Title
Realistic-Motion Activity Recognition

Author
Mortazavi, Jack Bobak

Publication Date
2014

Peer reviewed|Thesis/dissertation

University of California
Los Angeles

Realistic-Motion Activity Recognition

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Computer Science

by

Jack Bobak Mortazavi

2014

Abstract of the Dissertation

Realistic-Motion Activity Recognition

by

Jack Bobak Mortazavi

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2014

Professor Majid Sarrafzadeh, Chair

Sedentary behavior starting from a younger and younger age has led to many serious medical conditions including obesity, cardiovascular diseases and diabetes. These conditions have led not only to a significant health burden but a significant economic burden in treating such ailments. Advances in sensor technology have let us not only monitor human activity but use such activity in applications to help address this sedentary life style. Exergaming is the convergence of using physical activity monitoring techniques and video game design in order to create healthy video game activities.

This dissertation investigates the design and implementation of such exergames, culminating in the creation of one such game that encompasses multiple pieces to address the obesity epidemic and the several parallel conditions that must be met. New classification algorithms must be created to identify these fine-grain movements, often in a large, multi-class setting. The analysis of such an algorithm must incorporate two considerations, the accuracy as well as the responsiveness of such a system. This work investigates developing such an algorithm for these fine-grain detailed motions, achieving new, high levels of accuracy, as well as the appropriate trade-off between the complexity and accuracy of such systems and the responsiveness to detecting the appropriate movements in a responsive fashion, leveraging contextual information to strengthen the classification results for

a user-centric recognition experience. Further, clinical trials are run to investigate the energy expenditure levels of such a system, showing and guaranteeing levels of energy expenditure that can promote a more active and healthy gaming experience that is still enjoyable to the user.

The methods investigated in this dissertation show an activity level of moderate activity, and regression techniques that can predict these activity levels within an error of at most only 1 metabolic equivalent of task. User-specific models can improve the computational complexity and accuracy of such systems, reducing the delay needed to classify and improving classification results by about 12% when training a model for a specific user. Further, the multiple model approach investigated helps improve classification of difficult, fine-grain activities with greater than 90% accuracy and F-score, creating an end-to-end system for detailed physical motion recognition and a complex wireless health system and application.

The dissertation of Jack Bobak Mortazavi is approved.

Todd Millstein

Douglas Stott Parker

Greg Pottie

Majid Sarrafzadeh, Committee Chair

University of California, Los Angeles

2014

To my parents who made this work possible with their everlasting love,

to my sister for the support only a sibling can give,

to my brother-in-law for being the brother I never had,

and to my nephews for always cheering me up and helping me relax.

Table of Contents

List of Figures

List of Tables

Acknowledgments

I would first like to acknowledge Professor Majid Sarrafzadeh for his outstanding supervision throughout my graduate career. His invaluable insights into research, teaching, and the whole of academia, as well as industrial experience, in patenting, start up work, and consulting, are what have made me the researcher that I am today. It was with his help and mentoring that I have been able to find the thing I love doing most and succeeding at it. For that I will be forever grateful.

I would also like to acknowledge Professor Christian K. Roberts for his assistance in developing the clinical work. I would also like to acknowledge Professor Greg Pottie for his insight and contributions to my work, as well as his knowledge and supervision in mentoring young students in the WHI summer intern program. Further, I would also like to acknowledge my other committee members, Professor Stott Parker and Professor Todd Millstein, for their valuable insights and experience in further defining and refining this dissertation.

Finally, I would like to acknowledge Jonathan Woodbridge and Hagop Hagopian, for their contributions in developing this work, and Suneil Nyamathi, for his contributions to developing the demonstration applications. I would also like to acknowledge the countless lab-mates, friends, and Pahlavan brothers, all too many to list here, for all of their help, spanning the spectrum from contributing to work to providing emotional support. Each and every person was important in their own unique way in helping me get to where I am today.

Vita

2007	B.S. (Electrical Engineering and Computer Science) and B.A. (Applied Mathematics), University of California, Berkeley.
2009	M.S. (Electrical and Computer Engineering), University of California, Irvine.
2011–present	Teaching Assistant/Associate/Fellow, Computer Science Department, UCLA.
2012–2013	Lecturer, California State Polytechnic University, Pomona, Department of Electrical and Computer Engineering.

Publications

B. Mortazavi, S. Nyamathi, S.I. Lee, T Wilkerson, H. Ghasemzadeh, M. Sarrafzadeh. Near-Realistic Mobile Exergames with Wireless Wearable Sensors. IEEE Journal of Biomedical and Health Informatics (J-BHI). 2014

B. Mortazavi, M. Pourhomayoun, N. Alshurafa, M. Chronley, C.K. Roberts, M. Sarrafzadeh. Support Vector Regression for Exergaming Actions. International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). 2014

B. Mortazavi, M. Pourhomayoun, G. Alsheikh, N. Alshurafa, S.I. Lee, M. Sarrafzadeh. Determining the Single Best Axis for Exercise Repetition Recognition

and Counting in SmartWatches. IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN). 2014

S.I. Lee, H. Ghasemzadeh, B. Mortazavi, M. Sarrafzadeh. Pervasive Assessment of Motor Function: A Lightweight Grip Strength Tracking System. IEEE Transactions on Technology in Biomedicine (TITB). 2013

B. Mortazavi, N. Alshurafa, S.I. Lee, M. Lan, M. Chronley, C.K. Roberts, M. Sarrafzadeh. MET Calculations from On-Body Accelerometers for Exergaming Movements. IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN). 2013

S.I. Lee, H Ghasemzadeh, B Mortazavi, M Lan, M. Ong, M. Sarrafzadeh. Remote Monitoring Systems: What Impact Can Data Analytics Have on Cost? Wireless Health (WH). 2013

J. Woodbridge, B. Mortazavi, A. Bui, M. Sarrafzadeh. A Monte Carlo Approach to Biomedical Time Series Search. IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2012.

B. Mortazavi, K.C. Chu, X. Li, J. Tai, S. Kotekar, M. Sarrafzadeh. Near-Realistic Motion Video Games with Enforced Activity. 9th International Conference on Wearable and Implantable Sensor Networks (BSN). 2012

B. Mortazavi, H. Hagopian, J. Woodbridge, B. Yadegar, M. Sarrafzadeh. A Wireless Body-Wearable Sensor System for Designing Physically Interactive Video Games. International Conference on Biomedical Electronics and Devices (BioDevices). 2011

CHAPTER 1

Introduction

The increased prevalence of obesity in the world has brought attention to the causes of an overweight population and the impacts this epidemic has. Childhood and adult obesity are a serious concern to the future of health care around the world. Projections in 2008 placed approximately 90% of all American adults as overweight or obese by the year 2030 [WBL08b]. A similar study in 2012 found that by 2030 51% of the population will be obese and that 11% will be severely obese [FKT12]. The obesity epidemic is a worldwide problem as well. The World Health Organization estimates that more than 400 million adults are currently obese worldwide (9.8%) [HPT12]. European nations, such as the United Kingdom, believe that nearly a quarter of all adults are obese as of 2007 and that, if no action is taken, by 2050, 60% of men and 50% of women will be obese. This growing problem is now beginning with children, with the same trend being true in Switzerland [SSS04] and again the United Kingdom, where the same projections place 25% of all children as obese by 2050 [HPT12].

This growing burden not only poses a significant health risk to the world, but also a significant economic burden. In the United States, alone, prevalence of obesity in children and adolescents was 16.9% [OCK12], while in adults it was 35.5% in men and 35.8% in women [FCK12]. This corresponds to a direct medical cost of $1723 per person, or, on aggregate a cost of $113.9 billion, which is up to 10% of all US healthcare spending [TWG11]. Indeed, globally, studies and reviews have

shown that medical costs of obese individuals were approximately 30% greater than those individuals with normal weight [WA10]. Addressing this global issue is important socially, and economically, and fighting it must start at an early age. A 2012 study showed that if obesity were to remain even at 2010 levels, not even decrease, that $549.5 billion could be saved by 2030[FKT12]. Sedentary behavior is a root cause of several chronic conditions affecting health of adult and children in the United States and worldwide[WA00][BRL12]. Such physical inactivity often leads to overweight and obese populations [FKT12] that result in chronic conditions such as cardiovascular disease or diabetes. Cardiovascular disease and diabetes both present a significant economic burden in the United States as well as health, with cardiovascular disease accounting for an estimated $670 billion in health care costs in 2010 [HTK11] and estimates suggesting diabetes may approach $900 billion in health care costs by 2015 [Ass13]. As a result, many solutions using wireless wearable sensors have been presented to monitor activity [EPM08][BKW12].

The root factors of this obesity stem from, amongst other factors, sedentary behavior as a result of too much television and video game play. Studies in Xi'an city, China, show that amongst the most important factors in childhood obesity were television and video game time [YYC12]. Further increased sedentary behavior amongst children associated with television watching and video game playing [Rob99, SSS04, RVB08, DDD12] is resulting in an increasingly obese youth. Specifically with regards to video game use, the prevalence of obesity increased from around 5% in children spending no time playing video games to over 20% when playing three hours of electronic games per day [SSS04]. Indeed, video games are as popular as ever. In January 2010, Activision Blizzard Inc. reported that their game, Call of Duty®: Modern Warfare®2 had earned more than $1 billion in worldwide sales since its release in November of 2009, having sold enough

units to earn $550 million in only its first five days of retail availability [Act10]. Likewise, Electronic Arts Inc.'s FIFA Soccer 10, according to a press release published on December 1st, 2009, had already sold over 4.5 million units making it the fastest selling sports video game in the world. EA also estimated that players of FIFA 10 average somewhere around three (3) million games online daily [Ele09]. Since that time, video game sales have continued to grow. On October 4, 2011, EA reported that their newest edition, FIFA 12, sold a record 3.2 million units in less than one week and that fans had already played more than 200 million minutes online alone[Ele11b]. On September 20, 2011, EA reported that NHL 12, their popular ice hockey video game, had generated more than $27 million worldwide, up 19% from the previous year [Ele11a]. More recently, EA's FIFA 14's demonstration model had over 5 million users, before the full game was even released. Thus, video games, and in particular sports video games, seem like a natural target for addressing this childhood obesity. From this emerges the field of exergaming, making video games themselves a source of exercise and health education, so that the obesity epidemic may be slowed or stopped.

In order to implement proper exergaming solutions, many have researched the ability to use accelerometers on human bodies to turn the human body into a controller for video games [WCF07]. More recently, studies have been conducted that enforce the idea that popular commercial systems such as Nintendo's Wii do, in fact, promote physical activity and health [PMJ08, GPH09, MYO10, RVB08]. Wii Fit and Kinect games benefit from having a structure that targets a specific experience or activity [MB11, MEV11]. Furthermore, users are more motivated to continue playing multiplayer exergaming, having a drop-off of only 15% of people compared to 64% for single player games [PMJ08]. More generally, the tools are in place to use active games as not only an educational tool [Pap09] but as a source of moderate physical activity [GPH09, MYO10, Pap09, SHD09]

for healthy living, in college age students and adults [SHD09] as well as in children [GPH09]. Work from [MYO10] shows that activity can reach levels recommended by the American Heart Association for daily health requirements. Energy expenditure systems [BKV97, EPM08, PEK06] show the ability and accuracy of systems measuring particular human movements. However, while the activities themselves in [GPH09, MYO10, Pap09, SHD09] have shown the potential to be healthy, little has been done to guarantee levels of activity. While it is possible to achieve exercise levels of activity, it is not a requirement for the games and it is easy to play such games without actually exercising. None of the accelerometer systems address the importance of guaranteeing the physical activity that so many studies seem hopeful of. Improved activity monitoring, which has importance in many applications including medical monitoring systems [NAP03, PSM07, SMS05, VAS11], personal training [OF06], and personal gaming [ALM10b, BCF10, MHW11, Pap09, RKH11, SHD09, SGB09b], will be important for exergaming systems.

Exergaming systems have needs for monitoring that span from general human activity monitoring [CLA07, CBC10, HTM11, MHS01, VBV96, ZSF07], activity monitoring for gaming [SGB09b], to monitoring for health, through gaming with targeted health applications (hereafter known as exergaming) [ALM10b, BCF10, GPH09, MHW11, SHD09]. As exergaming expands, its application realm will grow, increasing the need for more immersive and responsive exergaming applications if health and accuracy are the combined goals. In [MCL12] enforcement of activities showed that more detailed movement recognition will allow for more accurate controls and energy expenditure calculations. In both the mobile domain and the stationary gaming environment the freedom to enact any movement in any environment will become increasingly important. These motions will achieve at least moderate levels of physical activity that will help address many of the

chronic medical conditions that are arising due to sedentary behavior [MYO10], though consideration needs to be made into the type of game and activity, to make sure the usage is not dropped due to the levels of activity [LTK12].

