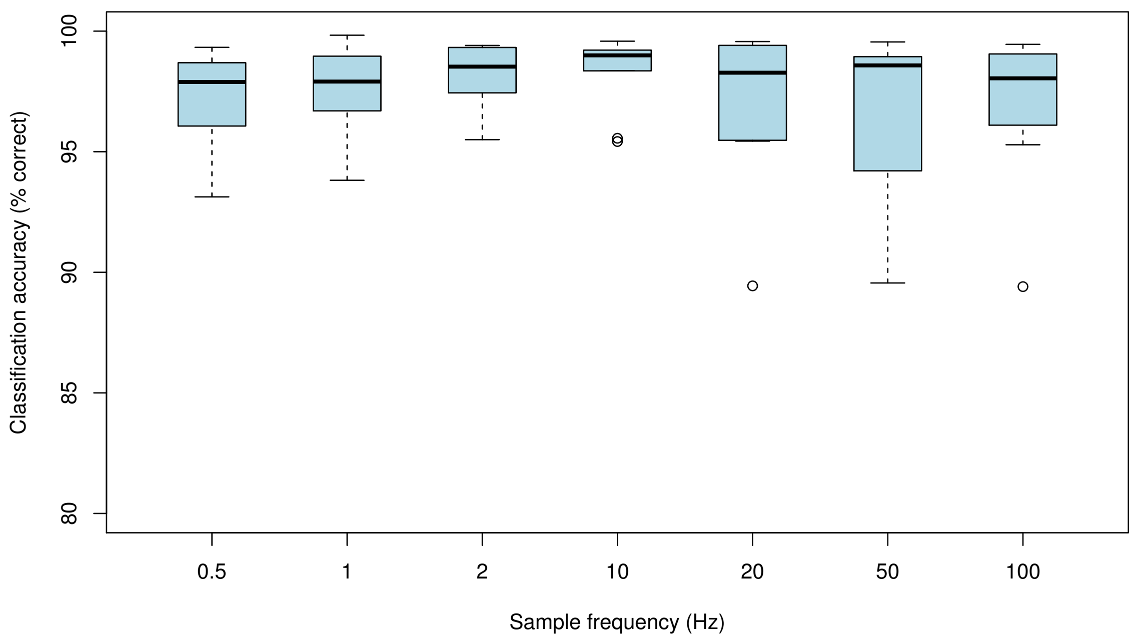

Figure 6.6: Effect of different sampling rates when classifying walking.

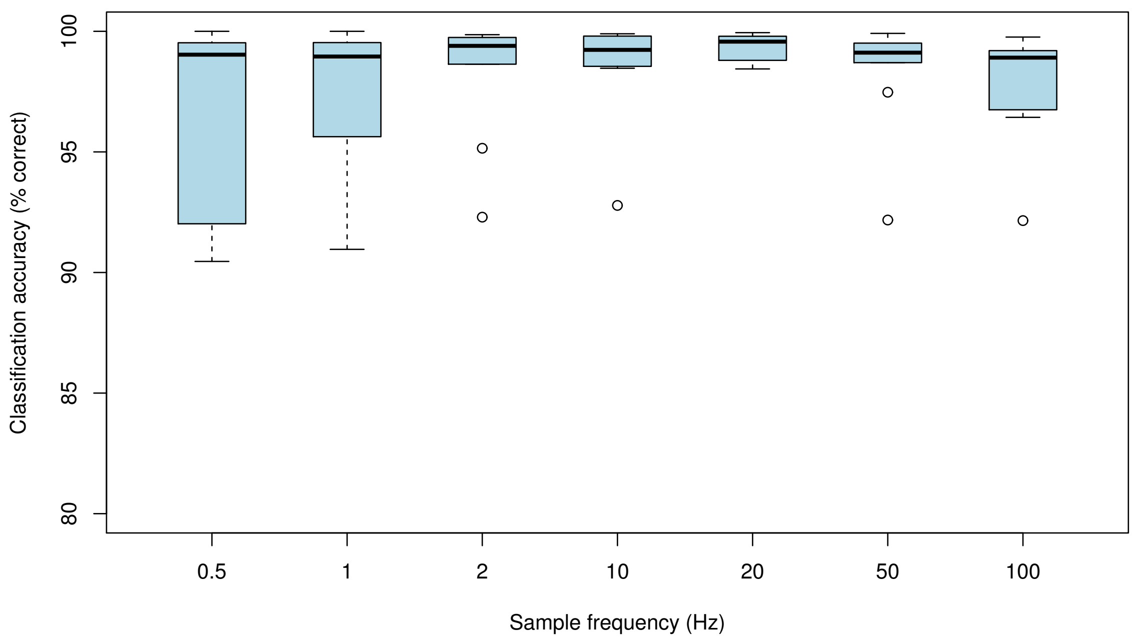

Figure 6.7: Effect of different sampling rates when classifying crawling.

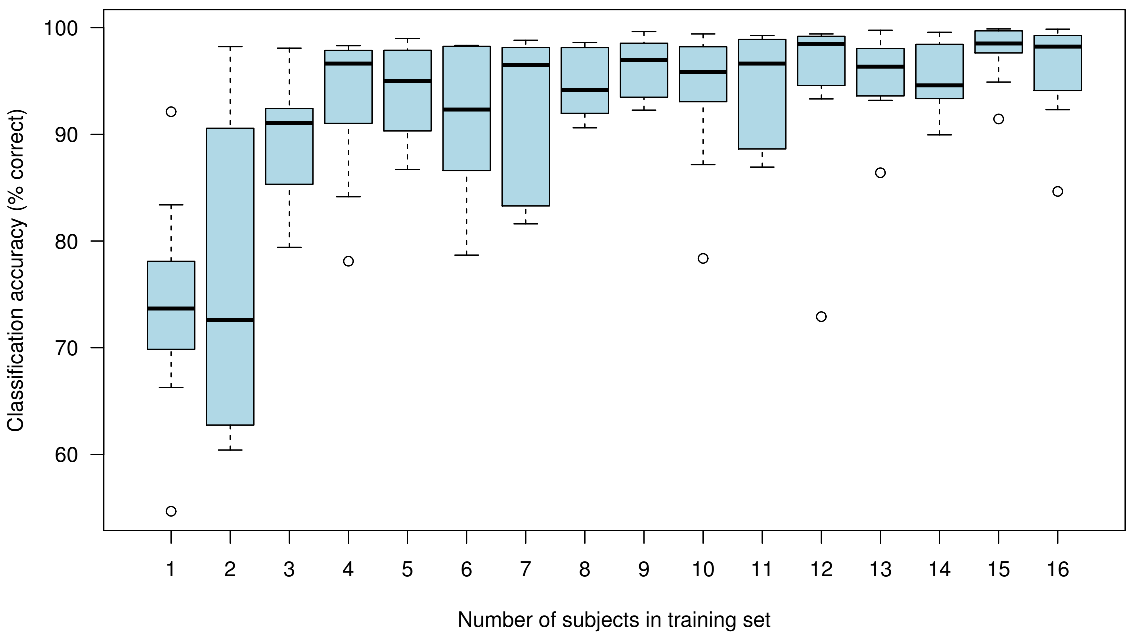

Figure 6.8: Classification accuracy when using different numbers of subjects in the training set.

insensitive to sampling rate, with similar median accuracies obtained at each of the rates tested. However, for walking, 10 Hz appears to provide the most consistent results and for crawling the same can be seen for 2 Hz to 50 Hz (represented by the smaller inter-quartile ranges in each case). The use of a 10 Hz sampling rate is thus justified for the postures considered in the work here. It is expected that other applications with different sets of postures will see similar results, though in an application requiring classification of both walking and running a higher sampling rate may be required.

6.8 Selection of training set size

Hypothesis 5 Obtaining a consistently high classification accuracy will require a certain minimum number of subjects in the training set. However, further increase in the dataset size will not show a matching increase in the classification accuracy.

The size of the training set used to train a machine learning based classifier affects its ability to generate rules that generalise well to unseen data. A larger training set thus increases the likelihood that the classifier will provide a high classification accuracy when deployed on a subject that was not involved in training. The disadvantage to having a large training set, though, is the time and expense required to obtain it. There has been little effort reported towards creating a “standard” database of posture classifier training data, and so one of the first tasks of the system implementer must almost always be to gather their own data. This section thus aims to determine the minimum appropriate size of such a training set in order to ensure high classification accuracy.

The investigation was performed using data sets D1, D2, D3, and D5 (7.5 hours of data from 17 subjects). The feature extracted from the data for this investigation was WVar over a 30 sample window.

Figure 6.8 on the preceding page shows the results obtained when training the classifier on N subjects (for $1 \leq N \leq 16$) and testing on an unseen subject. This was repeated 10 times for each N with training and test subjects randomly selected (without replacement). It can be seen that for the set of training subjects used there is no significant increase in classifier accuracy when training on more than eight subjects. The subjects available for this analysis did not cover all possible body builds as they were focused around the typical age range expected for EOD operatives (20 to 36 years old). This is to be expected in many applications however—monitoring of the daily activities of elderly people is a clear example where the range of subjects is limited. In applications that are expected to cover a wide range of people of all ages, heights and weights, it is likely that the minimum required number of subjects for the training set will be larger.

6.9 Classification of individual postures

Hypothesis 6 The posture classifier can classify specific targeted postures with a similar accuracy to the results found when classifying all postures.

The posture classification method developed here is intended to be generic and support classification of any defined group of postures with a high accuracy, including the case of classifying single individual postures. A system that is able to classify one specific posture or a small group of postures would be beneficial in monitoring a subject undergoing an at-home physical rehabilitation regime. For example, a system for monitoring an elderly person may only require classification of lying rather than classification of all of their daily activities. The system would, however, require the ability to distinguish the targeted postures from the remainder of the subject's activities.

An investigation was performed, targeting the case of classifying specific individual postures, using untruncated data collected from data set D5 (2.5 hours of data from 7 subjects). WVar was used as the extracted feature over a 30 sample window.

The classification performance for each posture in the set of eight used in this work was analysed in two ways:

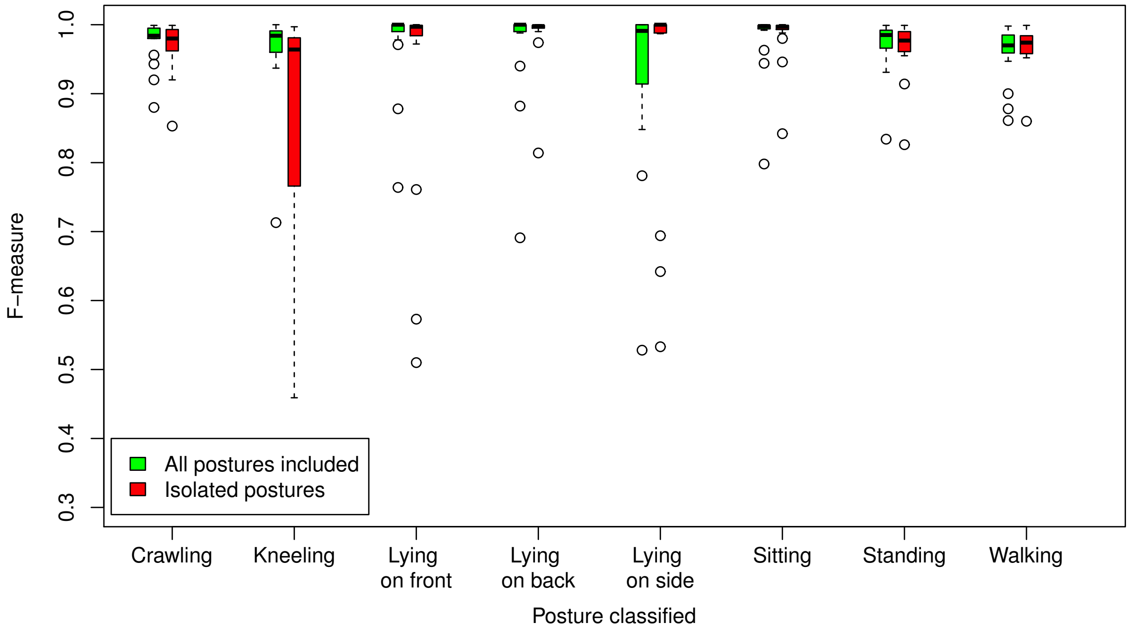

Figure 6.9: Classification F-measure when targeting specific postures.