As the field expands into more general gaming and even potentially personal training, a wider range of activity monitoring algorithms and methods are needed. Currently, many activity monitoring systems are based on either vision-based systems [BCD02, HTW04, MHK06, RKH11, SFC11] or body-worn sensors [CBC10, FSF99, HLJ08, SKV07]. These tools are used to classify varying activities, from detailed gestures [RKH11, SFC11] to general daily activities [BI04, KNM06, PEK06, RDM05], but not necessarily short, quick, continuous actions as often found in sports environments, which would be needed for sports exergames. General monitoring systems typically classify moves that are cyclical and distinct from each other, for example, related works that classify sitting, standing, climbing stairs, and running [BI04, HLJ08, RDM05]. The data for many of these monitoring systems have two things in common. First, each work has collected its own data and without a common collection of activities, it is difficult to draw comparisons between different works. Second, actions are often long and repetitive with individual samples that are generally between three and five seconds. These are not representative of the short burst continuous actions that often make up sports movements. These class of movements, which will be referred to as Fine-grain activities, can be classified with one short example, with no repetition or cycles. Fine-grain activities require fine, accurate segmentation while the cyclical, repetitive signals captured over longer periods of time do not. Examples of fine-grain activities include, for example, a soccer player that passes a ball and does not immediately have another ball to pass (no repetition). These movements are often related within an application domain with the difficulty often being that these actions come from the same body part with subtle differences. Further,

some of these actions can be standalone movements or beginning of longer, more complex movements. As a result a classification system is needed that can work on fine-grain movements and respond in a rapid enough manner to work for an exergaming application, to improve the physical requirements of exergaming applications, and ultimately, to allow for the design of a system that can guarantee energy expenditure by guaranteeing realistic, detailed motions. Finally, once such a classification system is developed, it must be shown to be adaptable to a number of sports-like applications and give responses in a real-time manner, allowing for the active and continuous control of a video game.

1.1 Research Goals

Algorithm 1: Algorithm for training and using a recognition algorithm in an Exergaming environment

Data: $D(t, C)$ where $D(t, C)$ is a data point at time t with C channels,
$expMove$ is the expected move

Result: $class$, classification label for move performed (if any)

```
begin
    Data ⟵ loadDataSet(); model ⟵ trainModel(Data);
    Buffer ⟵ ∅; missMove ⟵ ∅;
    while Playing Game do
        Add D(t, C) to Buffer; FeatureSet ← extractFeatures(Buffer);
        class ← classifyMove(FeatureSet);
        if expMove ≠ class then
            missMove ← missMove ⋃ class, expMove;
            if shouldUpdate(missMove) then
                Add missMove to Data with correct labels;
                model ← updateModel(model, missMove, Data);
    saveUser(model, Data)
```

The purpose of this dissertation is to investigate the design and implementation of an exergaming system, a framework of which is shown in Figure 1.1, which outlines each of the contributions of such a system. Such an exergaming system requires complex research into many different components working together to create a seamless system that is enjoyable to use and effective at achieving its goals. A truly mobile exergaming system that can recognize a large number of detailed, fine-grain, realistic motions must investigate several aspects, meeting the following required goals:

- Appropriate sensors for recording the appropriate motions
- Developing a game to match the motions required
- Collecting the data necessary to build such a system
- Validating the energy expenditure levels and guaranteeing the activity levels (preventing cheating)
- Developing an algorithm to identify these detailed, fine-grain motions
- Improving the realistic nature of such classification schemes by creating user-centric components, dynamically adjusting such an algorithm to improve the responsiveness of such systems

This dissertation investigates the development of such a system and the required components therein. It proposes the development of a system that achieves each of the above goals, by investigating and integrating a system that segments data, identifies the required duration of such movements. It then takes such movements, classifies them with high accuracy with detailed models that allow for activity recognition and modeling of energy expenditure. Algorithm 1 shows the detailed recognition system necessary for a detailed end-to-end system to work well. Such a system must then be dynamically adjusted to provide a detailed user

experience that responds in a fashion to provide the most accurate, detailed, but also quickly responsive recognition information to allow for active game play simply because, if it takes a system minutes to recognize actions that take seconds to play a game, no matter how accurate, users will not play it. Finally, to develop a successful wireless health system, clinical monitoring and validation are required to guarantee the desired outcomes, and to help the prediction/modeling of the energy expenditure approximated by such accelerometer-based systems. This dissertation investigates each of these components and from them, proposes the development of a single exergaming system that addresses the requirements listed to create a successful wireless health exergaming system.

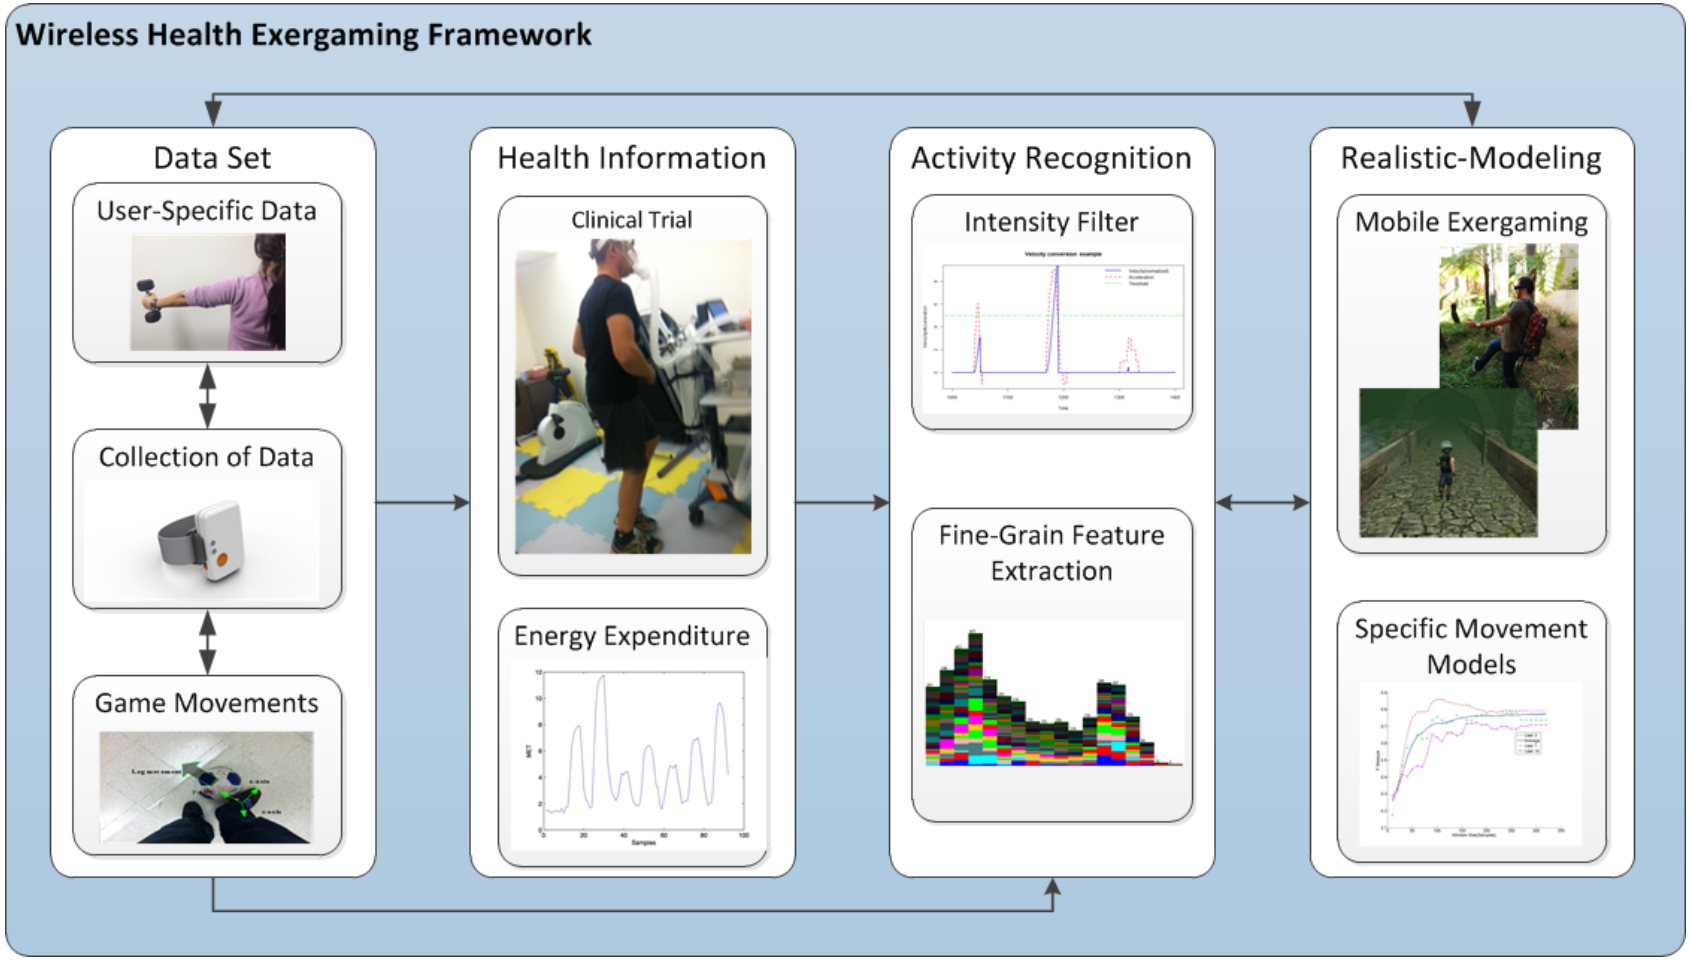

Figure 1.1: Desired exergaming framework

CHAPTER 2

Human Activity Classification and Exergaming

2.1 Manual Selection General Monitoring

Many activity recognition systems manually select features in order to better improve their classification results. In [BI04, CBC10, HLJ08, KNM06, RDM05, VAS11], a form of mean value for the accelerations given is used to determine a key feature to the energy expenditure of the movements. By taking the sum of accelerations, the mean values, or combination thereof in reducing a time series to a singular value, the feature no longer becomes time dependent. Often this leaves out important information about where in the movement certain behaviors occur. No longer can one tell if the foot twists firsts then swings forward or swings forward first then twists, now one can only tell the sum total energy expended. Both actions are the same.

In [BI04, RDM05], a method is presented that uses a combination of the mean of accelerations, the standard deviations, the energy or power expenditure, and a correlation between the channels of the accelerometer. All four of these are time independent values, with the correlation being the final comparison between channels. Features are extracted from windows of size 512 in [BI04] and 256 in [RDM05]. These windows, at their respective sampling rates, correspond to roughly five seconds or greater windows of movements. Movements that fit in much smaller windows could be problematic in such a scenario since their sliding

windows assume there is a cyclical or repetitive pattern to the data to allow for recognition. Finally, the data is not normalized per person. As a result, a factor like the power calculation can have extreme variations when carried out over data sets of a large number of individuals.

In [RDM05], a small data set of eight movements is used to achieve over 90% accuracy while in [BI04] the similar method is used for 20 people and the accuracy falls to 84%. The eight movements collected in [RDM05] are standing, walking, running, climbing up stairs, climbing down stairs, sit-ups, vacuuming, and brushing teeth, classified using mean, power, standard deviation, and correlation [BI04]. Ultimately, these calculations do not properly classify fine-grain movements. This method will be compared to the fine-grain method presented in this paper in the results section.

The results of many related works use accuracy as their primary measure of efficacy [BI04, CBC10, HLJ08, KNM06, RDM05]. Accuracy is not a good measurement for the effectiveness of a classification system because it can be made arbitrarily better by increasing the number of positive or the number of negative examples. [BI04, CBC10, HLJ08] show a confusion matrix and [RDM05] shows a representative confusion matrix, from which precision and recall can be calculated; however, these results are not presented as the primary results. Many works, including [RDM05] pollute their testing set by using the same people in training as in testing; where in real testing, variations across people are what the classifiers really need to be trained for. In [RDM05] four experiments with various set ups are run, the first three of which all contain polluted data. It is from those three settings that they report their highest accuracy, most cases above 80%, with the fourth being considerably lower, most below 64%. Exergaming work presented here will be compared to the method used by Ravi, Dandekar, Mysore,

and Littman [RDM05], hereafter referred to as RDML.

2.2 Commercial Gaming

Nintendo and Microsoft have developed systems that allow for human body movement to control games and have been successful in making exergaming popular [MYO10] [SHD09]. The Wii's handheld accelerometer method for controls, however, can be mimicked with minimal waggling and suffers from a lack of activity enforcement. This paper's system guarantee the activity levels achieved and promote health information for all game types.

Microsoft's Kinect is a camera-based system that allows for the human body to control games, as well, in particular with the use of a popular dance game, which classifies with skeleton animations [RKH11]. However, it does not address the enforcement of such activities. Furthermore, the system is easily affected by lighting changes in the environment and obscuring of body parts, potentially due to other body parts [HTW04, MHK06]. While some of these may be overcome, ultimately the Kinect is a stationary item that does not allow for mobile exergaming.

Another example of a popular video game being used for physical activity is Dance Dance Revolution (DDR), the popular dancing game by Konami. The game requires the use of a dance pad to sense pressure in specific directions and can double the heart rate of the user to exercise levels [Bro06]. Further studies on DDR have shown that promoting playing DDR lead to an increase in vigorous physical activity [MCK08], which further indicates the ability to promote video games that induce physical activity without having to specially tailor the game. Again, however, DDR along with Wii Fit are limited in scope due to their need of a physical pad as the control mechanism, limiting mobility. Further, many of

these devices enable movement but do not enforce it, and while they have shown to be promising for studies do not target any health statistics themselves.

2.3 Vision-Based Models

There exist many vision-based approaches. Microsoft Kinect camera system and SDK for human motion monitoring has proven effective [RKH11, SFC11] at monitoring activity. Skeletal gestures are combined with a decision forest [SFC11] and principal component analysis (PCA) [RKH11] to classify movements. The aforementioned vision-based work is restricted to particular environments, as are those surveyed in [AC97]. In [BCD02, MHK06] studies are compared to show that changes in lighting can adversely affect the performance of vision based models, which is a potential hazard in using cameras in a mobile environment. The goal of this work is to develop a general system that can be used in a mobile setting. Cameras would either need to be in a fixed position or their potential movements, if also on the body, must also be accounted for, further compounding the need for proper tracking via inertial sensors.

2.4 Targetted Exergaming

Systems using accelerometers for gaming, such as [SHD09], use gestures and do not involve continuous movement; they do not promote a moderate to intense physical work out. Other systems [Bro06], promote human activity and health knowledge, such as making a game specifically built to address Type 2 Diabetes [TBB]. The goal in [TBB] was to build a game whose content addressed dietary issues and diabetes understanding, with the intent of making the game fun to

play, but not involve actual movement. Meanwhile, other games [ALM10b], are developed to target specific health applications, like stroke rehabilitation. There have been games that give virtual rewards for being active [BCF10], however, the activity is not a necessity. None enforce the activity. One system does attempt to enforce activity, working with acceleration-based gaming with enforced running in [MHW11]. However, the system only prevents cheating on the running but makes no mention of requiring any of the other movements to have enforcement. The health aspects outputted are also based upon post-processing of data and stride length approximations, leaving room for improvement. The rest of this paper will describe a new system building upon these previous systems, targeting enforcement of activity, and will compare to the system in [MHW11].

Exergaming applications are an emerging set of recognition applications where the human body is used as a controller to video games in order to allow for energy expenditure and healthier life styles. In particular, [ALM10b] develops a game targeting stroke rehabilitation, using accelerometers for active motions of certain body parts in order to play games designed at helping stroke patients with their recovery. They, however, wear more sensors than are ideal; it is difficult to convince most subjects of a new system to attach many sensors to the body. Further, they validate their results from a medical point of view and do not discuss their actual activity recognition algorithm.

Further enhancements to exergaming can be made with the use of sensors by implementing them in mobile environments. Vision-based systems have problems with lighting for classification [BCD02, MHK06], allowing for the added possibility of taking body-worn sensors as an approach to mobile healthy exergaming. Some work has been completed on general mobile exergaming applications [WC08] but these do not clearly indicate any classification or movement types.

Work completed by [MHW11] begins work on fine-grain activity detection for a soccer video game in a healthy environment. However, the paper focuses more on the system description, health benefits, and the end application rather than reporting the accuracy, precision, or recall of their classifier. Further, they limit themselves to a small set of moves, targeted along each axis of the sensor.

Work from [MCL12] looked at the effects of exergaming upon taking a small subset of moves and enforcing the accurate action for those moves. By preventing the cheating of the exercise mechanism in games they show, through a 16 subject user-study, that the soccer exergame feels more realistic, enjoyable, and has an improved sense of exercise and energy expenditure than work in [MHW11] or standard video games. The method, however, uses distance metrics from mean templates of fixed-windowed moves. This classification technique does not extend well to multiple moves and the lack of a gyroscope keeps the work from attacking a significant number of motions that begin with rotations instead of accelerations. However, the accurate enforcement of soccer moves correlated with user experience and energy expenditure calculations is the basis for the work in this paper. By expanding upon the set of moves, and maintaining a level of accuracy via machine learning techniques instead of simple distance metrics, the realistic and healthy exergaming platform envisioned in [MCL12] is improved upon with the work completed in this paper. Ultimately, when increasing the number of potential classes, of which fine-grain movements could have many, the performance of such classification techniques degrades. A system is necessary that can work across many subjects and many movements.

2.5 Mobile Exergaming

Work for mobile exergaming systems in particular focus on the integration of standard exergaming systems onto mobile platforms. Works such as [WC08, GC09, PLL09, KPT10, GOO09, KMK11] validate the ability to place an exergaming system on a mobile phone and increase the areas in which a user can play such games and find exercise. In particular, systems such as [WC08] incorporate physiological data as a bonus onto their mobile system. However, their game does no activity recognition, and game play, while enhanced by an increased heart rate, does not depend on any exercise. In [GC09], a continuation of the [WC08] is built in which the sensors embedded within their phone device assist the physiological data. Namely, the gps and accelerometer can assist pace information in walking to correlate that data with heart beat. Ultimately, neither work expresses any kind of detailed activity recognition and instead simply quantitates whether a person is moving or not.

Further mobile games, such as [KMK11] that wish to address childhood obesity do so by using the mobile phone itself as the sensor. By placing the phone in the pocket users can jump to incorporate jumping into the exergame. While the orientation in which the user jumps is inferred, more complicated motions cannot be drawn from such calculations. Meanwhile, the game application itself only lasts three minutes at a time, hardly long enough to generate exercise unless repeatedly used. The direction of such systems seems to be using the phone as the controller rather as the computing platform, a direction this work plans to take.

CHAPTER 3

Soccer Exergaming and Energy Expenditure

3.1 Desired Outcomes

The first step in the development of a mobile exergaming platform was to target an application domain that was relatively untouched. Work presented focused primarily on arm movement so starting with an application involving the feet was a good idea. Work in this chapter presents the first steps toward an exergaming system targeted at playing FIFA 10, though the methods themselves are generalized for future exergaming platforms. The purpose being to show some form of health benefits can be quantitatively calculated using body-wearable sensors.

3.2 System Description

The body-wearable video game controller allows for physical activity at exercise levels that can be adaptable to a wide range of video games. While sedentary activities, such as video games, have previously been linked to the cause of obesity amongst children and adults, this system will actually turn such activities into entertaining exercise, promoting the health benefits of such games by using the human body as an active controller. The system can use tri-axis accelerometers attached to the hands and feet to monitor movement, pressure sensors to detect standing and running, and uses a software algorithm to further classify and interface with any PC Based game and in this chapter shows work adapted to the FIFA Soccer 10 PC Game.

3.2.1 Hardware

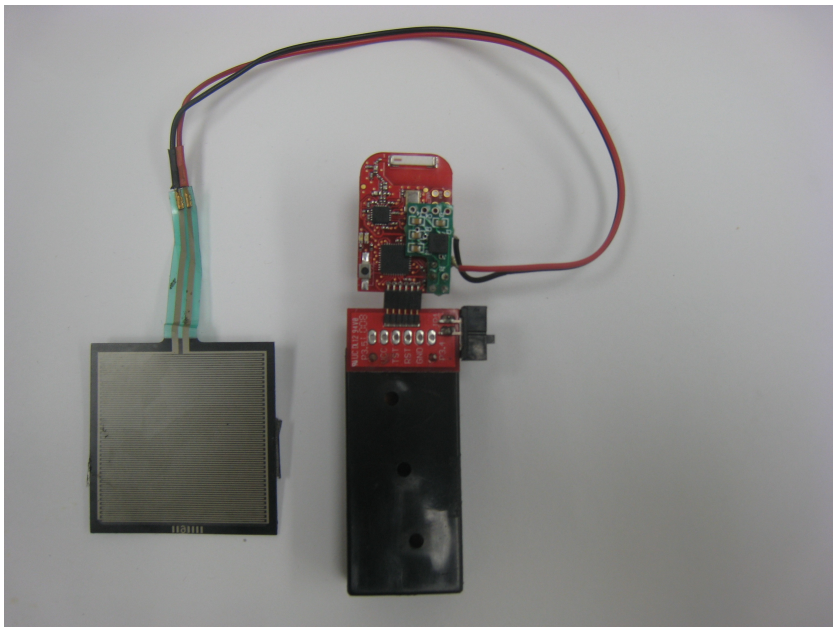

Figure 3.1: One hardware unit with pressure sensor

With the use of low-cost commodity hardware, enacting the desired controls necessary for the soccer game is possible, through the use of two Analog Devices 3-Axis Accelerometers (ADXL335), each attached to an MSP430 Development board by Texas Instruments, shown in Figure 3.1. The MSP430 Development boards allow for wireless communications of the accelerometer readings to the PC. Also attached to the MSP430 are two force sensitive resistor squares for pressure sensing. Each of these is sensors is strapped to a foot, allowing for two footed control of the players within the game. By strapping the accelerometer on the top of the foot, three primary actions need to be detected. The first is a forward shooting motion, in the y-axis direction of the accelerometer. The second, primarily in the x-axis direction, is the passing motion, while the z-axis helps determine strength and active running in place. These are the movements associated with the primary foot, while the secondary foot can determine crossing, through passing, and running in place respectively, thus being able to relate to the primary controls of FIFA Soccer 10. A Nintendo Wii-Remote is used as a

wireless direction pad for easy turning of players. This device is used only to indicate the direction of the running movement generated by the accelerometers and pressure sensors, and can easily be replaced by a gyroscope in future designs. Finally, a software algorithm written in the C# language, developed in Microsoft Visual Studio, helps further classify the continuous actions and interfaces.

3.2.2 Algorithm

The software algorithm consists of four primary components. A flowchart of the algorithm is shown in Figure 3.2. This algorithm detects the trained movements and constantly calculates health information based upon accelerometer+pressure sensor readings and consists of four primary pieces, the calibration, the segmentation, the support vector machine classification, and finally, the determination of action, strength of that action, and the updating of health benefits.

3.2.3 Calibration Phase

The top of a foot, however, is not generally flat and the angle with which the y-axis slopes downward on top of the foot, or rotation around the x-axis, also known as pitch, and the roll, the rotation around the y-axis associated with the side to side slope of a foot, varies from person to person and as such must be accounted for. Since a strap is used to place the device firmly on top of the foot, yaw, or rotation around the z-axis, is not considered a factor in affecting motion analysis, but can be accommodated in a similar fashion as the other two angles of rotation. Figure 3.4 shows the orientation of the sensor and the desired flat directions that the readings will be adjusted to. Since readings from the standard flat position of the accelerometer are known, it is easy to obtain a reading and normalize its values to that of 0 acceleration in the y-direction and 0 acceleration in the x-direction with an acceleration in the z-direction counteracting gravity. Using this as an initial

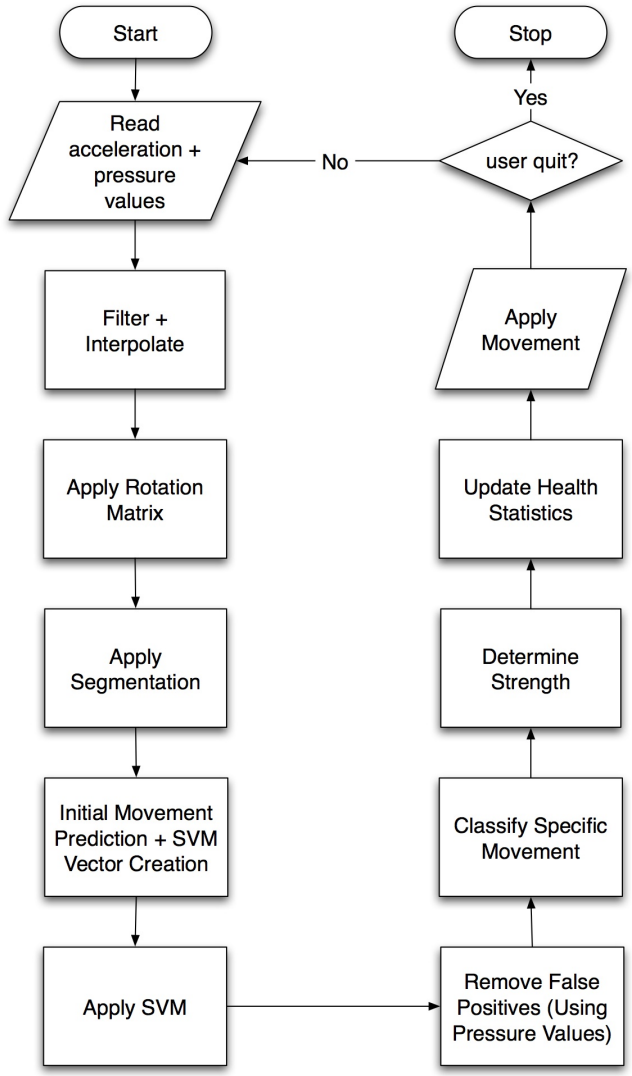

Figure 3.2: Flowchart of the Algorithm.

vector, the algorithm can begin by taking an initial set of readings while asking the user to remain still briefly to determine the acceleration readings do to the position of the foot and associated pull by gravity. Then, a simple rotation matrix is used to convert every 3-point (x, y, z) vector back into the normalized space to more accurately classify a set of movements that no longer becomes dependent on the person and the exact position of the accelerometers. The rotation can be solved via the following equation:

$$\hat{\theta} = 2\pi - \theta \quad , \quad \hat{\gamma} = 2\pi - \gamma \tag{3.1}$$

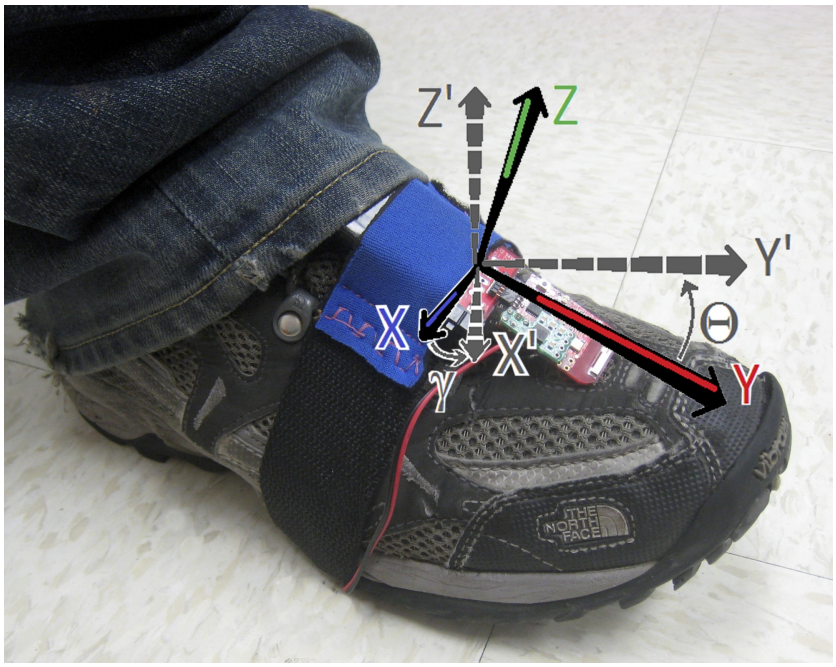

Figure 3.3: Sensor attached to foot. Rotation angle θ corresponds to the pitch of the accelerometer and γ corresponds to the roll

$$A = \begin{pmatrix} 0 & \cos(\hat{\gamma}) & -\sin(\hat{\gamma}) \\ \cos(\hat{\theta}) & -\sin(\hat{\theta})\sin(\hat{\gamma}) & -\sin(\hat{\theta})\cos(\hat{\gamma}) \\ \sin(\hat{\gamma}) & \cos(\hat{\theta})\sin(\hat{\gamma}) & \cos(\hat{\theta})\cos(\hat{\gamma}) \end{pmatrix} \tag{3.2}$$

$$A * \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} \tag{3.3}$$

Once all points are rotated back into the standard plain, the algorithm can move on to its initial classification.