1. the data was annotated as per the rest of the thesis using the annotation method described in Section 4.8 on page 69, and,

2. a separate set of annotations was performed that took each posture in turn and marked the rest of the data with “other”.

For annotation method 1, LOSOXV was used to find the F-measure for the targeted posture. For annotation method 2, LOSOXV was performed for each iteration (that is, for each posture being individually classified). F-measure was used in preference to accuracy as the classes were not equally represented (i.e. the data contained many more instances of the “other” class than instances of the targeted posture).

Figure 6.9 shows the classification F-measure for each posture using the two methods. The F-measure was similar for each posture when classified as one of the set of eight and when targeted individually except in the case of kneeling and (to a lesser extent) lying on one side. This demonstrates the capability of decision trees to distinguish a particular class from a set of aggregated classes at least as well as they distinguish it from the set of individual classes in most cases.

In conclusion, if the intention is to only classify a specific posture or small group of postures then the remaining postures can be simply annotated in a generic “other” class. This can greatly simplify the data gathering process as large amounts of data for the other postures can be collected without annotation being a concern.

6.10 Summary and design space definition

When building a machine learning based posture classification system there are a large number of options available that can have a significant impact on the accuracy of the classifier. Despite this, there appears to be little published investigation into the impact of the factors and little commonality between systems reported in the literature (as shown in Section 2.4.3 on page 32). This chapter presented an investigation of the design space for C4.5 classification trees and an evaluation of the impact of system configuration and training process factors on the accuracy of the posture/activity classifier. The design space is explored such that the results and discussion provide guidance for posture classification system design in a range of applications with differing constraints. From these results, the optimal configuration for the case study application is extracted. The factors considered were: extracted data feature choice, feature window size, number of sensors, training set size, sampling rate, and targeting of individual postures.

For the set of postures defined in this thesis the extracted data feature providing the highest overall accuracy was WVar. For this set of postures the main difficulty appeared to be distinguishing between similar static and dynamic postures, particularly standing and walking. The worst classification accuracy for each of the posture was generally produced by either Mean or EWMA.

The window size selected for feature extraction influences the classification results in two ways: introducing delays if the window is too large or not providing sufficient historical data if the window is too small. A window size of between 5 and 75 samples (0.5 to 7.5 seconds) appears to provide relatively consistent classification accuracy results when using WVar over data collected at 10 Hz. Window sizes above 75 samples show greater variation in the results along with a decreasing median accuracy as the window size increases. A period of incorrect classification is introduced following a posture transition based on the window size, though the length of this period does not increase rapidly.

Classifier accuracy is not significantly sensitive to the number of sensors used, towards a minimal configuration of 2 sensors (left thigh and calf). The accuracy degrades significantly, however, if only a single sensor is used—the mean accuracy decreases by 15.9% for one sensor on the left calf (the best performing individual sensor) compared to two sensors on right calf and thigh. The optimal position for a single sensor is the calf, with the arm sensors performing worst for this posture set. Sitting, standing and kneeling were not classified with an accuracy higher than 95% using any single sensor. For the set of postures monitored in this work, two correctly placed sensors provide classification at a high accuracy.

The sampling used will effect classification accuracy for dynamic postures. Walking is best classified at 10 Hz while crawling is best classified between 2 Hz and 50 Hz.

Classification accuracy was found to be highest when training on at least eight subjects, though little benefit was found beyond this number. The subjects selected were based on the EOD operative monitoring application and so did not cover all ages and body builds (other applications are expected to also have specific “target” populations in this way).

The ability to monitor individual postures is a possibility. It was found that, using the data gathering method and the classification algorithm described in this chapter, individual postures could be classified with a similar performance to when classified as part of the full set of eight.

In conclusion, the parameters selected as giving the highest accuracy for the case study scenario were: WVar as the extracted feature, a window size of 5 to 75 samples, a minimum of two sensors (on the thigh and calf), and 10 Hz sampling.

The next chapter concludes on the work presented in this thesis and summarises the answers to the research questions posed in Chapter 1.

Chapter 7

Conclusions and further work

This thesis investigated the feasibility of delivering high accuracy posture classification in real-time, through an autonomous, on-body, wireless end-to-end system. The work was concerned with several key aspects of this research domain:

1. Concepts, designs and implementations for successful deployable systems meeting generic and application-based requirements.

2. Solutions for handling transitions between postures in order to ensure high classification accuracy in the context of natural movement.

3. The definition and analysis of the design space for tree-based posture classifiers and identification of optimal design parameters.

Chapter 3 described a platform design for real-time on-body posture classification systems, including a set of generic requirements applicable to all applications. The platform is targeted at two basic usage scenarios: 1) a self-contained system providing real-time postural information to an external system and 2) an investigative instrument focused on remote reconfigurability. Two example applications were described, following from the usage scenarios, and a prototype posture classification system was implemented based on the platform design. This system was used in the evaluation of the algorithms developed during the work. One of the applications considered in the prototype system implementation—EOD operative monitoring—guided the development of the data gathering regimes used in Chapter 4.

Chapter 4 presented the chosen posture classification algorithm and detailed the method used to gather and process the data sets required for the investigations in this thesis. Given the application requirement for real-time classification using potentially resource constrained on-body nodes, a natural choice of posture classification algorithm was decision trees. The algorithm chosen was C4.5, and the WEKA toolkit was used to generate the classifiers using this algorithm. The trained classifiers provided an average accuracy of 96.3% when using WVar as a data feature. The results showed that the full set of required postures can be accurately classified if an extracted data feature such as WVar is used, and that a classifier trained on a subject wearing light clothing is still suitable when used with a bulkier suit, such as those used by EOD operatives.

The methods of gathering and handling data (further used for all of the analysis work in the thesis) were standardised and automated around a series of regimes guided by existing research in the area of EOD operative safety and on feedback from an EOD suit manufacturer. These regimes focused on: 1) the eight defined postures alone, 2) the eight postures combined with natural movement, and 3) mission-like activity. Twenty two experimental subjects were employed to provide sufficient coverage on sex, age, height and weight. A total of 22.7 hours of data was collected for use in training and testing the decision trees. Decision tree testing was performed via two methods: LOSOXV and direct comparison of a trained tree's output against the correct posture when tested on unseen data.

The end-to-end system resulting from the work was further evaluated when deployed and classifying posture in real-time for five subjects performing tasked activities in a realistic scenario. The classification accuracy was shown to be consistent with the results found during offline testing (between 94.2% and 98.5% for the five subjects), and the information yield was sufficient for the applications considered (above 99% in all cases).

Chapter 5 presented an investigation into transitions and described the method of handling them to reduce their negative impact on both accuracy and the number of posture change events generated. Three transition smoothing filters (a simple voting filter, an EWV filter, and a Bayes filter) were proposed and evaluated as a means of achieving two main goals:

1. an improvement of the accuracy of classification over data containing both stable postures and transitions, and,

2. a reduced number of generated posture updates, providing benefits in battery life and required network bandwidth as well as providing more stable input to an external subsystem.

.

All filters showed an increase in accuracy and a significant reduction in the number of posture change events compared to the unfiltered data. The EWV filter with $\alpha = 0.04$ had the best overall performance of the three post-processing filters with an overall increase of classification accuracy of up to 1.1% (when compared with classifier accuracy during natural movement regimes) and a 99.6% reduction in postural events generated (when compared with continuously reporting postural systems). The filter's performance was in-line with prediction based on offline testing, when deployed in realistic scenarios with high frequency of transitions within the tasked activities.

Chapter 6 presented an in-depth investigation into the design space of a supervised machine learning based posture classifier, with a focus on C4.5 decision trees. When building such a classification system there are a number of design parameters that will impact on the accuracy of the classifier. The factors identified and further analysed were: extracted data feature choice, feature window size, number of sensors and their potion, training set size, sampling rate, and targeting of individual postures.

The results showed that:

- WVar provides the highest overall accuracy for the set of postures considered here, and with the worst performing features being Mean and EWMA.
- A window size of between 0.5 to 7.5 seconds results in the best classification accuracy when using WVar.
- Acceptable classifier accuracy can be achieved with 2 sensors (on the right thigh and calf). Should only a single sensor be available, the optimal position is the calf for this posture set (with the arm sensors performing worst); however, a considerably lower classification accuracy results.
- Systems sampling at 10 Hz will deliver the best accuracy for walking. For crawling, the range is between 2 Hz and 50 Hz. Overall, 10 Hz provides the best classification accuracy for the set of 8 postures.
- There was no significant increase in classifier accuracy when training on more than eight subjects.

7.1 Research questions

The research questions posed by this thesis were as follows:

1. Can the defined set of postures, namely sitting, kneeling, crawling, standing, walking, and lying on front, back and one side, be accurately classified in real-time using an on-body wearable sensor-based system?
2. How do transitions between postures affect classifier accuracy, and can any negative impact be reduced or eliminated?
3. What is the design space for a posture classifier targeted at specific application requirements?

The following subsections present the answers to these questions as found in the relevant chapters.

7.1.1 Can the defined set of postures, namely sitting, kneeling, crawling, standing, walking, and lying on front, back and one side, be accurately classified in real-time using an on-body wearable sensor-based system?

Chapter 4 demonstrated that a C4.5 decision tree based posture classifier can provide a real-time, on-body, autonomous classification accuracy of 96.3% for the full set of eight postures considered here when WVar is used as a data feature. Realistic evaluation of the system, deployed on five subjects performing tasked activities, showed accuracies of between 94.2% and 98.5% (reducing to between 92.9% and 97.9% when transitions are included and transition filters are used). A data gathering process to support the training and testing of posture classifiers was detailed. Chapter 3 presented a design for a platform supporting real-time on-body classification of posture.

7.1.2 How do transitions between postures affect classifier accuracy, and can any negative impact be reduced or eliminated?

When classifiers trained on well defined, truncated posture data were tested on data representing natural, continuous movement, classifier accuracy was observed to decrease. Chapter 5 showed that a consistent degradation in the classifier accuracy should be expected, linked to the frequency of transitions: 2% accuracy reduction for each transition/minute.

Posture smoothing filters were presented and evaluated as a solution to this problem, with the EWV filter (voting with exponential weighting in favour of recent postures) delivering the best results. Using this filter the degradation in accuracy when including transitions in the testing data was less than 1% per transition/minute, with posture change events reduced by around 75.2%. Real-time evaluation with the deployed system confirmed these results.

7.1.3 What is the design space for a posture classifier targeted at specific application requirements?

Chapter 6 defined and assessed the following design space parameters (provided here together with their optimal values/ranges):

- Extracted data feature: WVar.
- Window size: between 5 and 75 samples (at 10 Hz sampling rate).
- Number of sensors: at least 2.

- Sensor positions: calf and thigh.
- Sampling rate: 10 Hz.
- Number of subjects for training: at least 8.

Additional results were also provided to aid in designing classifiers for other applications requiring similar postures or a subset of the eight postures considered here.

It is thus considered that all research questions have been answered by the author.

7.2 Further work

There are several avenues of further work that would benefit from investigation based on the results presented in this thesis.

Investigation of a larger number of postures

The analysis throughout this thesis was based on classification of eight specific postures: sitting, kneeling, crawling, standing, walking, and lying on front, back and one side. It would be beneficial to researchers pursuing other applications if the analysis were extended to include additional postures or activities. Furthermore, specific activities or variations on the basic postures could be analysed separately. An example of this would be if an application required classification of several broad movement speeds (such as standing, walking, jogging, running, and sprinting).