3.2.4 Segmentation + SVM

A segment state machine is built to identify changes in slop of the acceleration values over a sliding window of twenty (20) points. This allows for some history to indicate a trend, but is not big enough to introduce a delay in the actions of the

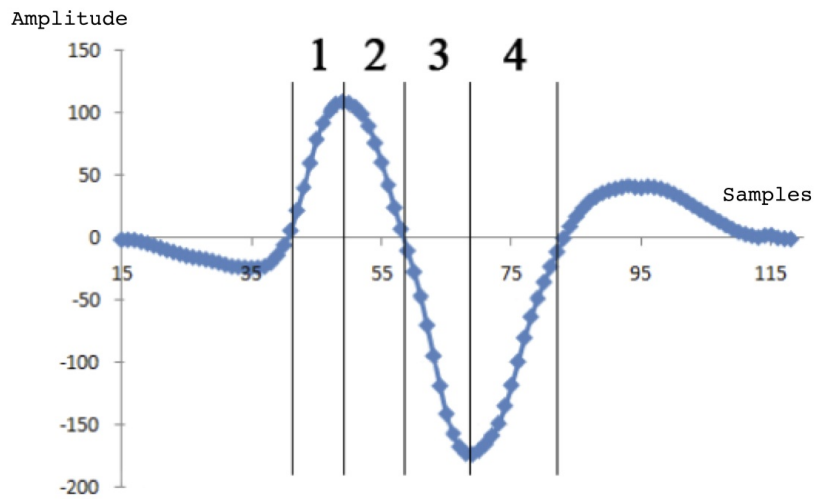

Figure 3.4: Samples of Acceleration in the Y direction for a kick with normalized amplitude. Numbers 1, 2, 3, and 4 correspond to the appropriate segmentation

user. This segmentation, adapted from an initial segmenting system presented in [Hag10] can accurately determine the slopes of acceleration curves and the associated state machine can determine rough potential movements based on a patter of acceleration and deceleration readings from standard kick, pass, cross, through pass, and running. Figure 3.4 shows a kick and how the state machine segments the kick. The state machine itself can classify a large set of actions but for movements that are not entirely dominant in any particular direction, a better pattern-matcher can be used for the boundary cases, like a support vector machine. A support vector machine (SVM), such as the open source libsvm [CL01] that was used, allows for multi-class classification based on extracted features from the segment state machine, namely, the slopes. The SVM, of course, is trained with the movements in a prior phase, but the calibration reduces the need to re-train the segment state machine or the SVM for each user. The final decision comes from the prediction of the state machine and the output of the SVM.

3.2.5 Movement Decision and Health Features

Once the classification is passed on, a final filter is run based upon the pressure profiles and health statistics are updated along with the movement being generated and passed on to the game. The health features are based upon the idea of the Metabolic Equivalent of Task [AHW00], in order to determine the activity level of the user. Different sports have different MET numbers resulting in different calorie burning. MET is expressed versus the cost of resting metabolic rates so it measures specifically based upon the increase in activity, where 1 MET is considered being at rest, while 3 MET is roughly walking and above 6 is more vigorous activity [HAS10]. Based upon the weight the user must input to start the system, and the level of activity and the pedometer activity based on the accelerometer+pedometer the following equation helps calculate the specific number of calories burned by the user:

$$Calories = \frac{MET * 3.5 * w}{200} * t \tag{3.4}$$

where w is the weight in kilograms and t is the total duration in minutes. This calorie information, along with the pedometer information based upon the pressure sensor, give a basic set of health information to the user as he/she is using our system. After the health statistics are updated the movements are decided. Based upon the strength of the movements, which can be calculated by various heuristics and is, in this case, simply the magnitude of the absolute sum of activity according to:

$$\|\overrightarrow{move}\| = \sum_{i=0}^{l} |move_i| \tag{3.5}$$

where $\overrightarrow{move}$ is the movement vector, l is the length of $\overrightarrow{move}$, and $move_i$ is the ith value of the vector. FIFA Soccer 10 allows a wide range of strengths for kicks, crosses, and a differentiation between standard running and sprinting. No movement of the player will occur if the user is not at least stepping in place. Running in place faster will allow the player to sprint, while soft movements of

the legs will generate soft kicks while stronger swings generate harder shots and finally wild movement will overpower the ball and send it out of play losing control. The system is able to continuously monitor not only the movements generated by the body but the strength of those movements in order to more properly determine the activity level MET and the actions within an associated game.

3.2.6 Accelerometer Vibration Issues and Enforcing Activity

Running in place, more so than most features, can cause sensor vibration which even the pattern-matching algorithms like an SVM can mistake for movement or cause it to miss a movement. As a result, the pressure profile of the user becomes a significant portion of the decision making process for movements in order to eliminate false positives as a result of acceleration vibration due to impact on the sensors. The pressure profile must work in association with the accelerometers in order to more accurately detect movement as shown in Figure 3.5. This sensor fusion is a quick and effective way of dealing with a significant problem with accelerometers when strapped to the body. The system, with the addition of the pressure sensors under the feet, is able to eliminate what might otherwise be mistaken as movements in simple accelerometer-only based systems. The important contribution to this system is the ability to control video games with the body in an active and continuous fashion in order to promote healthier game play. As a result, a primary component to this system is, in fact, to enforce that the user be actively playing the game. As the algorithm runs, if at any point the pressure profile decreases from what is expected as a standing person, it will pause and alert the user that he/she is no longer standing and actively playing. The algorithm does allow for the pressure to be lifted but times these intervals of reduced pressure and monitors signals to guarantee actions are occurring and the user is not, instead, attempting to circumvent the activity levels of the system. As a result, implemented is a simple but effective manner to ensure the user is up and

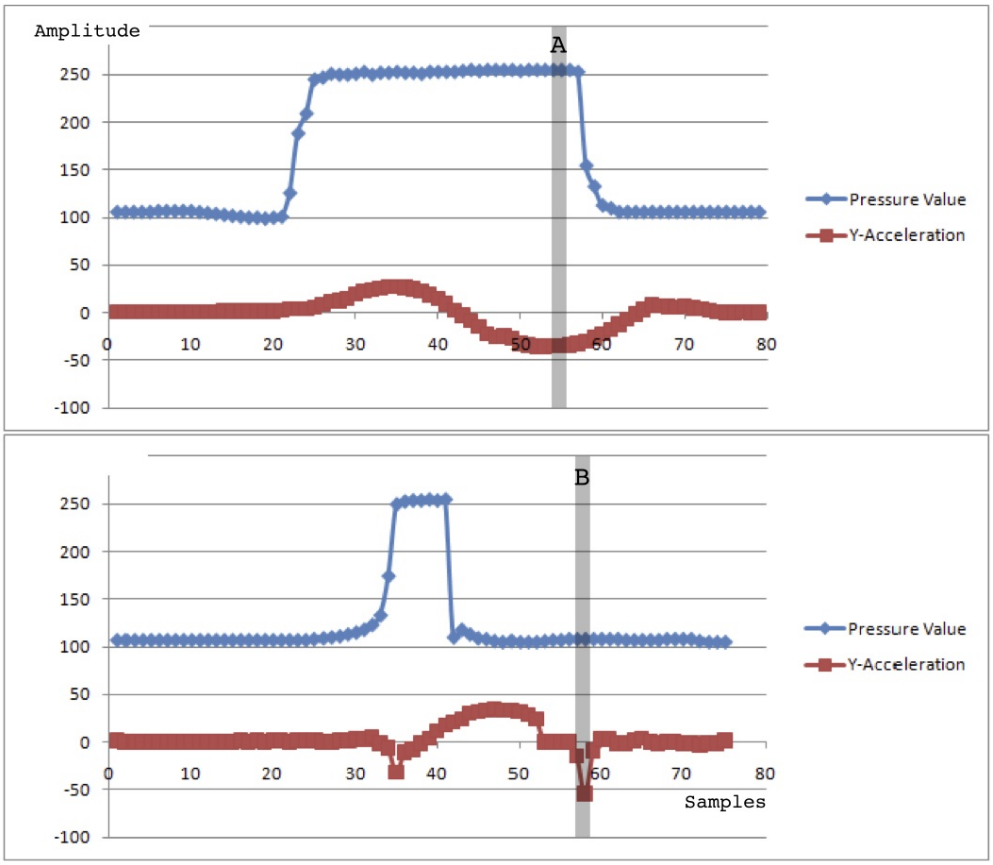

Figure 3.5: Normalized Amplitude of pressure sensor reading and Y-Acceleration. The top graph shows a proper kick at point A while the bottom graph shows an appropriate no kick at point B

about actively moving and, as a result, exercising.

3.3 Health Benefits

FIFA Soccer 10 was a suitable choice for the system because it allows for varying speeds of running and different leg movements for different actions; it has supported over 113 million total online games played [Ele09] and serves as a good example of a popular game that can be made active and healthy. The exergaming FIFA Soccer 10 system allows the user to run, sprint, kick and pass with one foot, and cross and through pass with the other. One can use this system to

Table 3.1: Table showing approximate calorie burn and heart rate (beats per minute) associated with 5 users playing FIFA Soccer 10 on PC with keyboard, with our system, and a comparison to walking 4.5 miles per hour and playing casual soccer [AHW00] over a 15 minute period.

Activity	Heart Rate	MET	Calories
FIFA 10 (PC)	73 bpm	1.0	19
Walking	–	4.5	85.7
FIFA 10 (Active)	144 bpm	4.5	85.7
Soccer	–	7.0	133.3

control any action in the game and can even take it online to challenge others. Table 3.1, which shows the average of five (5) users playing with the system over a fifteen (15) minute period, compares the level of exercise of various activities. Heart rate is based on the specific user, calorie burn is based on equation 3.4 for a 160 lb individual, MET for walking and soccer is a known value [AHW00], and MET for FIFA Soccer 10 PC and the system is based on activity level based on system equation 3.5. As indicated in the table, FIFA Soccer 10 as played with a standard keyboard is completely sedentary, whereas this system performs on par with moderate exercise. As a point of reference, actual soccer is also listed.

CHAPTER 4

Enforcement of Activity

4.1 Desired Outcomes

Work in the previous chapter showed a proof of concept in using a soccer video game as a basis for an exergame. However, beyond running, it did not introduce a significant exercise base. Once it was determined that a soccer exergame could be a viable exergaming solution, it became apparent the need to enforce such activity and improve upon the results was a necessity. Work in this chapter addresses enforcement of activity as a differentiating factor to future exergaming from related work.

4.2 Hardware Sensor

Figure 4.1 shows an example of how it would be worn around a shoe for use as a controller for the popular PC soccer game FIFA10. Much like the previous chapter, the sensor itself is attached to the foot to monitor actions, and a second sensor could be applied to the other foot, but for the goals of this work, only one sensor was applied. This system can be attached to other parts of the body as well but it is ideal that the pressure sensor be extended to a region of the body where it is easy to apply and remove pressure from the device.

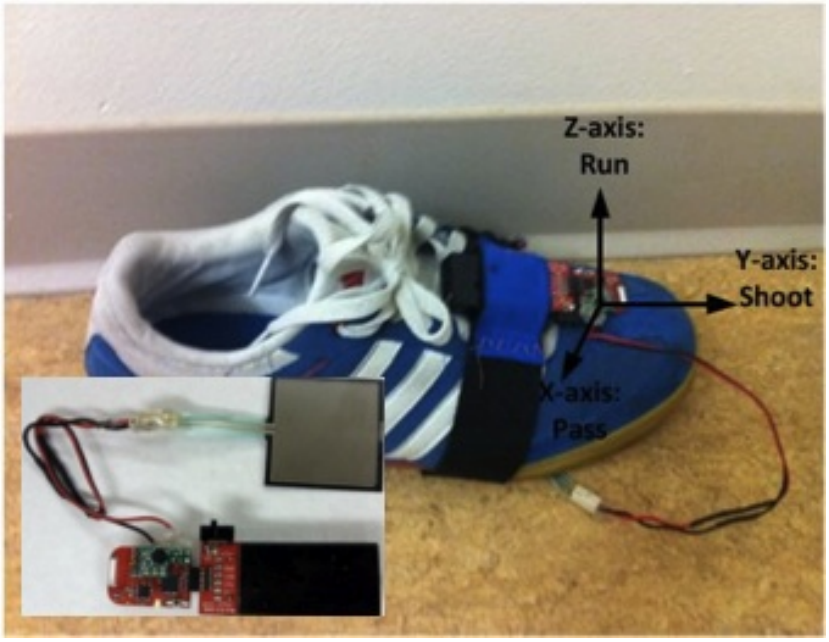

Figure 4.1: Hardware sensor (lower left) attached to shoe for use with FIFA 10, with implied directional axes

4.3 User Interface

Figure 4.2 shows four possible states of motion involved using sensors. Again, FIFA 10 was selected for analysis and for the user-study because of its popularity as a game, soccer's worldwide popularity as a sport, and because it gives us a good testing ground for foot movement within a video game. The users will also have a Nintendo Wii-Remote in their hands in order to facilitate directional changes. This was again chosen for this step of the work over a gyroscope for sensing rotations in order to allow the user to continuously look at the screen and never have to turn away. The directional inputs from the Wii-Remote, however, are only considered a running action comes from the sensors first. The combination of motions and directional inputs are then supplied to the game through a system on the host computer that also calculates health statistics.