Transition handling mechanism

The current approach to handling transitions is to define the postures prior to and following the transition to be “correct” and to attempt to maintain a stable output consisting of these two postures. There are, however, instances where additional postures should be considered as valid. For example, a subject transitioning from sitting to lying on their front is likely to kneel during the transition. Based on this, the output of the classifier during a transition could provide additional information regarding the final posture. To continue the previous example, if the subject was sitting and then enters a transitional period involving kneeling, then the transition filter could assign a higher weight to classifications of lying on their front soon after. This could increase classification accuracy and output stability by biasing the output towards the more likely postures. This is an approach that would benefit from the application of a Bayesian Network or Markov Model approach, where the links between particular postures can be modelled explicitly. Laerhoven and Cakmakci [73], for example, used this principle to model the links between activities on a broader scale.

Bibliography

[1] Mohammed Aloqlah, Rosa R. Lahiji, Kenneth A. Loparo, and Mehran Mehregany. A headband for classifying human postures. In *Proceedings of the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pages 382–385, Buenos Aires, Argentina, 31 August–4 September 2010.

[2] Oliver Amft. On the need for quality standards in activity recognition using ubiquitous sensors. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[3] K. Aminian, Ph. Robert, E. E. Buchser, B. Rutschmann, D. Hayoz, and M. Depairon. Physical activity monitoring based on accelerometry: validation and comparison with video observation. *Medical And Biological Engineering and Computing*, 37(3):304–308, May 1999.

[4] Javier Andreu and Plamen Angelov. Real-time human activity recognition from wireless sensors using evolving fuzzy systems. In *Proceedings of the IEEE World Congress on Computational Intelligence (WCCI'10)*, pages 1–8, Barcelona, Spain, 18–23 July 2010. IEEE Computer Society. ISSN: 1098-7584.

[5] E. K. Antonsson and R. W. Mann. The frequency content of gait. *Journal of biomechanics*, 18(1):39–47, 1985.

[6] D. K. Arvind, Martin J. Ling, Paul McEwan, and Alex D. Young. On-body wireless motion capture system for real-time animation. In *SIKTN Workshop on Wireless Sensing*, National Physical Laboratory, Teddington, UK, July 2008.

[7] H. Harry Asada, Phillip A. Shaltis, Andrew Reisner, Sokwoo Rhee, and Reginald C. Hutchinson. Mobile monitoring with wearable photoplethysmographic biosensors. *IEEE Engineering in Medicine and Biology Magazine*, 22(3):28–40, May-June 2003.

[8] Louis Atallah, Benny Lo, Rachel King, and Guang-Zhong Yang. Sensor placement for activity detection using wearable accelerometers. In *Proceedings of the 2010 International Conference on Body Sensor Networks (BSN '10)*, pages 24–29, Biopolis, Singapore, 7–9 June 2010. IEEE Computer Society, Washington, DC, USA.

[9] Louis Atallah, Benny Lo, Rachel King, and Guang-Zhong Yang. Sensor positioning for activity recognition using wearable accelerometers. *IEEE Transactions on Biomedical Circuits and Systems*, 5(4):320–329, August 2011.

[10] Ryan Aylward and Joseph A. Paradiso. A compact, high-speed, wearable sensor network for biomotion capture and interactive media. In *Proceedings of the 6th international conference on Information processing in sensor networks (IPSN '07)*, pages 380–389, Cambridge, Massachusetts, USA., 25–27 April 2007. ACM.

[11] Pierre Baldi, S/oren Brunak, Yves Chauvin, Claus A. F. Andersen, and Henrik Nielsen. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, 16(5):412–424, 2000.

[12] Stacy J. Morris Bamberg, Ari Y. Benbasat, Donna M. Scarborough, David E. Krebs, and Joseph A. Paradiso. Gait analysis using a shoe-integrated wireless sensor system. *IEEE Transactions on Information Technology in Biomedicine*, 12(4):413–423, July 2008.

[13] Danny Banks. *Microengineering, MEMS, and Interfacing: A Practical Guide*. CRC Press, 23 March 2006.

[14] Ling Bao and Stephen S. Intille. Activity recognition from user-annotated acceleration data. In *Proceedings of the Second International Conference on Pervasive Computing (PERVASIVE 2004)*, pages 1–17, Linz/Vienna, Austria, April 2004. Berlin, Heidelberg: Springer-Verlag.

[15] Pierre Barralon, Norbert Noury, and Nicolas Vuillerme. Classification of daily physical activities from a single kinematic sensor. In *Proceedings of the 27th IEEE Annual International Conference of the Engineering in Medicine and Biology Society (EMBS)*, volume 3, pages 2447–2450, Shanghai, China, 17–18 January 2005.

[16] Stephen J. Bellis, Kieran Delaney, Brendan O'Flynn, John Barton, Kafil M. Razeeb, and Cian O'Mathuna. Development of field programmable modular wireless sensor network nodes for ambient systems. *Computer Communications*, 28:1531–1544, 2005.

[17] Nagendra B. Bharatula, Stijn Ossevoort, Mathias Stäger, and Gerhard Tröster. Towards wearable autonomous microsystems. In Alois Ferscha and Friedemann Mattern, editors, *Pervasive Computing*, volume 3001 of *Lecture Notes in Computer Science*, pages 225–237. Springer Berlin / Heidelberg, 2004.

[18] Nagendra B. Bharatula, Mathias Stäger, Paul Lukowicz, and Gerhard Tröster. Empirical study of design choices in multi-sensor context recognition systems. In *Proceedings of the 2nd International Forum on Applied Wearable Computing (IFAWC)*, pages 79–93, Zürich, Switzerland, 17–18 March 2005.

[19] Nicola Bicocchi, Marco Mamei, and Franco Zambonelli. Detecting activities from body-worn accelerometers via instance-based algorithms. *Pervasive Mobile Computing*, 6(4):482–495, August 2010.

[20] Niranjan Bidargaddi, Antti Sarela, Lasse Klingbeil, and Mohanraj Karunanithi. Detecting walking activity in cardiac rehabilitation by using accelerometer. In *Proceedings of the 3rd International Conference on Intelligent Sensors, Sensor Networks and Information (ISSNIP)*, pages 555–560, Melbourne, Australia, 3–6 December 2007.

[21] Gerald Bieber and Christian Peter. Using physical activity for user behavior analysis. In *Proceedings of the 1st ACM International Conference on PErvasive Technologies Related to Assistive Environments (PETRA'08)*, pages 215–221, Athens, Greece, 2008. ACM.

[22] Subir Biswas and Muhannad Quwaider. Body posture identification using hidden markov model with wearable sensor networks. In *Proceedings of the 3rd International Conference on Body Area Networks (BodyNets'08)*, pages 1–8, Tempe, Arizona, 13–15 March 2008. Brussels: ICST Belgium.

[23] Subir Biswas and Muhannad Quwaider. Remote monitoring of soldier safety through body posture identification using wearable sensor networks. *Wireless Sensing and Processing III*, 6980, 03 April 2008.

[24] Ulf Blanke, Diane Larlus, Kristof Van Laerhoven, and Bernt Schiele. Standing on the shoulders of other researchers - a position statement. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[25] Cornelis P. Bogerd, Kurt Strässle, and Paul A. Brühwiler. The effects of wearing a full-face motorcycle helmet on cognitive performance in a warm environment. In John W. Castellani and Thomas L. Endrusick, editors, *Proceedings of the 13th International Conference on Environmental Ergonomics*, pages 553–557, Boston, USA, 3–7 August 2009.

[26] Johann Borenstein. Personal dead-reckoning (pdr) system for firefighters. In *Precision Indoor Personnel Location & Tracking International Technology Workshop*, WPI (Worcester Polytechnic Institute), Worcester, Massachusetts, 2011.

[27] Stephan Bosch, Mihai Marin-Perianu, Raluca Marin-Perianu, Paul Havinga, and Hermie Hermens. Keep on moving! activity monitoring and stimulation using wireless sensor networks. In *Proceedings of the 4th European conference on Smart sensing and context (EuroSSC'09)*, pages 11–23, Guildford, UK, 16–18 September 2009. Berlin, Heidelberg: Springer-Verlag.

[28] Carlijn V. C. Bouten, Karel T. M. Koekkoek, Maarten Verduin, Rens Kodde, , and Jan D. Janssen. A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity. *IEEE Transactions on Biomedical Engineering*, 44(3):136–147, March 1997.

[29] A.J. Bernheim Brush, John Krumm, and James Scott. Activity recognition research: The good, the bad, and the future. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[30] C4.5. C4.5 decision tree. `http://en.wikipedia.org/wiki/C4.5_algorithm`: Accessed 11-April-2012.

[31] J. Solá i Carós, O. Chételat, P. Celka, and S. Dasen. Very low complexity algorithm for ambulatory activity classification. In *Proceedings of the 3rd European Medical and Biological Engineering Conference and IFMBE European Conference on Biomedical Engineering (EMBEC'05)*, volume 11, Prague, Czech Republic, 20–25 November 2005. IFMBE.

[32] Yen-Ping Chen, Jhun-Ying Yanga, Shun-Nan Lioub, Gwo-Yun Leeb, and Jeen-Shing Wanga. Online classifier construction algorithm for human activity detection using a tri-axial accelerometer. *Special Issue on Advanced Intelligent Computing Theory and Methodology in Applied Mathematics and Computation*, 205(2):849–860, 15 November 2008.

[33] Stephen S. Cheung and Tom M. McLellan. Heat acclimation, aerobic fitness, and hydration effects on tolerance during uncompensable heat stress. *Journal of Applied Physiology*, 84(5):1731–1739, May 1998.

[34] Driss Choujaa and Naranker Dulay. Activity recognition using mobile phones: Achievements, challenges and recommendations. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[35] Davide Curone, Gian Mario Bertolotti, Andrea Cristiani, Emanuele Lindo Secco, and Giovanni Magenes. A Real-Time and Self-Calibrating Algorithm Based on Triaxial Accelerometer Signals for the Detection of Human Posture and Activity. *IEEE Transactions on Information Technology in Biomedicine,*, 14(4):1098–1105, July 2010.

[36] David DeStefano. On the Line: Crawl Before You Walk. `http://www.fireengineering.com/articles/2013/02/on-the-line--crawl-before-you-walk.html`: Accessed 05-June-2013.

[37] Thomas Degen, Heinz Jaeckel, Michael Rufer, and Stefan Wyss. Speedy:a fall detector in a wrist watch. In *Proceedings of the Seventh IEEE International Symposium on Wearable Computers*, pages 184–187, 18–21 October 2003.

[38] Lance Doherty and Dana A. Teasdale. Towards 100 In *Proceedings of the 3rd ACM international workshop on Performance evaluation of wireless ad hoc, sensor and ubiquitous networks*, PE-WASUN '06, pages 132–135, New York, NY, USA, 2006. ACM.

[39] Liang Dong, Jiankang Wu, and Xiang Chen. Real-time physical activity monitoring by data fusion in body sensor networks. In *Proceedings of the 10th International Conference on Information Fusion*, pages 1–7, Quebec, Canada, 9–12 July 2007.

[40] Elena Doran. *Power consumption and lifetime prediction in sensor network devices.* PhD thesis, Tehnical University of Cluj Napoca, Romania, 2011.

[41] Lucy E. Dunne. *The Design of Wearable Technology: Addressing the Human-Device Interface Through Functional Apparel Design.* PhD thesis, Faculty of the Graduate School of Cornell University, August 2004.

[42] Miikka Ermes, Juha Pärkkä, Jani Mäntyjärv, and Ilkk Korhonen. Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions. *IEEE Transactions on Information Technology in Biomedicine*, 12(1):20–26, 7 January 2008. ISSN: 1089-7771.

[43] Elisabetta Farella, Augusto Pieracci, Davide Brunelli, and Luca Benini. A wireless body area sensor network for posture detection. In *Proceedings of the 11th IEEE Symposium on Computers and Communications (ISCC'06)*, pages 454–459, Pula-Cagliari, Sardinia, Italy, 26–29 June 2006. Los Alamitos, CA: IEEE Computer Society Press. ISBN:0-7695-2588-1.

[44] Elisabetta Farella, Augusto Pieracci, Davide Brunelli, Luca Benini, Bruno Riccó, and Andrea Acquaviva. Design and Implementation of WiMoCA Node for a Body Area Wireless Sensor Network.

In *Proceedings of the 2005 Systems Communications (ICW '05)*, pages 342–347, Montreal, Canada, 14–17 August 2005. Los Alomitas, CA: IEEE Computer Society Press.

[45] Anthony Fleury, Norbert Noury, and Michel Vacher. A wavelet-based pattern recognition algorithm to classify postural transitions in humans. In *Proceedings of the 17th European Signal Processing Conference (EUSIPCO 2009)*, pages 2047–2051, Glasgow, Scotland, 24–28 August 2009.

[46] Korbinian Frank, María Josefa Vera Nadales, Robertson Patrick, and Michael Angermann. Reliable real-time recognition of motion related human activities using mems inertial sensors. In *Proceedings of the ION GNSS 2010*, Portland, Oregon, USA, 21–24 September 2010.

[47] Francine Gemperle, Chris Kasabach, John Stivoric, Malcolm Bauer, and Richard Martin. Design for wearability. In *Proceedings of the 2nd IEEE International Symposium on Wearable Computers (ISWC '98)*, pages 116–121, Washington, DC, USA, 1998. IEEE Computer Society.

[48] A Godfrey, A.K Bourke, G.M Alaighin, P Van de Ven, and J Nelson. Activity classification using a single chest mounted tri-axial accelerometer.(report). *Medical Engineering and Physics*, 33(9):1127–1129, May 2011.

[49] Eric Guenterberg, Hassan Ghasemzadeh, and Roozbeh Jafari. A Distributed Hidden Markov Model for Fine-grained Annotation in Body Sensor Networks. In *Proceedings of the Sixth International Workshop on Wearable and Implantable Body Sensor Networks (BSN '09)*, pages 339–344, Washington, DC, USA, 3–5 June 2009. IEEE Computer Society.

[50] Eric Guenterberg, Allen A. Yang, Hassan Ghasemzadeh, Roozbeh Jafari, Ruzena Bajcsy, and S. Shankar Sastry. A Method for Extracting Temporal Parameters Based on Hidden Markov Models in Body Sensor Networks With Inertial Sensors. *Information Technology in Biomedicine*, 13(6):1019–1030, November 2009. ISSN: 1089-7771.

[51] Gumstix, Inc. Gumstix Motherboard IO. `http://www.gumstix.com/`; Accessed 11-April-2012.

[52] Zhenyu He and Lianwen Jin. Activity recognition from acceleration data using AR model representation and SVM. In *Proceedings of the Seventh International Conference on Machine Learning and Cybernetics*, pages 2245–2250, Kunming, China, 12–15 July 2008.

[53] Zhenyu He and Lianwen Jin. Activity Recognition from acceleration data Based on Discrete Consine Transform and SVM. In *Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics*, pages 5041–5044, San Antonio, TX, USA, 11–14 October 2009.

[54] Ernst A. Heinz, Kai Kunze, Matthias Gruber, David Bannach, and Paul Lukowicz. Using Wearable Sensors for Real-time Recognition Tasks in Games of Martial Arts - An Initial Experiment. In *Proceedings of the 2nd IEEE Symposium on Computational Intelligence and Games (CIG)*, pages 98–102, Reno/Lake Taho, Nevada, USA, 23–24 May 2006.

[55] Weiyao Huang, Jun Zhang, and Zhijing Liu. Activity recognition based on Hidden Markov Models. In *Proceedings of the 2nd international conference on Knowledge science, engineering and management (KSEM'07)*, pages 532–537, Melbourne, Australia, 28–30 November 2007. Berlin, Heidelberg: Springer-Verlag.

[56] Tâm Huýnh, Ulf Blanke, and Bernt Schiele. Scalable recognition of daily activities with wearable sensors. In *Proceedings of the 3rd international conference on Location-and context-awareness (LoCA'07)*, pages 50–67, Oberpfaffenhofen, Germany, 20–21 September 2007. Berlin, Heidelberg: Springer-Verlag.

[57] Tâm Huýnh and Bernt Schiele. Analyzing features for activity recognition. In *Proceedings of the 2005 joint conference on Smart Objects and Ambient Intelligence (sOc-EUSAI '05)*, pages 159–163, Grenoble, France, 12–14 October 2005. New York, NY, USA: ACM.

[58] Tâm Huýnh and Bernt Schiele. Towards Less Supervision in Activity Recognition from Wearable Sensors. In *Proceedings of the 10th IEEE International Symposium on Wearable Computers*, pages 3–10, Montreux, Switzerland, 11–14 October 2006. ISSN: 1550-4816.

[59] Nuri Firat Ince, Cheol-Hong Min, Ahmed Tewfik, and David Vanderpool. Detection of early morning daily activities with static home and wearable wireless sensors. *EURASIP Journal on Advances in Signal Processing*, 2008(31):11, January 2008. ISSN: 1110-8657.

[60] Roozbeh Jafari, Wenchao Li, Ruzena Bajcsy, Steven Glaser, and Shankar Sastry. Physical Activity Monitoring for Assisted Living at Home. In *Proceedings of the 4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, volume 13, pages 213–219, RWTH University, Aachen Germany, 26–28 March 2007.

[61] Ah-young Jeon, In-cheol Kim, Jae-hee Jung, Soo-young Ye, Jae-hyung Kim, Ki-gon Nam, Seoung-wan Baik, Jung-hoon Ro, and Gye-rok Jeon. Implementation of the personal emergency response system. *International Journal of Biological and Medical Sciences*, 4(1):61–65, 2009.

[62] Shanshan Jiang, Yanchuan Cao, Sameer Iyengar, Philip Kuryloski, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, and Stephen Wicker. CareNet: an integrated wireless sensor networking environment for remote healthcare. In *Proceedings of the 3rd ICST International Conference on Body Area Networks (BodyNets '08)*, volume 9, pages 1–3, Tempe, Arizona, 13–15 March 2008. Brussels, Belgium: ICST.

[63] Emil Jovanov, Aleksandar Milenkovic, Chris Otto, and Piet de Groen. A wireless body area network of intelligent motion sensors for computer assisted physical rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2:6, 1 March 2005.

[64] Dong Won Kang, Jin Seung Choi, Jeong Whan Lee, Soon Cheol Chung, Soo Jun Park, and Gye Rae Tack. Real-time elderly activity monitoring system based on a tri-axial accelerometer. *Disability and Rehabilitation: Assistive Technology*, 5(4):247–253, July 2010.

[65] Dean M. Karantonis, M. Rajaram Narayanan, Merryn Mathie, Nigel H. Lovell, and Branko G. Celler. Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring. *IEEE Transactions on Information Technology in Biomedicine*, 10(1):156–167, January 2006.

[66] John Kemp. *Body Sensor Networks for Health Monitoring: A Safety Critical Mission Application*. PhD thesis, Coventry University, 2010.

[67] John Kemp. *Body Sensor Networks for Health Monitoring: A Safety Critical Mission Application*, chapter 2 Literature Review, pages 13–49. PhD thesis, Coventry University, 2010.

[68] Adil Mehmood Khan, Young-Koo Lee, Sungyoung Y. Lee, and Tae-Seong Kim. A triaxial accelerometer-based physical-activity recognition via augmented-signal features and a hierarchical recognizer. *IEEE Transactions on Information Technology in Biomedicine*, 14(5):1166–1172, September 2010.

[69] A.M Khan, Y.K Lee, and S. Y. Lee. Accelerometer's position free human activity recognition using a hierarchical recognition model. In *Proceedings of the 12th IEEE International Conference on e-Health Networking Applications and Services*, Healthcom, pages 296–301, Lyon, 1–3 July 2010.

[70] James F. Knight, Chris Baber, Anthony Schwirtz, and Huw W. Bristow. The comfort assessment of wearable computers. In *In Proceedings of the International Symposium of Wearable Computers (IEEE ISWC)*, pages 65–72, Seattle, WA, USA, 7–10 October 2002. IEEE Computer Society.

[71] Narayanan C. Krishnan, Dirk Colbry, Colin Juillard, and Sethuraman Panchanathan. Real Time Human Activity Recognition Using Tri-Axial Accelerometers. In *Proceedings of the Sensors Signals and Information Processing Workshop (SenSIP 2008)*, pages 3–6, Sedona, AZ, 11–14 May 2008.

[72] Martin Kusserow, Oliver Amft, and Gerhard Tröster. BodyANT: Miniature wireless sensors for naturalistic monitoring of daily activity. In *Proceedings of the 4th International Conference on Body Area Networks (BodyNets '09)*, pages 1–8, Los Angeles, California, 1–3 April 2009. New York: ACM Press. ISBN: 978-963-9799-41-7.

[73] Kristof Van Laerhoven and Ozan Cakmakci. What shall we teach our pants? In *Proceeding of the Fourth International Symposium on Wearable Computers*, pages 77–83, Atlanta, Georgia, 16–17 October 2000.

[74] Kristof Van Laerhoven, David Kilian, and Bernt Schiele. Using rhythm awareness in long-term activity recognition. In *Proceedings of the 12th International Symposium on Wearable Computers (ISWC2008)*, pages 63–66, Washington, DC, USA, 28 September–1 October 2008. Los Alamitos, CA: IEEE Computer Society Press.

[75] Patricia A. LeDuc, Matthew J. Reardon, James L. Persson, Siobhan M. Gallagher, Shannen L. Dunkin, Daniel R. Fuller, and Gerald E. Lucier. Heat stress evaluation of Air Warrior Block I MOPP0 and MOPP4 ensembles with and without microclimate cooling. Technical Report USAARL Report No. 2002-19, U.S. Army Aeromedical Research Laboratory, August 2002.

[76] Qiang Li, John A. Stankovic, Mark A. Hanson, Adam T. Barth, John Lach, and Gang Zhou. Accurate, Fast Fall Detection Using Gyroscopes and Accelerometer-Derived Posture Information. In *Proceedings of the Sixth International Workshop on Wearable and Implantable Body Sensor Networks (BSN '09)*, pages 138–143, Berkeley, California, USA, 3–5 June 2009. Washington, DC, USA: IEEE Computer Society.

[77] LifeShirt system. `http://vivonoetics.com/products/sensors/lifeshirt/`: Accessed 04-December-2012.

[78] Clemens Lombriser, Nagendra B. Bharatula, Daniel Roggen, and Gerhard Tröster. On-body activity recognition in a dynamic sensor network. In *Proceedings of the ICST 2nd international conference on Body area networks (BodyNets '07)*, pages 1–6, Florence, Italy, 11–13 June 2007. Brussels, Belgium: ICST.

[79] Xi Long, Bin Yin, and Ronald M. Aarts. Single-accelerometer-based daily physical activity classification. In *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Soc (EMBC'09)*, pages 6107–6110, Minneapolis, USA, 2–6 September 2009. ISBN: 978-1-4244-3296-7.

[80] Paul Lukowicz, Stephen Intille, and Jamie A Ward. How To Do Good Research In Activity Recognition: Experimental methodology, performance evaluation and reproducibility (Workshop in conjunction with Pervasive 2010). `http://ubicomp.lancs.ac.uk/workshops/activity2010/`.

[81] Paul Lukowicz, Jamie A. Ward, Holger Junker, Mathias Stäger, Gerhard Tröster, Amin Atrash, and Thad Starner. Recognizing workshop activity using body worn microphones and accelerometers. In *Pervasive Computing: Proceedings of the 2nd International Conference*, pages 18–22, Vienna, Austria, 21–23 April 2004. Berlin, Heidelberg: Springer-Verlag.

[82] Andreas Lymberis and André Dittmar. Advanced wearable health systems and applications - research and development efforts in the european union. *Engineering in Medicine and Biology Magazine, IEEE*, 26(3):29–33, May–June 2007.

[83] A. Lynch, B. Majeed, B. O'Flynn, J. Barton, F. Murphy, K. Delaney, and S. C. O'Mathuna. A wireless inertial measurement system (WIMS) for an interactive dance environment. *Journal of Physics: Conference Series*, 15:95–100, September 2005.

[84] M. Mathie, A. Coster, N. Lovell, and B. Celler. Detection of daily physical activities using a triaxial accelerometer. *Medical and Biological Engineering and Computing*, 41(3):296–301, May 2003. ISBN: 0140-0118.