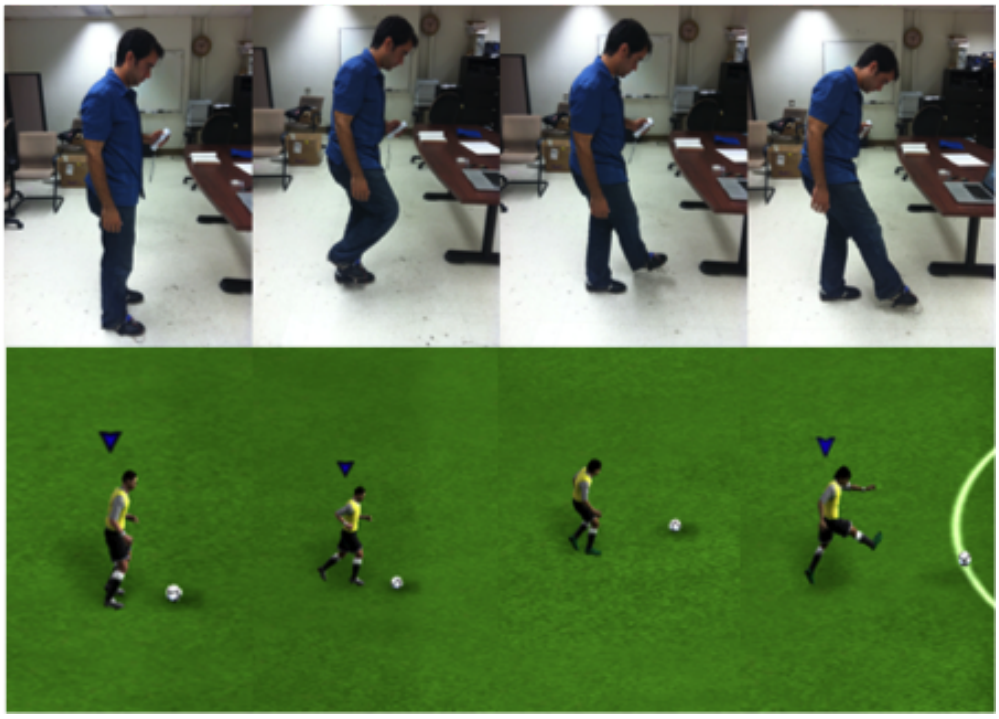

Figure 4.2: Selection of user actions (top) and game response(bottom), left to right: stationary, running, passing, shooting

4.4 Velocity Algorithm

The algorithm is split into two portions. The first portion of the work receives data from the sensor and runs the necessary calibrations and integrations. This data is then passed to the front-end portion, the MET [AHW00, MYO10] online calculator and the movement state machine that determines the actions to be sent to the associated game. This state machine can vary from simple state machines to being replaced by complex classifiers.

4.4.1 Signal Processing and Velocity

Systems that use accelerations only, often lead to the waggle problem, which allows the user to cause spikes in acceleration to mimic movements without actually doing those movements. Even in more advanced systems where accelerations must

meet some defined pattern or template as in the previous chapter [MHW11], acceleration spikes can be created with little or no movement. As a result, velocities can be used to determine the appropriate limb movements. The achieved velocity of a movement not only gives a direct representation of the strength of the movement but also an assurance that the limbs are reaching certain speeds, ensuring exercise-levels of movement.

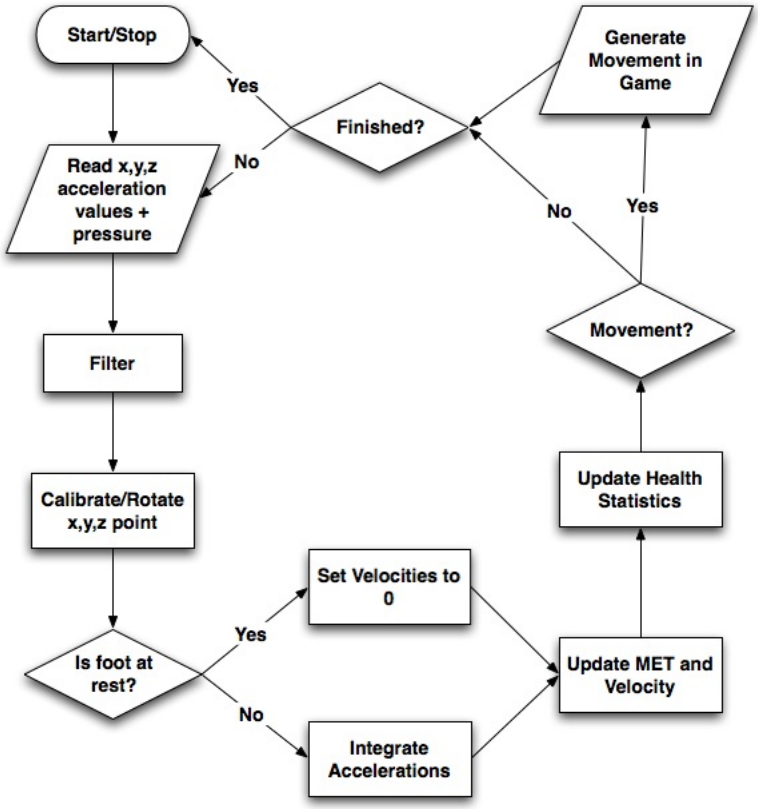

Figure 4.3: Algorithm Loop

In Figure 4.3, the flowchart for the algorithm is shown, which incorporates a low-pass filter. Then there is the calibration method, based upon a rotation matrix of [MHW11], to properly orient a sensor lying on a foot that is not perfectly flat. Once calibrated, the sensor's acceleration readings are integrated to determine the velocity. Velocities determined by this method are often noisy and unreliable. Most notably, it is difficult to determine precisely when the acceler-

ations have returned to a zero value and when velocities should stabilize. Even a little noise in acceleration will be amplified when viewing the velocity values. Without a ground-truth value for zero movement, it is often difficult to properly determine movements by velocities. The pressure sensor, however, can give this information, helping address the drift problem. In the example of the FIFA 10 application, pressure applied to the sensor indicates the user is standing fully on the sensor and one can determine that the velocities should be zero. For situations in which the pressure sensor does not make sense, for example upper arm movements for a basketball game, other options that present themselves include a long period of very low acceleration as indication of noise and no movement or pressing on sensor with the hand.

This static correction provided by the pressure also eliminates unrealistic actions and further improves accuracy. If each movement begins and ends with the foot returning to a static resting position of zero velocity, then the interval window in which each action occurs becomes readily apparent. This eliminates the possibilities of generating multiple movements within one swing of the leg, for example, trying to cheat the system into doing more actions with fewer movements. Further, the Euclidean distance of this windowed move to known move templates can be determined, selecting the minimum distance movement as the appropriate classification. The known move template is the mean move, where

$$\mu = \langle \frac{1}{m}\sum_{i=1}^{m} x_{i1}, \frac{1}{m}\sum_{i=1}^{m} x_{i2}, ..., \frac{1}{m}\sum_{i=1}^{m} x_{iw} \rangle \tag{4.1}$$

is the component-wise mean for each channel of data (e.g., x-axis, y-axis, z-axis accelerometers, and pressure sensor reading), resulting in a w-point vector for each channel of data. Then the Euclidean distance of a move x and this mean template is calculated as

$$d(\mu, x) = \|\mu - x\| = \sqrt{\sum_{i} (\mu_i - x_i)^2} \tag{4.2}$$

A velocity threshold can help guarantee the energy expenditure of a given move and filter false positive moves. Figure 4.4 shows three situations: velocity movement non-detection where an acceleration spike does not generate a large enough velocity, velocity movement based on a large acceleration for a larger period of time, and the final acceleration noise from stepping that does not generate any velocity because the pressure sensor is the ground truth for no movement.

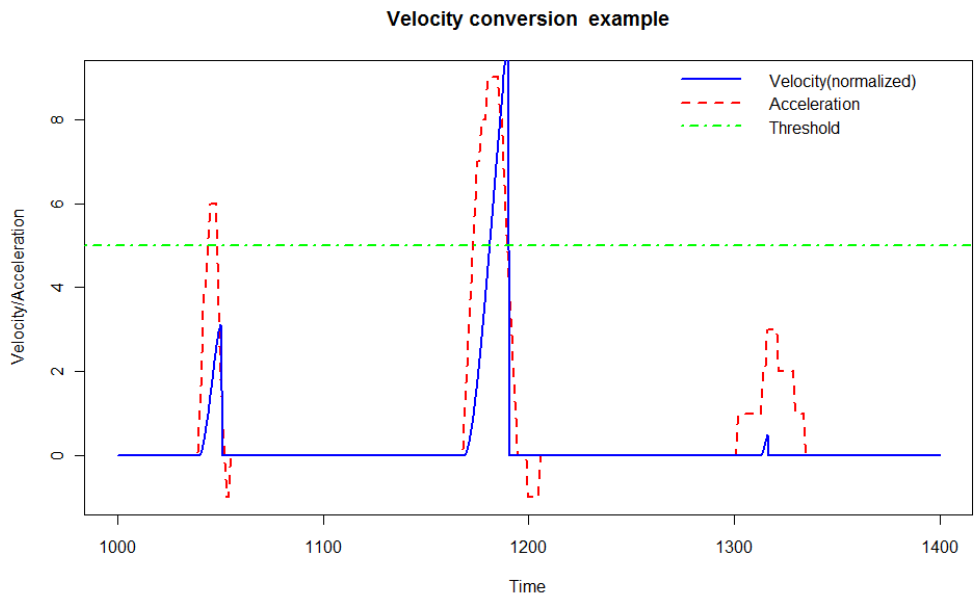

Figure 4.4: Threshold and calculated velocities from accelerations. Normalized velocities shown in m/s while time is listed in samples

As a result, a method for better movement classification and cheating prevention by requiring a certain level of activity is presented. Users cannot mimic actions in small scales with this system, they must put forth the effort, ensuring a higher level of activity than past systems.

4.4.2 Video Game Interface

The velocity algorithm allows us to interface to various video games. For the application to the FIFA 10 Soccer game, it is possible to give the user control

over running and sprinting along with passing, shooting, and lobbing the ball all with varying strengths. These strengths also prevent users from cheating the system by attempting to sit or put in minimal effort. There are numerous sports video games where this application becomes apparent and for each the direction of movements and associated necessary strengths are adjusted on a per application basis. Further, multiple sensors can be worn on the body to compare movement on multiple limbs at the same time, however, FIFA 10's actions lend themselves to a single sensor implementation.

4.5 Online MET Calculator

In order to accurately calculate the amount of exercise the system generates, an integral method to calculate the summation of the integrals with respect to time of the moduli of accelerometer output (IMA) [BKV97] is used and then one can calculate the energy expenditure and MET values from this number. The equation for this formula is

$$IMA = \int_{t_0}^{t_f} |a_x| dt + \int_{t_0}^{t_f} |a_y| dt + \int_{t_0}^{t_f} |a_z| dt \tag{4.3}$$

Work completed in [EPM08, Pap09] shows that placing an accelerometer on various parts of the body will also result in accurate findings. This is particularly helpful in determining the energy expenditure of movements of different limbs for various applications. From this method, each individual activity can have its MET [AHW00, MYO10] calculated online from the values gathered from the accelerometer. This online method is preferable to post-processing methods and comparisons such as the one in [MHW11] because that compares all actions to walking and then uses estimated stride length of actions to estimate MET value; a difficult task while running in-place. This method can find calorie consumption

by the formula presented in equation (4),

$$Calories = \frac{MET * 3.5 * w}{200} * t \tag{4.4}$$

where the mass of the user (in kilograms), the MET value of the exercise, and the total duration of that exercise results in calculated calorie expenditure [MHW11].

4.6 User Study and Validation

In order to validate the accuracy and cheating prevention of the system, a user study of 16 individuals not involved with this project was run. These ranged from undergraduate to graduate students of various majors, ages 21-33. The user study was run on the Electronic Arts game FIFA 10 and chose students with varying levels of experience playing video games and playing actual soccer. Users were asked to rate the system, along with a comparison to the system in the previous chapter [MHW11] and to the commercial version of the game running on a computer with a keyboard interface and/or joystick/gamepad interface. A rating of 1 is the worst while 5 is the best. The order in which the three systems were played was randomized. An example run would start with a user being given five minutes to play the commercial system. After playing the standard game they were trained with the system in [MHW11] for two minutes, and then played that system for five minutes. Afterward they were trained on the velocity system for two minutes, played the game for five minutes.