[85] Merryn. J. Mathie and Branko G. Celler. A system for monitoring posture and physical activity using accelerometers. In *Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 4, pages 3654–3657, Istanbul, Turkey, 28 October 2001.

[86] Merryn J. Mathie, Adelle C. F. Coster, Nigel H. Lovell, and Branko G. Celler. Accelerometry: providing an integrated, practical method for long-term, ambulatory monitoring of human movement. *Physiological Measurement*, 25(2):1–20, 2004.

[87] Uwe Maurer, Asim Smailagic, Daniel P. Siewiorek, and Michael Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *Proceedings of the International*

Workshop on Wearable and Implantable Body Sensor Networks (BSN 2006), pages 113–116, Cambridge, MA, 3–5 April 2006. IEEE Computer Society. ISBN: 0-7695-2547-4.

[88] Yao Meng and Hee-Cheol Kim. Wearable systems and applications for healthcare. In *Proceedings of the First ACIS/JNU International Conference on Computers, Networks, Systems and Industrial Engineering (CNSI)*, pages 325–330, 23–25 May 2011.

[89] David L. Mills. Internet time syncronization: The network time protocol. *IEEE Transaction on Communications*, 39(10):1482–1493, 1991.

[90] David Minnen, Thad Starner, Irfan Essa, and Charles Isbell. Discovering Characteristic Actions from On-Body Sensor Data. In *Proceedings of the 10th IEEE International Symposium on Wearable Computers (ISWC)*, pages 11–18, Montreux, Switzerland, 11–14 October 2006. IEEE Computer Society.

[91] Andreas F. Molisch. *Wireless Communications*. John Wiley & Sons, 2 edition, 2012.

[92] K. Montgomery, C. Mundt, G. Thonier, A. Tellier, U. Udoh, V. Barker, R. Ricks, L. Giovangrandi, P. Davies, Y. Cagle, J. Swain, J. Hines, and G. Kovacs. Lifeguard - a personal physiological monitor for extreme environments. In *Proceedings of the 26th Annual International Conference of the IEEE EMBS*, San Francisco, CA, USA, 1–5 September 2004.

[93] K. Motoi, K. Ikeda, Y. Kuwae, T. Yuji, Y. Higashi, M. Nogawa, S. Tanaka, and K. Yamakoshi. Development of an ambulatory device for monitoring posture change and walking speed for use in rehabilitation. In *Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS '06)*, pages 5940–5943, New York, USA, 30 August–3 September 2006.

[94] M. N. Nyan, Francis E. H. Tay, M. Manimaran, and K. H. W. Seah. Garment-based detection of falls and activities of daily living using 3-axis MEMS accelerometer. *Journal of Physics: Conference Series*, 34:1059–1067, May 2006.

[95] Daniel Olguín Olguín and Alex (Sandy) Pentland. Human activity recognition: Accuracy across common locations for wearable sensors. In *Proceedings of the 10th International Symposium on Wearable Computing (Student Colloquium Proceedings)*, pages 11–13, Montreux, Switzerland, 11–14 October 2006.

[96] Julien Pansiot, Danail Stoyanov, Douglas McIlwraith, Benny Lo, and G. Yang. Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems. In *Proceedings of the 4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, volume 13, pages 208–212, RWTH Aachen University, Germany, 26–28 March 2007. Berlin, Heidelberg: Springer-Verlag. ISSN: 1680-0737.

[97] Juha Pärkkä, Miikka Ermes, Panu Korpipää, Jani Mäntyjärvi, Johannes Peltola, and Ilkka Korhonen. Activity classification using realistic data from wearable sensors. *IEEE Transactions on Information Technology in Biomedicine*, 10(1):119–128, January 2006.

[98] Shyamal Patel, Hyung Park, Paolo Bonato, Leighton Chan, and Mary Rodgers. A review of wearable sensors and systems with application in rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 9(1):21, 20 April 2012.

[99] Thomas Plötz. How to do good research in activity recognition. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[100] Benjamin Poppinga and Susanne Boll. Activity recognition research is more than finding the ultimate algorithms or parameters. In *How To Do Good Research In Activity Recognition (Workshop in conjunction with Eighth International Conference on Pervasive)*, Helsinki, Finland, 17–20 May 2010.

[101] PSM Responder system. `http://www.zephyr-technology.com/solutions/psm-responder`: Accessed 04-December-2012.

[102] John Ross Quinlan. *C4.5: Programs for machine learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 1993.

[103] Muhannad Quwaider, Anthony Plummer, Jayanthi Rao, Mahmoud Taghizade, and Subir Biswas. Real-time Posture Detection using Body Area Sensor Networks. In *Proceedings of the 13th International Symposium on Wearable Computers (ISWC 2009)*, Linz, Austria, 4–7 September 2009.

[104] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L. Littman. Activity recognition from accelerometer data. In *Proceedings of the Seventeenth Conference on Innovative Applications of Artificial Intelligence (IAAI)*, volume 3, pages 1541–1546, Pittsburgh, Pennsylvania, 9–13 July 2005. AAAI Press.

[105] N. Ruchansky, C. Lochner, E. Do, T. Rawls, N. Hajj Chehade, J. Chien, G. Pottie, and W. Kaiser. Monitoring workspace activities using accelerometers. In *Proceeding of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1837–1839, Prague, Czech Republic, 22–27 May 2011.

[106] Annapurna Sharma, Young-Dong Lee, and Wan-Young Chung. High accuracy human activity monitoring using neural network. In *Proceedings of the Third International Conference on Convergence and Hybrid Information Technology (ICCIT '08)*, pages 430–435, Busan, Korea, 11–13 November 2008.

[107] SHIMMER Research. Shimmer - wireless sensor platform for wearable applications. `http://www.shimmer-research.com/`; Accessed 10-February-2012.

[108] Pekka Siirtola, Heli Koskimäki, and Juha Röning. Periodic quick test for classifying long-term activities. In *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pages 135–140, Paris, France, 11–15 April 2011.

[109] Pekka Siirtola, P. Laurinen, E. Haapalainen, J. Roning, and H. Kinnunen. Clustering-based Activity Classification with a Wrist-worn Accelerometer Using Basic Features. In *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM '09)*, pages 95–100, Nashville, TN, USA, 30 March–2 April 2009.

[110] Bonnie G Steele, Basia Belza, Kevin Cain, Catherine Warms, Jeff Coppersmith, and Joellen Howard. Bodies in motion: Monitoring daily activity and exercise with motion sensors in people with chronic pulmonary disease. *Journal Of Rehabilitation Research And Development*, 40(5 Supplement 2):45–58, September/October 2003.

[111] Thomas Stiefmeier, Daniel Roggen, Gerhard Tröster, Georg Ogris, and Paul Lukowicz. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing*, 7(2):42–50, April-June 2008.

[112] Michael Sung, Rich DeVaul, Silvia Jimenez, Jonathan Gips, and Alex (Sandy) Pentland. Shiver motion and core body temperature classification for wearable soldier health monitoring systems. In *Proceedings of the Eighth International Symposium on Wearable Computers (ISWC '04)*, pages 192–193, Arlington, VA, USA, 31 October–3 November 2004. Washington, DC, USA: IEEE Computer Society.

[113] Jaakko Suutala, Susanna Pirttikangas, and Juha Röning. Discriminative temporal smoothing for activity recognition from wearable sensors. In *Proceedings of the 4th international conference on*

Ubiquitous computing systems (UCS'07), pages 182–195, Tokyo, Japan, 25–28 November 2007. Berlin, Heidelberg: Springer-Verlag.

[114] Shin'ichi Takeuchi, Satoshi Tamura, and Satoru Hayamizu. Human Action Recognition Using Acceleration Information Based On Hidden Markov Model. In *Proceedings of the 2009 APSIPA Annual Summit and Conference*, pages 829–832, Sapporo, Japan, 4–7 October 2009.

[115] D.M. Tanner. Mems reliability: Where are we now? *Microelectronics Reliability*, 49(9–11):937–940, 2009.

[116] Emmanuel Munguia Tapia, Stephen S. Intille, William Haskell, Kent Larson, Julie Wright, Abby King, and Robert Friedman. Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. In *Proceedings of the 11th International Conference on Wearable Computers (ISWC '07)*, pages 37–40, Boston, MA, 11–13 October 2007. ISBN: 978-1-4244-1452-9.

[117] C. D. Thake, M. J. Zurawlew, M. J. Price, and M. Oldroyd. The effect of heat acclimation on thermal strain during explosives ordnance disposal (EOD) related activity in moderate and hot conditions. In *Proceedings of the 13th International Environmental Ergonomics Conference (ICEE 2009)*, pages 113–117, Boston, USA, August 2009. University of Wollongong:Australia. ISBN: 978-1-74128-179-8 (online).

[118] C. Douglas Thake and Mike J. Price. Reducing uncompensable heat stress in a bomb disposal (EOD) suit: a laboratory based assessment. In Igor B. Mekjavic, Stelios N. Kounalakis, and Nigel A.S. Taylor, editors, *Proceedings of the 12th International Conference on Environmental Ergonomics (ICEE 2007)*, pages 229–232, Piran, Slovenia, 19–24 August 2007. Ljubljana: Biomed d.o.o. ISBN: 978-961-90545-1-2.

[119] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA, 2005.

[120] University of Waikato. Weka 3: Data Mining Software in Java. `http://www.cs.waikato.ac.nz/ml/weka/`; Accessed 12 April 2012.

[121] Susan M Watkins. *Clothing: The Portable Environment*. Iowa State University Press, 1995.

[122] Min Xu, A. Goldfain, A.R. Chowdhury, and J. DelloStritto. Towards accelerometry based static posture identification. In *Proceedings of the 1st IEEE International Workshop on Consumer eHealth Platforms, Services and Applications, (CCNC)*, pages 29–33, Las Vegas, NV, 9– 12 Jauary 2011.

[123] Yang Xue and Lianwen Jin. A naturalistic 3D acceleration-based activity aataset & benchmark evaluations. In *Proceeding of the IEEE International Conference on Systems Man and Cybernetics (SMC)*, pages 4081–4085, Istanbul, Turkey, 10–13 October 2010.

[124] Che-Chang Yang and Yeh-Liang Hsu. A review of accelerometry-based wearable motion detectors for physical activity monitoring. *Sensors MDPI journal*, 8:7772–7788, 20 August 2010.

[125] Guang-Zhong Yang. *Body Sensor Networks.* New York: Springer-Verlag, 1 edition, 2006.

[126] Jhun-Ying Yang, Jeen-Shing Wang, and Yen-Ping Chen. Using acceleration measurements for activity recognition: An effective learning algorithm for constructing neural classifiers. *Pattern Recognition Letters*, 29(16):2213–2220, December 2008.

[127] Hong Ying, C. Silex, A. Schnitzer, S. Leonhardt, and M. Schiek. Automatic step detection in the accelerometer signal. In *Proceedings of the 4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, volume 13, pages 10–85, RWTH Aachen University, Germany, 26–28 March 2007. ISSN: 1680-0737.

[128] Alex D. Young, M. J. Ling, and D. K. Arvind. Orient-2: a realtime wireless posture tracking system using local orientation estimation. In *Proceedings of the Fourth International Workshop on Embedded Networked Sensors (EmNets '07)*, pages 53–57, Cork, Ireland, 25–26 June 2007. New York, NY, USA: ACM.

[129] Shumei Zhang, Paul McCullagh, Chris Nugent, Huiru Zheng, and Matthias Baumgarten. Optimal model selection for posture recognition in home-based healthcare. *International Journal of Machine Learning and Cybernetics*, 2(1):1–14, March 2011.

Appendix A

Instrumentation system

This appendix hardware details for the instrumentation systems implemented during the course of the research described in this thesis. The hardware implementation (a mix of off-the-shelf and in-house components) is described, along with the method of calibrating the accelerometers used (Section A.1). The implemented instrumentation systems are evaluated in terms of battery life, communication range, communication bandwidth, end-to-end latency, information yield and wearability (Section A.2). A supplementary data gathering system based on SHIMMER hardware is also described (Section A.3).