After the trial, they were given a series of actions randomly generated to mimic in order to test the accuracy of the system and to enhance the mean templates. The useres were asked to repeat the test asking them to fake or cheat the movements by swing their legs in ways they felt could be accepted as valid movements. Users also tried to trick the pressure sensor by leaving their feet in the air for long period of time or play the video game while sitting down. Using

the sensor in hand to cheat was not deemed practical because it is too difficult to apply full weight on the pressure sensor at the correct periods of time. It was deemed acceptable to randomly generate the movements instead of defining a set of actions to test since it would allow for analyzing more realistic aspects of game play. Ultimately the desired outcome was an experience that was closer to actually playing soccer instead of a standard video game.

4.7 Results

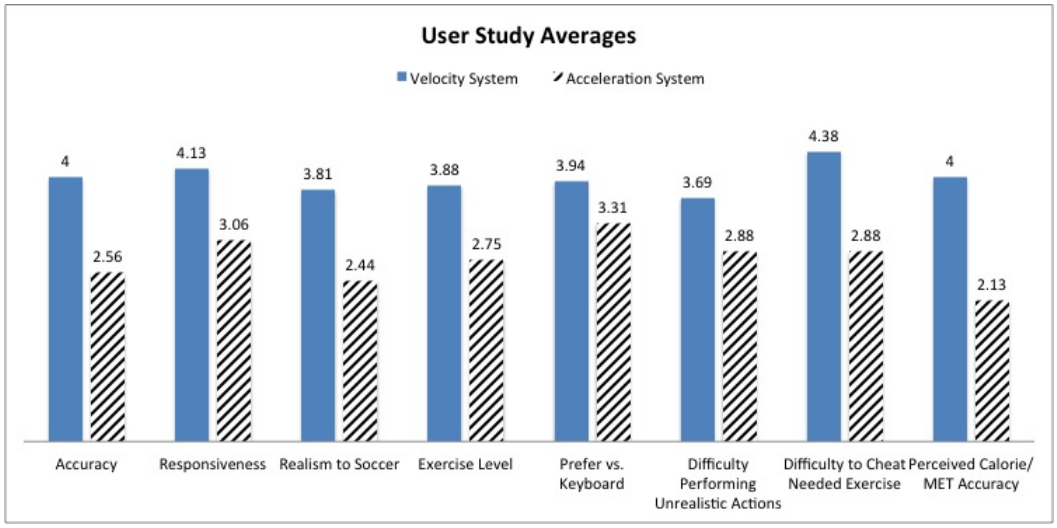

Figure 4.5: Mean values of user study scores. Scores could range from 1 (worst) to 5 (best)

Figure 4.5 shows the results of the user study amongst all of the categories rated. The average score per question was calculated and the results plotted, showing the velocity system has improved on every aspect of exergaming deemed important to measure. An important result in the comparison against the system in [MHW11] is how much more difficult it is to cheat the new velocity algorithm system out of exercising, and the improvement in responsiveness. The results also show that users generally think the system does not allow for any additional unrealistic actions to map into game movements. As a result, the newer system

allows only for the movements desired and requires the users to actually perform those movements in order to burn calories and exercise.

4.7.1 Accuracy

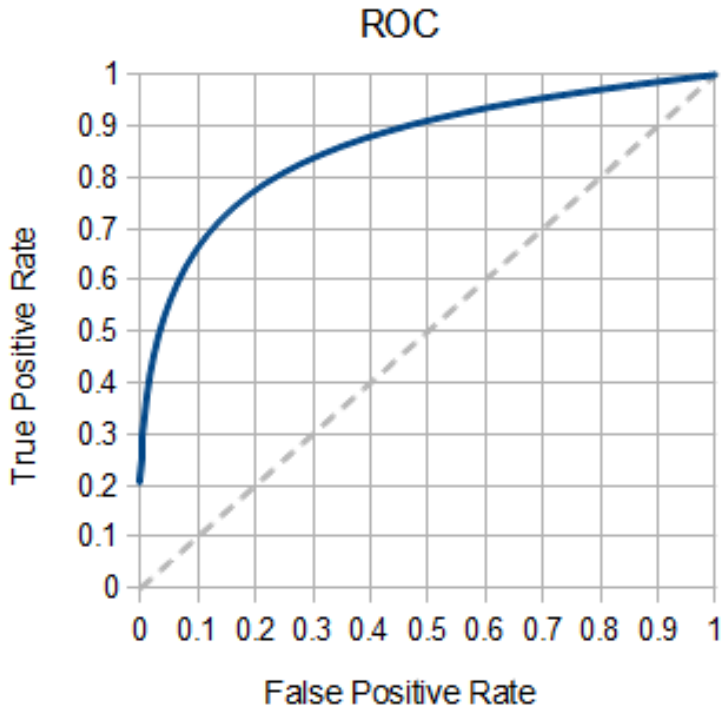

Figure 4.6: ROC Curve for accuracy of actions in FIFA 10

Figure 4.6 is the Receiver Operating Characteristics (ROC) curve of the users in the velocity system. The curve is a plot of threshold for the decision between classified actions. As more actions are classified as any of the movements, the possibility for false positives also increases. As plotted here, the system's test of the accuracy of movements falls in line with the user ratings. 12 of the 16 people in the user study rated the velocity system as either accurate (4) or very accurate (5), the highest two ratings possible.

4.7.2 Preference of System

The adoption of any user interface requires some level of novelty that individuals playing the game would prefer. The work in [MHW11] shows a range of preference between the commercial game and the new user-controlled method, with an average preference score just above 3, meaning somewhat preferred over commercial system. However, in the velocity system, 14 of the 16 users preferred this exergaming system versus the commercial game. None totally preferred the standard interface to using the body to control the game. Since the user interactions seem appealing then the other aspects of the exergaming system can be considered, since it is reasonable to assume a level of adoption in gameplay.

4.7.3 Realism and Exercise

Another goal for these exergaming applications is to mimic realistic sports. In order to obtain an MET for actual soccer, 7.0 [AHW00], the system must require realistic inputs. It is these realistic set of inputs that will then allow for the user to exercise while playing the video games. The realism results show that while the acceleration-based system could be realistic it could also feel too similar to a PC-based game. As mentioned, systems such as [MHW11] can be realistic, but with the focus on enforcement it can also be treated as any other video game, thus the improvement for the velocity system.

Figure 4.7 shows how much users believe they actually exercised. The spread in values rated by the users in the acceleration system comes from how honestly they attempted to play the game. Some continued to exercise while others determined they could lower their level and still play the game. Since the velocity system now requires more enforced activity beyond the basic running, the users agreed that it does make them feel like they are indeed exercising. Candid feedback from the users included mentioning an increased heart rate and actually

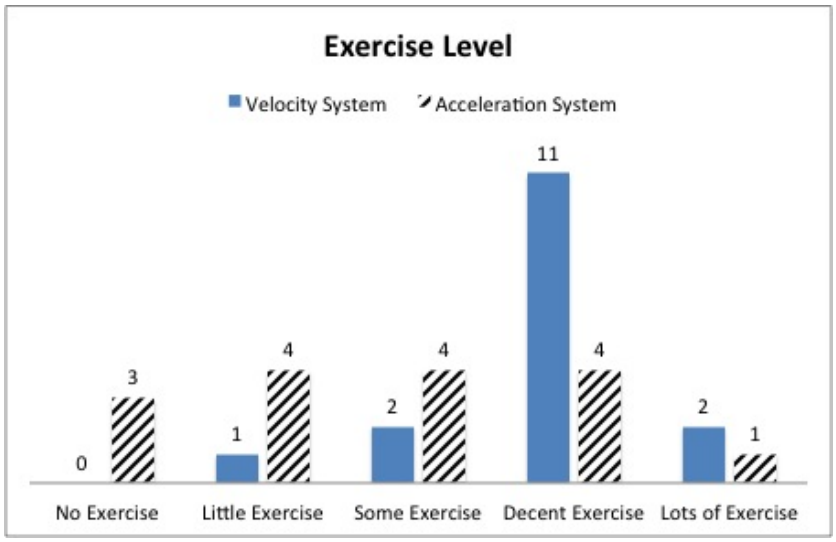

Figure 4.7: Number of users per category on impressions of amount of exercise required

sweating after playing with the sensor system, which led to their preference of this system over [MHW11] and the standard game.

4.7.4 Cheating Prevention

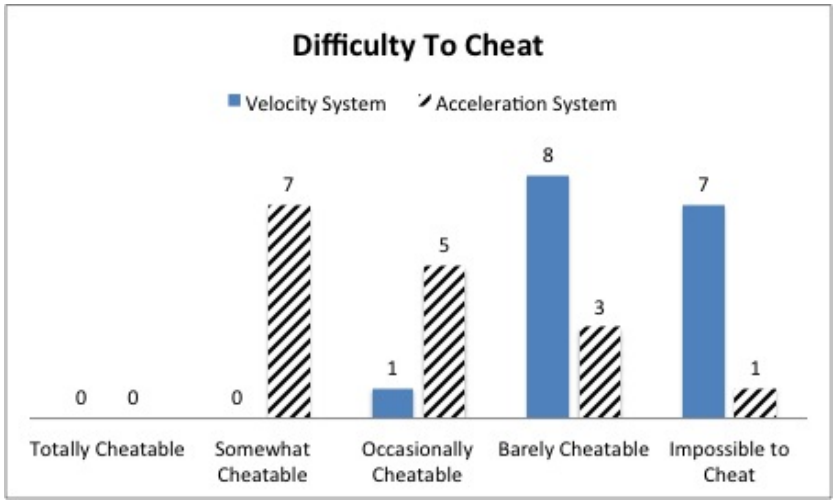

Figure 4.8: Users per rating on ability to cheat necessary exercise

If being able to achieve exercise is the ideal goal, then enforcing the activity and not allowing users to cheat their way out of moves is the necessary corollary to that

goal. Figure 4.7 shows the results of the user study with regards to how difficult it is to cheat the movements required and play the game without exercising. 12 of the 16 users gave the system in [MHW11] a rating of only occasionally cheatable or somewhat cheatable. The velocity-based system is a further improvement to the attempt at enforcement; only one user found that the system was occasionally cheatable.

4.7.5 MET and Calorie Expenditure

The system takes the weight of the user as an input and outputs a set of health statistics from steps taken, to the MET value and calories burned. Figure 4.5 results show the online calculations of the MET based upon the IMA value from [BKV97] was perceived to be significantly more accurate and useful with a rating of 4 than the previous versions with only a 2.13 rating. The system achieved an average MET value similar to best of [MHW11], being 4.5 METs, the equivalent of a run at 4.5 miles per hour, and an average calorie consumption of approximately 30 calories over a 5-minute period. The primary difference is that users genuinely felt they had to expend this much energy playing the game with the new interface.

4.8 Discussion

There are several goals that need to be addressed and accomplished in order to expand this work into a viable set of video games that are indeed healthy for users to play. The first is to run similar user studies with this system applied to other sports games and other games in general with multiple sensors. The second is a comparison to the ability of detecting movements versus a camera system with visual algorithms more complex than those presented in current works as [RKH11]. Finally, the addition of further sensors, such as a gyroscope, and more complex algorithms can potentially improve the number of motions and preciseness needed

to repeat them accurately.

CHAPTER 5

Framework for Mobile Exergaming with Fine-Grain Activity

5.1 Desired Outcomes

Exergaming has emerged as a potentially valuable tool of wireless health to help with regular exercise for healthy individuals [PHL12], collect health information [GKG10], and help as treatments for rehabilitation [APK11]. Obesity [MCL12] [WJN10] [Dal09] is of particular importance,becoming a significant health burden and impacting world-wide economies[FKT12][WA10]. This trend has the potential to increase the obese population by 33% and severely obese population by 130% in the United States alone [FKT12] by 2030, resulting in a population of which at least half are obese [WBL08b]; such a cost, if curbed over the next 20 years, can save the US economy almost $550 billion in medical expenditures. Children are becoming significantly more overweight and obese as a result of sedentary behavior associated with television and video games[SSS04][RVB08][Rob99]. As a result, pervasive sensing technologies to monitor physical activity have become increasingly prevalent[PBW13], particularly through the usage of body-wearable sensor networks [KLH10][LGJ12][CCB06b]. This has lead to the usage of these accelerometer systems particularly to provide input from the human body to video games[WCF07].

Exergaming, through the use of pervasive body wearable sensor or camera-based systems, has been further evaluated to see if it can reduce sedentary be-

havior and increase exercise in children and adults alike, both healthy and obese [Dal09] [SHD09]. Indeed, exergaming achieves light-to-moderate physical activity [PLC11] and impact overweight children [MFM11]. Lately, work has emerged showing a guaranteed level of energy expenditure in exergames, good for healthy adults as well as obese children [MAL13], by monitoring activity levels through accelerometers attached to the body.

As the field expands, the use of mobile technology to assist in the development of exergames has grown [KPT10][PLL09]. This ubiquitous nature of gaming is important to expand the use from both a multiplayer environment as well as an immersion environment. Such games allow for better interaction with others as well as incorporating the real world into the game to increase the enjoyment. These mobile games, however, often use the mobile device for controller input and thus, do not provide a realistic gaming experience. When evaluating what makes a successful exergame, one must consider the elements necessary. Those seem to be the requirements of fast, intense activities [MB11] as well as an enjoyable experience that does its best to mask the exertion through the usage of a multiplayer setting as well as the appropriate set of activities [MEV11].

This work attempts to unify several fragmented exergaming ideas into one system development for successful exergaming work. First, unlike previous work, this paper will present a soccer exergame developed based on movements set in reality instead of mapping moves to preexisting games. This work will first take a realistic sports environment and attempt to translate it into a ubiquitous game platform by collecting an appropriate soccer data set. Further, this work will present a truly pervasive platform for mobile exergaming with near-realistic motions through wearable sensors. The fast, intense exercises will provide caloric expenditures and create a level of enjoyment through the gaming platform. The work will be evaluated from an algorithmic context with regards to the precision and recall of the activity system, as well as present a qualitative result with the

use of a 30-person user experience study. This information is then presented in a structured format to assist in the development of future mobile sports exergames.