A.1 Prototype on-body hardware implementation

This section describes the hardware used in the implementation of the prototype system described in Section 3.4 on page 49. The choice of system hardware platform for the prototype was influenced partly by the prototype implementation of Medusa2. One of the future goals is the tight integration of the two systems, which would be likely to include dual-use of the physical nodes (as both systems include upper-body and lower-body nodes). Furthermore, the platform used for the Medusa2 system (the Gumstix Verdex embedded computer [51]) is relatively powerful and is thus capable of handling complex data processing algorithms that may be deployed during system development and during posture-related investigations.

The node hardware platform is described in Section A.1.1, the accelerometer sensor boards are described in Section A.1.2, and the accelerometer calibration procedure given in Section A.1.3. Note that the DG1 and DG2 systems used for data gathering (without the ability to perform on-body classification) were also built using the same hardware platform.

A.1.1 Node hardware platform

The Class-act system implemented consists of two on-body processing nodes (the Primary Node and Secondary Node) providing data to an external system. Gumstix Verdex XM4-bt devices (shown in

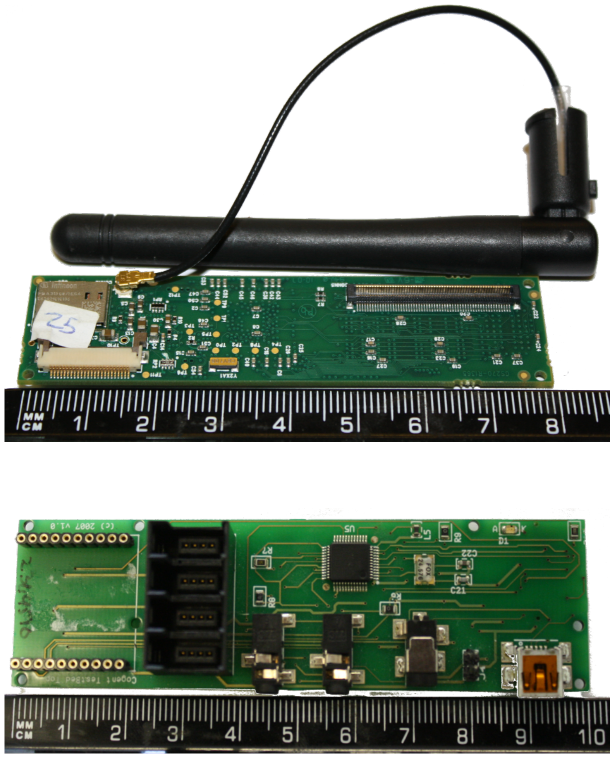

Figure A.1: Node hardware platform—Gumstix Verdex device (top) and expansion board (bottom).

Figure A.1 (top)) are used as the main processing and communications platform for the on-body nodes. These devices are fully-functional single-board computers with a footprint of $80 \times 20 \times 6.3$ mm^3 and a weight of 8 grams (14 grams with the Bluetooth antenna attached). The devices include a 400 MHz Marvell PXA270 XScale CPU and a Bluetooth radio. At the same time, the Verdex devices are small and light enough to be easily carried in a pouch or pocket. The device is enhanced with an expansion board (shown in Figure A.1 (bottom)) with a footprint of $100 \times 30 \times 17$ mm^3 and a weight of 16 grams. This board provides connections for I^{2}C devices and a battery.

A.1.2 Accelerometer sensor boards

A bespoke acceleration sensor board was designed and produced, as shown in Figure A.2 on the facing page. Several of these boards are connected to each Verdex device via the expansion board. The I^{2}C bus allows more than one sensor to share the same bus, reducing the number of wires and I/O ports needed.

Each sensor board consists of a Microchip PIC24FJ64GA002 microcontroller, an Analog Devices ADT75A temperature sensor, and a STMicroelectronics LIS3LV02DQ triaxial accelerometer (with a selectable range of ± 2 g or ± 6 g). The board was designed as a low-cost, small size, low-power wearable solution based on commodity components, and has a footprint of $41 \times 30 \times 7$ mm^3 and a weight of 4 grams.

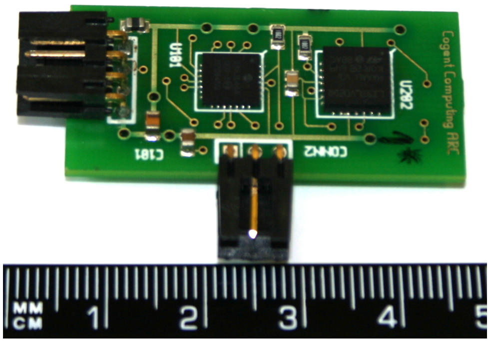

Figure A.2: Acceleration sensor board, which includes a PIC microcontroller, triaxial accelerometer, and digital temperature sensor.

Table A.1: Calibration constants for the set of sensors used in the prototype system. *Offset is given in raw sensor output units (a correctly calibrated sensor outputs a value of 1024 when experiencing an acceleration of 1 g).*

		X		Y		Z	
	Sensor	Slope	Offset	Slope	Offset	Slope	Offset
L1	legrd	0.979	21.50	0.951	2.67	0.941	20.67
L2	legru	0.967	-36.33	0.964	-7.17	0.936	-97.17
L3	leglu	0.964	20.83	0.953	-30.00	0.933	40.33
L4	legld	0.976	77.00	0.942	-20.83	0.926	-0.50
L5	hip	0.974	3.50	0.942	-31.17	0.941	-95.50
L6	ankle	0.968	-18.33	0.948	-24.33	0.933	-19.67
U1	armld	0.972	18.00	0.949	-21.17	0.942	13.33
U2	armlu	0.957	-28.17	0.957	-27.50	0.934	-58.50
U3	body	0.967	1.00	0.951	-20.83	0.942	-29.00
U4	armrd	0.968	-18.33	0.948	-24.33	0.933	-19.67
U5	armru	0.970	-15.33	0.943	-19.50	0.939	-11.83

The accelerometer range used in this work was ± 2 g. The provision of a temperature sensor in the board design eases the future integration of the Medusa2 and Class-act systems by reducing the total number of body sensors required—reducing overall system complexity and helping to increase wearability.

A.1.3 Accelerometer static calibration and unit conversion

For the purpose of accelerometer calibration, each sensor board was mounted in turn on a camera stand. The stand provided angle indications to allow accurate orientation of the sensor (as is shown in Figure A.3). Mounting was performed using a bespoke fixture. Each axis was oriented to observe $+1$ g, -1 g, $+\sin(45)$ g, $-\sin(45)$ g and zero acceleration. A single reading was obtained in each orientation by averaging one minute of samples taken at 10 Hz (600 samples total). Calibration was performed as a linear transformation based on the measurements in the $+1$ g, -1 g, $+\sin(45)$ g, $-\sin(45)$ g and zero

Figure A.3: Sensor board mounting for accelerometer calibration

orientations. Adjustment of raw readings is thus performed as

$$x' = mx + c$$

where x' is the calibrated acceleration value in sensor units, x is the raw measured value of acceleration in sensor units, m is the calculated slope constant for the axis, and c is the calculated offset constant for the axis. Table A.1 shows the calculated slope and offset for the accelerometers used in the system implemented here.

Once raw readings have been adjusted based on the calibration constants found, they are converted into SI units (ms^{-2}) by dividing the values by 1024 (the output given by a correctly calibrated accelerometer of this type when experiencing an acceleration of 1 g) and multiplying by 9.8.

A.2 Class-act prototype instrument evaluation

This section provides a functional evaluation of the prototype instrumentation system described in this appendix with regard to: battery life (Section A.2.3), communication range (Section A.2.4), communication bandwidth (Section A.2.5), end-to-end latency (Section A.2.1), information yield (Section A.2.2) and wearability (Section A.2.6).

A.2.1 End-to-end latency

End-to-end latency is defined here as the time taken between gathering a set of data samples at the sensor and receiving either the data or the processed postural information at the base station. In the EOD

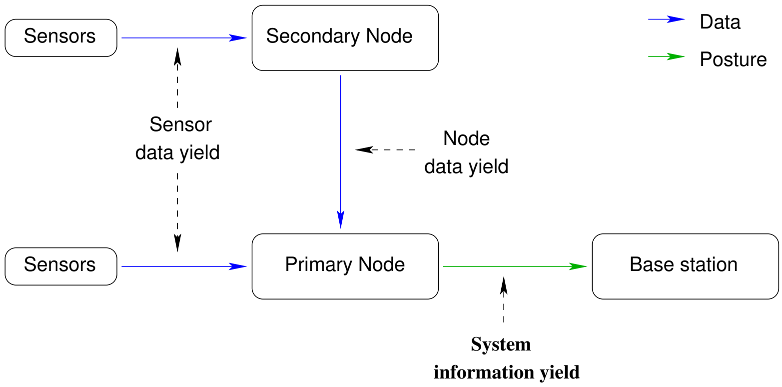

Figure A.4: Yield measurement locations for the prototype system.

application, for example, low latency is important to ensure that the Medusa2 system (external system) uses, in its prediction and control algorithm, the most recent postural information available—incorrect postural information will increase the error in the output of the predictive modelling performed. The importance in evaluating the end-to-end latency of the instrument is in demonstrating that the system provides real-time on-body classification—part of the contribution of Chapter 3.

For latency measurement purposes, the clocks on each node were synchronised to the base station via NTP [89]. This generally results in synchronisation to within around 10^{-2} seconds—at least an order of magnitude smaller than the sampling rate. The timestamp attached to the data (set when it was collected at the Secondary Node) was compared with the time that the data was received at the Primary Node and then at the base station. The time taken to perform a classification was also tested separately via a benchmarking algorithm that performed a large number of classifications in order to get an accurate estimate of the average time per classification. Experimentally, the average latency measured was 30 ms for Secondary Node to Primary Node and Primary Node to base station communication and 0.8 ms to perform posture classification on the Primary Node. Thus the end-to-end latency is 60.8 ms. This latency is satisfactory as it is around half the sampling period (at 10 Hz), meaning that the samples received at the base station will be the most recently gathered samples (that is, there is not a more recently gathered sample that has yet to be received). This meets the real-time requirement as defined in Section 1.1 and demonstrates that the system provides real-time on-body classification of posture.

A.2.2 Information yield

Yield is an end-to-end system measure and is defined here as the number of samples received by the Medusa2 system or base station as a percentage of the number of samples expected in a given time period. In this case, a "sample" refers to a sensor data vector or a classified posture as required by the transmission mode in use. Figure A.4 on the previous page shows the points in the system data flow at which various types of yield may be measured. However, as Class-act is intended to be deployed as an end-to-end posture monitoring solution, only the system information (posture) yield of the system is considered here (in bold in Figure A.4).

The information yield for the system was determined by analysing the experimental data gathered through deployment of the prototype system during the development and evaluation process. Specifically, the system was analysed in the "analysis" transmission mode (transmitting both data and posture). In this mode, the system achieved a yield of 83%. Yield in the "information" transmission mode (transmitting only posture) was demonstrated in Section 4.12 on page 76 to be greater than 99%. The conclusion drawn is that the low yield for the system in the "analysis" transmission mode is related to the Gumstix Verdex platform—the Primary Node has only one Bluetooth chip and so it has to switch between transmitting and receiving modes as required, which results in delays in receiving data from the Secondary Node and potentially dropped packets or missed classifications. The EOD application requires use of the "information" transmission mode, for which a high yield is provided.

A.2.3 Battery life

In this section, the battery life of the system is evaluated in relation to the minimum acceptable system life for the two application scenarios. A basic requirement is that the battery life of the system must be longer than the duration of the mission/activity/experiment engaged with by the subject. For the EOD application, the system battery life should be at least 1 hour since the typical EOD mission is expected to be carried out in stages of up to that duration, after which time the EOD operative returns to the support team (which would allow a change of batteries). In the investigative instrument application (taking, as an example, deployment during the work performed here) the battery is also required to last up to 1 hour in order to perform data gathering experimentation.

To maintain consistency with expected normal usage of the system when evaluating battery life:

1. Acceleration data was sampled at 10 Hz from each sensor axis by both the Primary and Secondary Nodes. Five sensors were connected to the Primary Node and four sensors were connected to the Secondary Node.