5.2 Related Works

5.2.1 Exergames

Several exergaming solutions target full-body motion gaming [MCL12] [PHL12] [GLN12]. Work by Gerling et. al [GLN12] developed a full-body motion-based game with four static gestures and four dynamic gestures. The goal of this work was to develop a fun game for adults based upon the Kinect platform. Such a platform, however, is limited to the environment and region the camera can caption, and as a result, is not a truly mobile system with intense activities. Further, work in [HTW04] surveys Kinect and other camera-based motion tracking and points to potential lighting change problems in certain recognition systems. As a result, this work will use wearable sensing to capture the activities.

Work by Park et. al [PHL12] introduced a competitive exercise based game. Users performed actions on a stationary bike, with a jump rope or a hula-hoop in order to race each other in teams of two. Authors used a motivating scenario of taking standard exercise mechanisms and applying a gaming environment to them to make the games more enjoyable. While this is a great way to encourage adoption of such a platform, ultimately the actions do not correlate to the gaming environment movements. This work builds upon such an idea by taking actions that map directly to games to create a virtual soccer environment based upon a realistic one. Work by Mortazavi et. al [MAL13] guarantees that these actions will achieve moderately intense exercise levels.

Work by Mortazavi et. al [MCL12] involved creating a realistic soccer exergame based upon a gaming platform. Using wearable sensors to create a per-

vasive controller, this work showed creating intense exercise, and guaranteed a level of realistic motions, created an enjoyable interface for existing soccer games. This work builds upon the idea of creating a realistic sports gaming environment. However, instead of mapping to an existing game, this work will delve into creating a new game for an experience that immerses users in the game completely. For example, the soccer exergame presented maps to FIFA 10[MCL12], however, no defensive moves are recognized, most likely due to the inability of users to constantly slide tackle in place. This work will expand upon that as well as develop a more comprehensive survey, but consider the same cheating and near-realistic techniques involved.

5.2.2 Mobile Exergames

Mobile exergames, meanwhile, tend to fall into the realm of being a health game on a mobile platform [GKG10] or games in which the mobile platform is the controlling mechanism for the exergame [KMK11][MR12]. While work by Grimes, Kantroo, and Grinter [GKG10] introduced a new realm in exergaming by developing a ubiquitous, mobile platform, where the gameplay itself does not require any exercise. The other games, which use the mobile device as the controller have limited motions since the location and direction of movement are considered the important input [KMK11][MR12].

Work by Wylie and Coulton [WC08] used wearable sensors to help enable persuasive mobile exergaming. Their work considered measuring the heart rate of a user in order to motivate both running movements in and around locations in order to defend a healthy heart from attacks of viruses. This work incorporates pervasive sensing technologies and realistic running with exergaming. As a result, this work falls short in developing a platform for many kinds of realistic motion-based exergaming, but is one that validates it as an area of research.

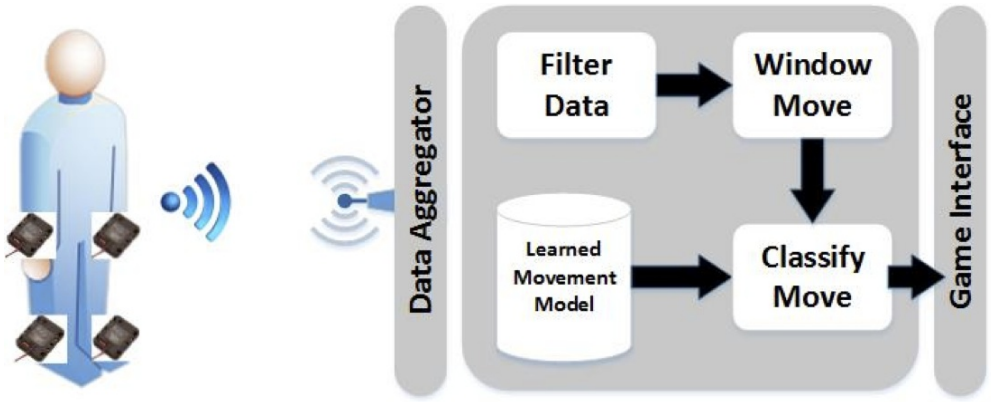

Figure 5.1: Full Mobile Exergame System with Wireless Body Area Network of Sensors with Data Aggregator and Classification System for Gameplay

5.3 System Description

This section describes the methodology involved in developing a tablet-based mobile exergame with wearable body sensors as shown in Figure 5.1. The development of a soccer exergame based upon realistic motions took several steps. The first step was to determine the movements necessary to make a realistic game. For example, when training on an actual soccer field, players are taken through a practice process in which they run from place to place on a field and perform certain actions, such as passing or shooting. Based upon this idea, this work presents a gaming framework based upon realistic sports scenarios in order to develop a directly realistic comparison between gameplay and a real sports environment. In this case, an obstacle course-like soccer game with soccer actions is such an environment. After collecting a comprehensive soccer data set, movements were selected and used for training a recognition algorithm. Then, an appropriate evaluation of such a game is necessary. This evaluation should be two-fold. The first method should analyze the accuracy of such a recognition system, and use this information to generate health statistics. The second is to gather user input on the finished product, since, ultimately, such a game needs to meet certain features

in order to be adopted as an actual exergaming solution [MB11][MEV11].

5.3.1 Data Set and Creation

A full list of soccer movements collected from users is shown in Table 6.2 with all their descriptions, referenced from the right foot. This data was collected using a Memsense Wireless IMU (Inertial Measurement Unit), attached to the right shoe. This device contains a 5g triaxial accelerometer, 600 degree per second gyroscope, and a magnetometer for a total of nine degrees of freedom, and operates at 100 Hz. All nine sensor readings are combined in a single packet of information transmitted at 100 Hz, so that each reading has 9 values per point in time. 24 users collected 10 repetitions of each movement. These movements were selected because they spanned the realm of offensive soccer movements, with multiple types of passes and shots, along with a few trick plays.

A short duration game in which users must run between obstacle courses and then participate in a given action mimics the simulated gameplay work conducted by Mortazavi et. al [MAL13], in which a guaranteed moderate level of physical activity is measured in a simulated soccer exergame situation. With that in mind, a subset of these movements were selected for the game, as they were deemed necessary for a game that would stand in place.

5.3.2 Recognition

In order to develop a game around these movements, a recognition system must be built first. As stated earlier, the system in this work is extended from work by Mortazavi et. al [MCL12] for a soccer exergame. This system uses a nearest neighbor technique for a few movements off of mean templates of movements. Due to the increased number of classes, and the similarity of some of these movements (e.g. the square pass versus the through pass) a gyroscope was added to each

sensor in the system in order to provide more information for each movement.

5.3.2.1 Developing Mean Templates for Training

After all the data was run through a moving average filter (of length 20 points - selected to allow highest classification accuracy), each movement was analyzed in the training set and a window size of 300 points was initially selected to encapsulate each movement. The training set was manually annotated with the midpoint of each move. Thus, any window size that incorporates the midpoint is important and adjustable window sizes are possible. For example, a spin move might take the full 300 points but a pass might only take 120. However, the individual actions have no immediate movement around them so filling the window to 300 points does not alter the activity recognition greatly. Given a particular window size, a vector can be made of each movement in which each channel of data is concatenated together:

$$\mathbf{m} = \langle \mathbf{a_x}, \mathbf{a_y}, \mathbf{a_z}, \mathbf{g_x}, \mathbf{g_y}, \mathbf{g_z}, \parallel \mathbf{a} \parallel, \parallel \mathbf{g} \parallel \rangle \tag{5.1}$$

where each component, such as $\mathbf{a_x}$, is the window-sized time-series vector for that axis of acceleration. A vector for the move is created from each sensor on the body. Finally, the final two channels in this 8-channel signal are the magnitude of the acceleration and the magnitude of the gyroscope, calculated as:

$$\parallel \mathbf{a} \parallel = \sqrt{\mathbf{a_x}^2 + \mathbf{a_y}^2 + \mathbf{a_z}^2} \tag{5.2}$$

$$\parallel \mathbf{g} \parallel = \sqrt{\mathbf{g_x}^2 + \mathbf{g_y}^2 + \mathbf{g_z}^2} \tag{5.3}$$

Once an individual move m is arranged, each training example is ordered and a mean-template for each move is created,

$$\mu = \langle \frac{1}{n} \sum_{i=1}^{n} \mathbf{a_{xi}}, \frac{1}{n} \sum_{i=1}^{n} \mathbf{a_{yi}}, ... \rangle \tag{5.4}$$

where each component in the vector is the associated index in vector $\mathbf{m}$ and each i is a training example from the training set. For example, if the moves are 300

points by 8 channels, μ is a 300 point, 8 channel signal averaged over the range of i training examples. Thus, the training set becomes:

$$Training = \{\Theta_j | j = 1, ..., 15\} \tag{5.5}$$

where each j represents one of the fifteen moves in this training set, and:

$$\Theta_j = \{\mu_1, \mu_2, ..., \mu_s\} \tag{5.6}$$

Θ_j represents the set of means for an individual move across each of the s sensors worn on the body (e.g., for example s could be 4, for a sensor on each limb).

5.3.2.2 Classification Algorithm

Once the training set is created, the online classification algorithm runs. This is based off of a nearest neighbor algorithm in which each move is evaluated and compared against the average movements stored. A sliding window of size w is set, where w is the same size used in the training set above. The window overlaps point by point. This is because a euclidean distance will result in a large gap until the windows line up well. This allows for the move to center nicely in the window to be classified correctly. Two factors in classification change from work by Mortazavi et. al [MCL12]. First, with the addition of the gyroscope, the velocity comparison is removed, primarily to reduce computation time in the mobile environment. Further, work from Mortazavi et. al [MAL13] shows the energy expenditure calculations for each movement in a soccer exergame provides a formula to guarantee intensity. This calculation was used for both health information and for the cheating prevention that was velocity-based in previous work. Once the sliding window adjusts for the new point, the euclidean distance between the point and the mean-templates are calculated, and a class is selected:

$$d(\mathbf{m}, \mu) = \|\mathbf{m} - \mu\| = \sqrt{\sum_w (m_w - \mu_w)^2} \tag{5.7}$$

and then an overall distance value D is calculated as a weighted sum of the distances between each sensor template as in:

$$D = \sum_{i}^{k} p_i \times d(\mathbf{m}, \mu_\mathbf{i}) \tag{5.8}$$

where the weights, p_i, are adjustable if particular limbs are more important for some movements over the others. Finally, with this overall distance value, a class can be selected where:

$$c(m) = \{class(\Theta_j)|\Theta_j = argmin_j\{D(m, \Theta_j)\}, D < \tau\} \tag{5.9}$$

where c is the class of template Θ_i, $D(m, \Theta_i)$ is the euclidean distance between each sensor of move m and each sensor mean μ_i of the template move Θ_i that results in the minimum distance. This distance must be smaller than a predefined threshold τ or the movement is simply considered as no class. This τ is adjustable based upon the distance and the energy expenditure calculated [MAL13], and can be altered given the intensity desired in the final application. In a system with multiple sensors, each sensor runs such a classification. A majority voting scheme is used to determine the classification result and in the event of a tie the foot sensor on the right foot makes the ultimate decision as the game in this paper is primarily based upon movements of the right foot.

5.3.3 Improvement with PCA

As the number of movements and classes increases, a more robust algorithm is needed aside from the mean templating presented. An obvious expansion from the mean templating is to use those templates and determine the eigenvectors and eigenvalues of the system and run a principal component analysis. Then, each move window is decomposed and reconstructed with the top eigenvectors and the minimum reconstruction error results in the chosen class.

5.3.3.1 Model Generation

A window around each of the channels of data is selected to create a fixed-sized move similar to above. Then, the training system calculates the mean for each of the channels. If m moves are represented as c-channel vectors, then the mean move is a component-wise mean, as in:

$$\mu = \left(\frac{1}{m} \sum_{i=1}^{m} \mathbf{x}_{1i}, \frac{1}{m} \sum_{i=1}^{m} \mathbf{x}_{2i}, ..., \frac{1}{m} \sum_{i=1}^{m} \mathbf{x}_{ci} \right) \tag{5.10}$$

where the sum of the $\mathbf{x}_j$ are the sums of window size w vectors $\mathbf{x}_j$, where, in this system $\mathbf{x_1} = \mathbf{a_x}$, $\mathbf{x_1} = \mathbf{a_y}$ and so on. The mean template, μ, is a c-channel signal of the same window size, w, as each of the individual moves. The principal component analysis (PCA) of each move is run, using the sorted order of eigenvalues to find the most significant eigenvectors. Thus, the training set consists of a mean move μ and a set of eigenvectors e_i per μ, across all moves.