2. Acceleration data and extracted features were transmitted from the Secondary Node to the Primary Node, where posture classification was performed.

Three transmission modes from the Primary Node to the base station were tested: data and posture, posture only, and posture updates only. The effective life of the system was defined as the life of the Primary Node, which was the first node to fail in all cases due to the additional transmission/processing requirements compared to the Secondary Node. For each transmission mode, the test was performed five times and the results were averaged to obtain the reported figure.

The resulting system battery life (using four 1100 mAh batteries) was: 1 hour 50 minutes when transmitting data and posture, 2 hours when transmitting only posture, and 2 hours 20 minutes when transmitting only posture updates. This meets the requirement for both applications.

Scenarios other than those considered here may require longer operating lifetime. This could be obtained via the use of more batteries, however the batteries are the heaviest items in the system and thus can be expected to have an impact on wearability/comfort. Instead, solutions must be found in the system design to reduce power consumption. Two possible avenues for this are related to the sensor devices and radio communications. As described previously, the battery life testing performed essentially provides a lower bound on the expected lifetime as it includes the full set of nine sensors and the radio is always on. The first solution to increasing lifetime is therefore to reduce the number of sensors used, and thus their total power draw. As shown in Section 6.6 on page 111, posture classification can be performed using only two sensors (on the calf and lower leg). Doran [40] showed that the use of two sensors as opposed to nine provides an increase in battery life of around 30 minutes. The use of only two sensors would, additionally, mean that only one on-body node is required, further reducing the total weight of the system by eliminating a node and the associated battery pack. The possible second solution for increasing battery life is to power off the node radios when not in use. This would primarily be of use in an event-based system where transmissions would be expected to be less frequent. Doran [40] demonstrated a 40% reduction in power consumption when the radio is deactivated for nodes of the type used here. The disadvantage of this method is the additional latency incurred in establishing the communications link when the radio is powered on, the impact of which is application-dependent. This would not be suitable for the EOD application, for example, as the posture information is required for predictive modelling of heat stress and so any delay decreases the accuracy of the prediction and the effectiveness of any countermeasures employed.

A.2.4 Communication range

Communication range was tested to determine the distance at which the base station or gateway is still capable of receiving data from the monitored subject.

Table A.2: Node transmission rates for the prototype system.

	Transfer rate (kB/s)
Primary Node to Secondary Node	106.7
Secondary Node to Primary Node	116.9
Primary Node to base station	62.3
Secondary Node to base station	54.0

Only on-body communication is required in the EOD application (around 30 cm) as Class-act only needs to transmit to the Medusa2 processing node, which is located on the operative's body in close proximity to the Class-act nodes. The investigative instrument application, on the other hand, places importance on the communication range of the system due to the inherent mobility of the subjects. The communication range in the investigative instrument application must be around 10 m, as in the intended usage the subject would be in the same room as the base station.

The specified communication range for a Bluetooth class 2 device (as used in the Class-act and Medusa2 systems) is approximately 10 m when working at a maximum permitted broadcast power of 2.5 mW. Testing was performed in a variety of environments similar to those that may be encountered in the two applications. In all cases, the base station was fixed in a static location, while the nodes were gradually moved away until communication was observed to become unreliable (for example, missing data or high latency spikes). In summary, the experimentally observed ranges were: 14 m with several walls and some light office machinery between the nodes and base station, 49.4 m within a straight corridor with line of sight, and 63.5 m outdoors with clear line of sight. The requirements for both application examples are thus clearly met.

A.2.5 Communication bandwidth

Bandwidth is a measure of the number of bits transmitted per second. It is an important metric since it is one of the factors that determines the system's ability to provide postural information in real-time—if the available bandwidth is exceeded then the output will be delayed. For the EOD application, the communication bandwidth required is 100 bytes/second (at a sampling rate of 10 Hz) since the system is only required to transmit posture information. In the investigative instrument application, any of the transmission modes could be used for different types of experimentation and thus the mode considered for evaluation here combines data and postural information—the most bandwidth-consuming mode, at around 1.6 kB/s (at a sampling rate of 10 Hz).

The theoretical available transfer rate for Bluetooth 2.0+EDR is 2.1 Mb/s (263 kB/s), which is sufficient to transmit the data required. However the available bandwidth is lower than the theoretical maximum. This was tested by transferring a large (10 MB) file from the Primary Node to Secondary Node, from Secondary Node to Primary Node and from both nodes in turn to the base station. This test was repeated five times for each sender/receiver pair. Table A.2 on the facing page shows the measured average transfer rate. While lower than the theoretical maximum, the transfer rate is still sufficient to transmit the full set of data and posture when sampling acceleration at 10 Hz.

A.2.6 Wearability

Wearability was discussed previously in Section 2.2 where several aspects related to establishing the wearability of a system for a specific target user group are described, including the physical dimensions of the devices in the system, the way in which it affects movement and so on. While no explicit requirements have been specified here per-scenario for wearability, it can be assumed that in both cases, subject comfort is a concern.

In terms of system weight, a full deployed prototype system with nine sensing locations weighs a total of 483 g, with the majority of this weight being the on-body nodes and battery packs (which are carried in a pouch worn by the subject). Each of the sensor boards weighs only 4 g. As the majority of the weight of the system is carried in the pouch it is less noticeable than, for example, the case where each sensor has a separate battery pack (as implemented in other systems described in the literature [39, 49, 59, 62]).

For the prototype system presented, the focus was on data gathering and algorithm development. Thus a thorough investigation into wearability was not pursued at this stage. However, at the conclusion of each experiment performed, the subject was asked about their general comfort and whether the system restricted their movement or was felt to be cumbersome. It was found that the subjects could move freely without discomfort caused by the system. However, care needed to be exercised for particular postures which were likely to apply pressure upon the nodes (in the postures studied here this was primarily the lying on one side and lying on front postures). The sensors are easy to place and their small footprint allows for them to be taped to a subject with minimal effort. The shape of the underside of the sensor board however means that padding was required to be applied onto the subject's skin prior to the sensors being mounted to avoid bruising. Typical set-up time (inclusive of start-up and fitting to subject) for the system is around 18 minutes.

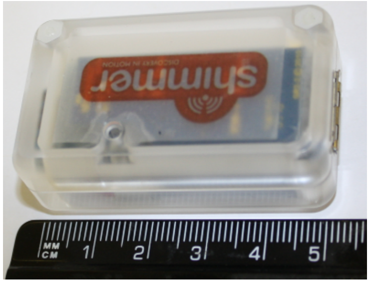

Figure A.5: SHIMMER node.

A.3 SHIMMER-based data gathering instrument

In addition to the prototype monitoring instrumentation described in this chapter, an additional instrument was implemented by the author to support experimental data acquisition at a higher sampling rate (required for the sampling rate investigation in Section 6.7 on page 116). This system is based around SHIMMER devices [107], one of which is shown in Figure A.5. Seven of these devices were used, communicating wirelessly through Bluetooth with each device transmitting data directly to the base station (a laptop computer). The SHIMMER devices include a three-axis accelerometer, the Freescale MMA7260Q with a range of ±6g, along with an MSP430F1611 microcontroller running at up to 8 MHz and a 280 mAh battery. The devices are packaged into a 53.3×31.8×17.8 mm^3 enclosure.

At the time that the experimentation was performed, there was no existing implementation of the NTP time synchronisation method for the SHIMMER devices when using Bluetooth. An alternative method was thus used to allow the samples from each device to be time-aligned for processing. The following process was performed for each device at the start of a data gathering experiment:

1. The average latency for the connection is determined.

   (a) A message is sent to the SHIMMER device, which immediately responds. The latency in each direction is assumed to be half of the total time between transmission of the message and reception of the response.

   (b) This is repeated 20 times and the one-way latencies found are averaged.

2. The offset between the time on the SHIMMER device and the base station is determined.

   (a) A message is sent to the SHIMMER device requesting the current timestamp.

   (b) The average latency found in step 1 is added to the timestamp received.

   (c) The difference between this modified timestamp and the base station time is the offset.

3. The determined offset is stored and applied to each sample received.

The data gathered using this system is processed in the same way as described in Section 4.9 on page 72 for the Class-act system.

A.4 Summary

In this chapter the hardware details for the prototype system implemented during the work in this thesis was described. The implementation is based on the Class-act platform architecture described in Chapter 3 and uses the Gumstix Verdex hardware platform. The hardware platform selected provides several benefits including provision of sufficient processing power to allow real-time classification using a variety of algorithms during system development and ease of future integration with other systems such as the one implemented by Kemp [66] to monitor EOD operatives. The system implementation was evaluated in terms of battery life, communication range, communication bandwidth, end-to-end latency, information yield and wearability. An additional system for data gathering was also developed based on SHIMMER devices that allow data gathering at 100 Hz. This alternate system was used in the analysis in Section 6.7 on page 116.

Appendix B

Publications

B.1 Journal publications

Edge mining the Internet of Things

Elena Gaura, James Brusey, Michael Allen, Ross Wilkins, Dan Goldsmith, Ramona Rednic. Edge mining the internet of things. IEEE Sensors Journal May 2013 , To appear

This paper examines the benefits of edge mining— data mining that takes place on the wireless, battery-powered, smart sensing devices that sit at the edge points of the Internet of Things. Through local data reduction and transformation, edge mining can quantifiably reduce the number of packets that must be sent, reducing energy usage and remote storage requirements. Additionally, edge mining has the potential to reduce the risk to personal privacy through embedding of information requirements at the sensing point, limiting inappropriate use. The benefits of edge mining are examined with respect to three specific algorithms: Linear Spanish Inquisition Protocol (L-SIP), ClassAct, and Bare Necessities (BN), which are all instantiations of General SIP (G-SIP). In general, the benefits provided by edge mining are related to the predictability of data streams and availability of precise information requirements; results show that L-SIP typically reduces packet transmission by around 95% (20-fold), BN reduces packet transmission by 99.98% (5000-fold) and ClassAct reduces packet transmission by 99.6% (250-fold). Although energy reduction is not as radical due to other overheads, minimisation of these overheads can lead to up to a 10-fold battery life extension for L-SIP, for example. These results demonstrate the importance of edge mining to the feasibility of many IoT applications.

A Web-Based System for Home Monitoring of Patients with Parkinson's Disease Using Wearable Sensors

Bor-rong Chen, Shyamal Patel, Thomas Buckley, Ramona Rednic, Doug McClure, Ludy Shih, Daniel Tarsy, Matt Welsh, and Paolo Bonato. A Web-Based System for Home Monitoring of Patients With

Parkinson's Disease Using Wearable Sensors. IEEE Transactions on Biomedical Engineering (TBME) Letters Special Issue on Emerging Technologies in Point-of-Care Health Care, 58(3):831–836, March 2011.

This letter introduces MercuryLive, a platform to enable home monitoring of patients with Parkinson's disease (PD) using wearable sensors. MercuryLive contains three tiers: a resource-aware data collection engine that relies upon wearable sensors, web services for live streaming and storage of sensor data, and a web-based graphical user interface client with video conferencing capability. Besides, the platform has the capability of analyzing sensor (i.e., accelerometer) data to reliably estimate clinical scores capturing the severity of tremor, bradykinesia, and dyskinesia. Testing results showed an average data latency of less than 400 ms and video latency of about 200 ms with video frame rate of about 13 frames/s when 800 kB/s of bandwidth were available and we used a 40% video compression, and data feature upload requiring 1 min of extra time following a 10 min interactive session. These results indicate that the proposed platform is suitable to monitor patients with PD to facilitate the titration of medications in the late stages of the disease.

Classifying transition behavior in postural activity monitoring

James Brusey, Ramona Rednic, and Elena Gaura. Classifying transition behaviour in postural activity monitoring. Sensor & Transducers Special Issue, 17(10):213–223, October 2009.

A few accelerometers positioned on different parts of the body can be used to accurately classify steady state behaviour, such as walking, running, or sitting. Such systems are usually built using supervised learning approaches. Transitions between postures are, however, difficult to deal with using posture classification systems proposed to date, since there is no label set for intermediary postures and also the exact point at which the transition occurs can sometimes be hard to pinpoint. The usual bypass when using supervised learning to train such systems is to discard a section of the dataset around each transition. This leads to poorer classification performance when the systems are deployed out of the laboratory and used on-line, particularly if the regimes monitored involve fast paced activity changes. Time-based filtering that takes advantage of sequential patterns is a potential mechanism to improve posture classification accuracy in such real-life applications. Also, such filtering should reduce the number of event messages needed to be sent across a wireless network to track posture remotely, hence extending the system's life. To support time-based filtering, understanding transitions, which are the major event generators in a classification system, is key. This work examines three approaches to post-process the output of a posture classifier using time-based filtering: a naïve voting scheme, an exponentially weighted voting scheme, and a Bayes filter. Best performance is obtained from the exponentially weighted voting scheme although it is suspected that a more sophisticated treatment of the Bayes filter might yield better results.

Postural activity monitoring for increasing safety in bomb disposal missions

James Brusey, Ramona Rednic, Elena I. Gaura, John Kemp, and Nigel Poole. Postural activity monitoring for increasing safety in bomb disposal missions. *Measurement Science and Technology*, 20(7):075204 (11pp), July 2009.

In enclosed suits, such as those worn by explosive ordnance disposal (EOD) experts, evaporative cooling through perspiration is less effective and, particularly in hot environments, uncompensable heat stress (UHS) may occur. Although some suits have cooling systems, their effectiveness during missions is dependent on the operative's posture. In order to properly assess thermal state, temperature-based assessment systems need to take posture into account. This paper builds on previous work for instrumenting EOD suits with regard to temperature monitoring and proposes to also monitor operative posture with MEMS accelerometers. Posture is a key factor in predicting how body temperature will change and is therefore important in providing local or remote warning of the onset of UHS. In this work, the C4.5 decision tree algorithm is used to produce an on-line classifier that can differentiate between nine key postures from current acceleration readings. Additional features that summarize how acceleration is changing over time are used to improve average classification accuracy to around 97.2%. Without such temporal feature extraction, dynamic postures are difficult to classify accurately. Experimental results show that training over a variety of subjects, and in particular, mixing gender, improves results on unseen subjects. The main advantages of the on-line posture classification system described here are that it is accurate, does not require integration of acceleration over time, and is computationally lightweight, allowing it to be easily supported on wearable microprocessors.

B.2 Conference proceedings

Long-term Behavioural Change Detection Through Pervasive Sensing

John Kemp, Elena Gaura, Ramona Rednic, James Brusey, Long-term behavioural change detection through pervasive sensing. In Proceedings of the 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2013), Honolulu Hawai, U.S.A., 1–3 July 2013.

The paper proposes an information generation and summarisation algorithm to detect behavioural change in applications such as long-term monitoring of vulnerable people. The algorithm learns the monitored subject's behaviour autonomously post-deployment and provides time-suppressed summaries of the activity types engaged with by the subject over the course of their day to day life. It transmits updates to external observers only when the summary changes by more than a defined threshold. This technique substantially reduces the number of transmission required by a wearable monitoring system, both through summarisation of the raw data into useful information and by preventing transmission of duplicated or predictable data and information. Based on evaluation using simulated activity data, the proposed algorithm results in an average of one transmission per month following an initial convergence period (reaching less than 1 transmission per day after only three days) and detects a change in behaviour after an average of 1.1 days.

Fielded Autonomous Posture Classification Systems: Design and Realistic Evaluation

Ramona Rednic, John Kemp, Elena Gaura, James Brusey. Fielded autonomous posture classification systems: Design and realistic evaluation. In Proceedings of the 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2013), Honolulu Hawai, U.S.A., 1–3 July 2013.

Few Body Sensor Network (BSN) based posture classification systems have been fielded to date, despite laboratory based research work confirming their theoretical suitability for a range of applications. This paper reports and reflects on two algorithms which i) improve the accuracy of real-time, multi-accelerometer based posture classifiers when dealing with natural movement and transitions and ii) maximize a wearable system's battery life through distributed computation at nodes. The EWV transition filters proposed here increase the classification accuracy by 1% over unfiltered results in realistic scenarios and significantly reduces spurious classifier output in real-time visualizations. A 200 fold transmission reduction from the on-body system to an outside system was achieved in practice by combining the transition filters with an event-based design. Furthermore, a method of reducing transmissions between on-body data gathering nodes based on distributed processing of the classifier rules (but maintaining a one-way flow of communications during system use) is also described. This provides a 3.3

fold reduction in packets and a 13.5 fold reduction in data transmitted from one node to the other in a two-node wearable system.

Wearable posture recognition systems: factors affecting performance

Ramona Rednic, Elena Gaura, James Brusey, and John Kemp. Wearable posture recognition systems: factors affecting performance. In Proceedings of the IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI 2012), pages 200–203, Shenzhen, China, 5–7 January 2012.

This paper presents an investigation into the design space for real-time, wearable posture classification systems; specifically, it analyses the impact of various factors/design choices on classification accuracy when using C4.5 decision trees. The factors can be broadly divided into: 1) system factors (such as sensor sampling rate and number of sensors used) and 2) algorithm and training factors (such as quantity of training data and temporal data features used). These factors are analysed in the context of a case study involving postural activity monitoring of Explosive Ordinance Disposal (EOD) operatives. The case study involves classifying a set of eight postures commonly encountered in EOD missions: sitting, walking, crawling, lying (on all sides) and kneeling. Design guidelines and generic lessons for a wider class of applications can be drawn from the work.

Networked Body Sensing: Enabling real-time decisions in health and defence applications

Ramona Rednic, John Kemp, Elena Gaura, and James Brusey. Networked body sensing: Enabling real-time decisions in health and defence applications. In Proceedings of the Annual International Conference on Advance Computer Science and Information Systems 2011 (ICACSIS 2011), pages 17–24, Jakarta, Indonesia, 17–18 December 2011

This paper presents the application scenario, conceptual overview and implementation of a monitoring system targeted at monitoring EOD suit wearers during missions. The system's aim is to deliver prediction of heat stress risk in the operative and provide actuation of a cooling system integrated within the suit. Prior work established that such prediction requires real-time autonomous processing of skin temperature and body acceleration data, and thus a system implementation is presented based on two interacting subsystems that perform the required sensing and data processing. Posture classification is performed with an accuracy of 96.1%, and a heat stress prediction algorithm is demonstrated with an overall accuracy of 88.5% when predicting the occurrence of heat stress within the next 2 minutes

Home monitoring of patients with Parkinson's disease via wearable technology and a web-based application

Shyamal Patel, Bor-rong Chen, Thomas Buckley, Ramona Rednic, Doug McClure, Daniel Tarsy, Ludy Shih, Jennifer Dy, Matt Welsh, and Paolo Bonato. Home monitoring of patients with Parkinson's disease via wearable technology and a web-based application. In Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pages 4411–4414, Buenos Aires, 31 August–4 September 2010. IEEE.

Objective long-term health monitoring can improve the clinical management of several medical conditions ranging from cardiopulmonary diseases to motor disorders. In this paper, we present our work toward the development of a home-monitoring system. The system is currently used to monitor patients with Parkinson's disease who experience severe motor fluctuations. Monitoring is achieved using wireless wearable sensors whose data are relayed to a remote clinical site via a web-based application. The work herein presented shows that wearable sensors combined with a web-based application provide reliable quantitative information that can be used for clinical decision making.

MercuryLive: A Web-Enhanced Platform for Long-Term High Fidelity Motion Analysis

Bor-rong Chen, Thomas Buckley, Ramona Rednic, Shyamal Patel, Paolo Bonato, and Matt Welsh. MercuryLive: A Web-Enhanced Platform for Long-Term High Fidelity Motion Analysis. In Proceedings of the 7th Annual IEEE Communications Society Conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON), pages 1–2, Boston, USA, 21–25 June 2010.

We present MercuryLive, a web-enhanced extension to a body sensor network platform for continuous home-based body motion sensing, interactive supervised data collection sessions, and long-term activity data analysis. The major goal of MercuryLive is to enable practical long-term health monitoring in a home setting and henceforth reduce the effort and cost for collecting clinically relevant quantitative measures on patients' health conditions during daily activities. MercuryLive contains three tiers: a central web server for streaming and storage of sensor data, a sensor data collection engine, and a user-friendly web-based GUI client. The platform is currently used in clinical studies on Parkinson's disease.

ClassAct: Accelerometer-based Real-Time Activity Classifier

Ramona Rednic, Elena Gaura, and James Brusey. ClassAct: Accelerometer-based real-time activity classifier. In Sensors & Instrumentation KTN: Wireless Sensing Demonstrator Showcase (WiSIG), 2 July 2009.

In enclosed bomb disposal suits, posture affects the air flow and is thus a key indicator for predicting the onset of Uncompensable Heat Stress (UHS). In order to allow the exploration of this effect, a system was developed to monitor the posture of human subjects during bomb disposal missions using only low cost accelerometers. Decision trees are used to identify in real-time, within the suit, eight mission-like postures: standing, kneeling, sitting, crawling, walking and lying on front, back, and one side. A variety of time domain features were explored to aid differentiation between static and dynamic postures. An average classification accuracy of 97.2% over the nine postures are obtained when using windowed variance and nine accelerometers. Similar performance was obtained with as little as two accelerometers, whilst a single hip accelerometer was shown to classify standing, walking and sitting with an average accuracy of 96.4%. Overall the instrument exhibits a suitable level of performance for the application at hand, in terms of wearability, accuracy, timeliness and data yield. The classification technique developed could be extended to the classification of other task oriented activities.

Wireless sensor networks for activity monitoring in safety critical applications

Ramona Rednic, Elena Gaura, James Brusey, and John Kemp. Wireless sensor networks for activity monitoring in safety critical applications. In Proceedings of NSTI Nanotech 2009, volume 1—Fabrication, Particles, Characterization, MEMS, Electronics and Photonics, pages 521–525, Houston, Texas, USA, May 3–7 2009. ISBN: 978-1-4398-1782-7.

The ability to monitor posture is essential to many application areas, including virtual reality, health, and sports applications. The work here focuses on the use of postural monitoring in safety critical missions such as explosive ordnance disposal (EOD) missions. The operatives undertaking these missions are commonly placed under a high level of physical and psychological strain due to the weight of the protective armoured suit and the potential risk of their work. Remote monitoring of posture may allow a better understanding of the operative's status. When combined with additional health information, posture can enhance the accuracy of operative's global state estimation. Previously, a Body Sensor Network-based (BSN) posture monitoring system consisting of nine accelerometers was designed and implemented by the authors here. The system was able to recognise six specific postures (sitting, kneeling, crawling, and three variations of lying on the ground) with high accuracy. However, the system was unable to consistently distinguish between a subject standing, walking or running. In order to counteract this limitation, a new prototype utilising additional sensors and an augmented data processing method has been implemented and evaluated and is reported here.

B.3 Technical reports

COGENT.006: Posture Determination Using a Body Sensor Network

Ramona Rednic, John Kemp, Elena Gaura, James Brusey. Posture Determination Using a Body Sensor Network. Technical Report COGENT.006, Coventry University, 2008.
http://cogentee.coventry.ac.uk/tech_reports/COGENT.006.pdf

Due of the large number of degrees of freedom of the human body, posture monitoring of human during activity regimes presents many research challenges. Several research groups world wide have engaged with the development of low-power wireless body sensor networks that are capable of providing real-time posture tracking for a variety of applications, such as dance and sports. The work reported here is concerned with the development of a wireless body sensor network that, as opposed to posture tracking, can: a) provide the identification and classification of eight human postures (standing, kneeling, sitting, crawling, walking, lying down on front and back, and lying on one side) in real-time and b) is able to relate this information wirelessly to a remote monitoring point. Posture information is an essential part of monitoring operatives in safety critical missions. The work sits within a larger project aiming to increase general safety of operatives in bomb disposal missions.
The goal of the posture body sensor network developed here is to identify the eight named postures using data from nine accelerometers placed at various sites on the human body. A prototype implementation which fulfills the goal has been produced and evaluated and is reported here.

Selected publications follow.