5.3.3.2 Activity Recognition

The models generated are made over various window sizes and are stored. Each is then used during the classification processes. To classify a move, a w-sized test window is decomposed and reconstructed using the top k eigenvectors and the mean μ for each of the training moves. The number of eigenvectors can be chosen based upon performance and speed, and as such, 10 are chosen for this particular application. The Euclidean distance between the original test move, $\mathbf{x}$, and the reconstructed move, $\mathbf{r}$ is calculated as

$$err\left(\mathbf{x}, \mathbf{r}\left(\mu\right)\right) = \left\|\mathbf{x} - \mathbf{r}\left(\mu\right)\right\| = \sqrt{\sum_{i} \left(x_i - r_i\left(\mu\right)\right)^2} \tag{5.11}$$

where i represents each channel of the data, (e.g., eight in this case), and $\mathbf{r}$ is a

Figure 5.2: User from trial playing the mobile soccer exergame (left)

function of μ and its associated eigenvectors. This reconstruction error is computed over all of the training means μ and their associated eigenvectors. To classify a move $\mathbf{x}$ of type C based upon the minimum error, the class with minimum reconstruction error is chosen, as in

$$C\left(x\right) = \left\{class\left(\mu\right)|\mu = argmin_{\mu}\left\{err\left(x, r\left(\mu\right)\right)\right\}\right\} \tag{5.12}$$

If a single μ is returned, then simply classify move x as the same class as the mean move μ. In the event that a subset of two or more classes of moves results in the same minimum reconstruction error, then the classifier randomly chooses the class from this subset of moves.

5.3.4 Implementing on Mobile Device

The game was developed as a side-scrolling field. As the user runs, he/she approaches each action node, a prompt appears to complete the action, as seen in Figure 5.2. As the action is completed, the field will scroll and the next movement indicated, with the current move indication disappearing. The field progression and movement indicators are how the user knows that the actions are being performed. Based upon how well the users are able to complete the actions they are giving a final time score, intended to encourage competition. The game was developed on the Android development platform for a Nexus 7 tablet (quad-core with 1 GB of ram). Several computational adjustments needed to be made in order to guarantee a real-time and intense experience. A trade off between multi-threading the classification and running each sequentially must be analyzed on a case by case basis to avoid delay. Too many threads causes too much context switching, particularly because the bluetooth-enabled sensor is transmitting at 100Hz. Too few and the number of moves must be reduced. This is ultimately how the limit of 15 moves was selected for this game. The game was targeted at being around 3 minutes per round, so that it could be played quickly if needed, or repeated to make a longer gaming experience. Running between each action guaranteed around 7 metabolic equivalents[MAL13][AHH11] of energy expenditure, being light to moderate intensity activity.

5.3.5 Survey Creation

In order to evaluate the success of such a game, three aspects must be analyzed. The first is the recognition algorithm, the second is the measurement of caloric expenditure, and the third is a user experience trial to determine if such a game can achieve the tasks desired. However, determining the appropriate requirements for a user survey might not be immediately obvious. Several works

discuss the need to have the right set of questions depending on the goal. These works consider the appropriate techniques to evaluate user experience for gaming environments[VLR10][SVS12], namely:

- Multiplayer approach[Nac09]
- Fast, intense, accurate motions [MB11]
- Encouragement of physical activity [BIM08][BI12]
- Focusing on details of gameplay [BIM08]
- Realism and cheating prevention of exergames [MCL12]
- Health Information [MCL12]
- Comparison to other games and general enjoyment [MCL12][SVS12]

In that light, this work presents ExerSurvey; an initial attempt at developing a comprehensive, unifying survey for user experience trials for future exergame work and user experience research. This survey attempts to address each category, where each category is generally ranked from 1 to 5, where 1 would be very bad and 5 very good. First, general statistics are collected regarding whether they have played video games, mobile video games, the sports, or those video games related to the sport. Then, how often the users exercise, to gauge their general level of physical activity. Questions on each related category (the scoring for those not immediately obvious will be listed after the question) are:

- **Enjoyment of Games**
  - Q1) Enjoyment of Comparison Games
  - Q2) Enjoyment of Mobile Exergame
  - Q3) Enjoyment of Exergame Over Comparison Game (1. great prefer standard game, 2. prefer standard game 3. neutral 4. prefer exergame, 5. greatly prefer exergame)
  - Q4) Enjoyment of Exergame over Motion Input to Comparison Games (1. great prefer mobile motions, 2. prefer mobile motions 3. neutral 4. prefer exergame, 5. greatly prefer exergame)

- **Statistics Regarding Exergame**
  - Q5) Perceived Accuracy of Exergame
  - Q6) Level of Fatigue (1. great fatigued, 2. fatigued 3. neutral 4. rested 5. completely rested)
  - Q7) Realism of Exercise
  - Q8) Ability to Cheat Game (1. Don't have to perform actions 2. Easy to avoid most actions 3. neutral 4. almost impossible to cheat 5. absolutely impossible to cheat)
- **Enjoyment of Games**
  - Q9) How often per week would you play standard mobile games?
  - Q10) How often per week would you play mobile exergames?
  - Q11) How did you perceive the duration of the game (1. way too short 2. too short 3. about right 4. too long 5. way too long)?
- **Multiplayer Aspects**
  - Q12) How well does the scoring motivate you?
  - Q13) How do the multiplayer aspects of the game affect your opinion?
- **System Goals**
  - Q14) How helpful was the health information provided?
  - Q15) How did you feel about the sensor system?
  - Q16) How would you rate the mobility of the game?
  - Q17) How would you rate the intensity of activity reached?
  - Q18) How well did the exergame achieve its intended goals?

5.4 Results

The mobile soccer exergame developed was evaluated on 30 users, 6 females and 24 males, ages ranging from 18 to 58, with varying levels of experience in video

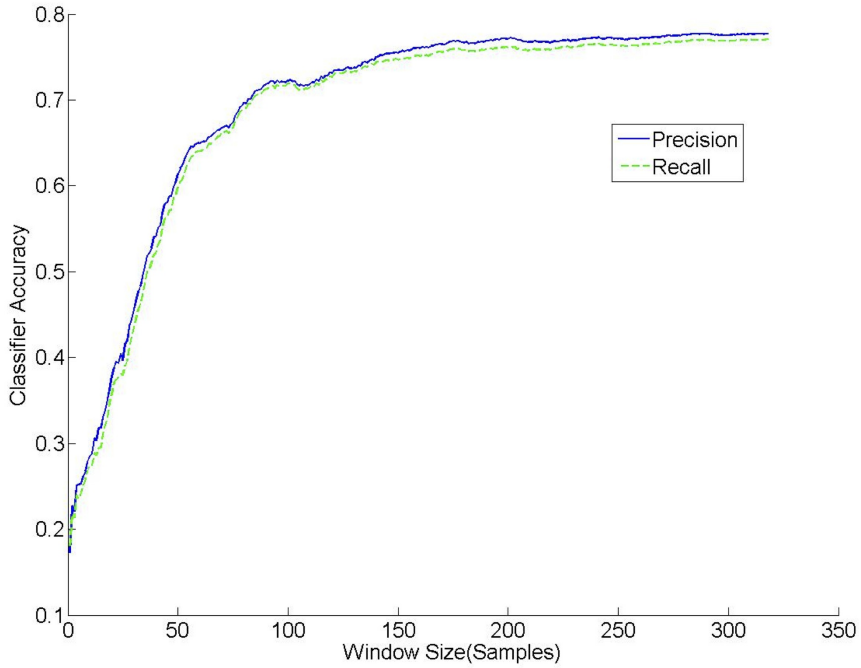

Figure 5.3: Accuracy vs. Window Size

games, soccer, and mobile games. One sensor was used, and after attaching the sensor to the user's right foot, users were given three games to play. Users were shown all the movements and then were given a minute to practice and ask questions before playing the games. Those games were EA Sports' FIFA 10 on the tablet played with the simulated joystick, played with tablet controls (using the accelerometer for running by tilting the tablet), and the soccer exergame developed in this paper with PCA as its classification method. The order in which they played the game was randomized. They were allowed to play the game in any environment they desired (e.g., a lab, a park, an apartment, etc.). Each half of FIFA was set to four minutes and the soccer exergame took about three to four minutes to complete, depending on the user's abilities. After playing each game, the users were given the survey. The short duration of the game, while not ideal for prolonged exercise, mimics many user's behaviors with mobile games

that can be played in any environment with any break during his/her day or repeated if desired. Health statistics were calculated by the method in [MAL13], were validated there.

5.4.1 Classification Results

A leave-one-person-out cross validation was used on the training set in order to test the accuracy of the algorithm, for both the mean template and for PCA. It was compared with a common general daily-activity monitoring system, that extracts from a movement window the mean, standard deviation, power, and correlation across each axis and supplies this information to a support vector machine, as in work by Ravi et. al [RDM05]; the classification results are shown in Table 5.2. The results show that the algorithm with the gyroscope added is very strong for the movements used in Mortazavi et. al [MCL12] and while it diminishes a bit on the full set of movements selected, it significantly outperforms the general daily-living activity system, showing the need for gaming movement-specific classification algorithms. As seen in Figure 5.3, the precision and recall reaches a high level a little over half-way through the window. Since each move ranged between 80 and 300 points, the window size was dropped to 200 points. This is because the last 100 points of activity are generally the user's foot returning back to the ground, and so, do not provide a significant classification difference; however, the reduced window size gives an extra second with which to compute the classification and give the user more responsive behavior. This result can be measured in the perceived accuracy, realism, and cheatability of the game as measured in the survey.

5.4.2 Survey Results

The entire user survey is shown in Table 5.3 and the results are generally positive. The users felt the system was accurate, realistic, generated exercise, and was hard to cheat. In particular, the categories under the scores of 3 all still show signs of being appropriate given the scenarios presented. For question 1, users did not seem to like the control scheme for FIFA on the tablet. For question 6, the rating just under 3 indicates users feeling slightly fatigued when playing the game, a desirable outcome (the survey should be corrected to make this a score of 4, a desirable outcome). Question 10, while indicating users would likely only play this game a few times a week, has a rating higher than Question 9, which indicates how often they would play FIFA in a mobile setting. Question 11 indicates the game should be extended in length, a simple parameter to adjust. Question 12 shows that the timing score did not motivate users, showing the need for a better feedback mechanism; however, that same score still drove multiplayer competitiveness, as shown in question 13. Finally, question 14 was listed as not-applicable due to its addition to the survey after the trial had begun. Wii and Kinect FIFA were not used because the game on each console uses standard controls.

5.5 Discussion of User Experience Trial

The mean results of most questions were 3 or higher. In particular, only several questions had a mean score of 2.5 or below. As a result, the survey shows a fairly consistent experience, with obvious room for improvement. In particular, the results with lower mean and higher standard deviation related to the classic FIFA game, the lack of fatigue and short duration of the exergame. The former is a beneficial result while the latter two show the need to build a more complex game, as will be discussed in the future work section.

5.5.1 Enjoyment of Games

First and foremost, users must enjoy a game, otherwise it will not be adopted. As already demonstrated in Table 5.3, users enjoyed the exergame. Those that had a higher experience playing soccer exergames were more demanding; however, they still thoroughly enjoyed it. As seen in Figure 5.4a, most users either enjoyed both games or thoroughly enjoyed the exergame. The few users who did prefer FIFA all indicated that they played some form of the mobile FIFA game on a consistent basis. However, when asked if they preferred the exergame's motions, the survey results were generally stronger, with users indicating that they preferred using their body to tablet motion for movement. Participant P4, for example, who considers himself an experienced video gameplayer and soccer player, indicated that

> "The Exergame was better than FIFA. In particular it motivates movement and physical activity (especially using the foot) is very exciting."

P10, a 32 year old male who plays a lot of video games, soccer video games, and mobile games but is not an actual soccer player, and states

> "The movements of the exergame were more applicable to actually playing. Moving the tablet just seemed like an awkward way to control players."

5.5.2 Exergaming Accuracy

Figure 5.4b shows that the system was deemed generally fairly accurate, though a little more neutral than one would desire. One can infer from the score question (question 12) having a mean under 3 that the system did not produce enough feedback with regards to the specific actions. Further, it became apparent that the cheating prevention aspects of the activities, the specificity and intensity required

to perform the actions realistically, were hampering the perceived accuracy of the system. As shown in Figure 5.4c, the system was generally considered near impossible to cheat and the actions very realistic. Future exergame work must do a better job providing direct feedback to each action, so that the user knows whether the system is inaccurate or if she/he is performing the action incorrectly, as was often the case with this game.

5.5.3 Multiplayer Aspects

As seen in Figure 5.5a, the users did not greatly enjoy the feedback from the system but still deemed it enough to compare between users and develop a competitive aspect to the game. Participant P21, a male aged 39, stated simply

> "Beating other people is the only reason I play games."

5.5.4 System Design

Figure 5.5b shows the user perception on mobility of the system, deemed very mobile. However, the perception of the sensing device itself saw some negative feedback. In particular, certain types of shoes did not lend themselves to attaching the sensor properly. Further, as participant P14 indicated

> "It wouldn't work barefoot,"

such a system requires a certain type of clothing as is currently implemented. Figure 5.5c shows that this irritated a few users. As a result, future generations of games should consider even less invasive sensing technology.

5.5.5 Future Developments

While this work has presented the development of a truly virtual sports environment game with pervasive sensing, it opens the door to significantly more

work. A more complex game can be designed, that can last longer and present more information to the user through graphical means, be it health information or graphics to better immerse the user in the activities. Competitive aspects to such a game can be expanded upon, as can the multiplayer aspects and feedback. Further, these movements can take further advantage of the mobility and open environment by increasing the range of movements and the sensors embedded and incorporated. As a larger user base adopts such a game, further questions on the classification algorithm itself must be addressed. These include questions regarding the modeling of large multiclass problems, skill levels, and better ways of combining data from various sensors. Once these goals are achieved, longer studies must be conducted, as is suggested by [WJN10], for long-term adoption clinical benefits.

5.6 Discussion

This paper has introduced a unique development approach to truly mobile exergames with near-realistic activity recognition. This work presents a merger of ideas to present a framework for a truly mobile, ubiquitous, and pervasive gaming platform using wearable sensors to monitor human activity. By first creating a recognition system for realistic soccer actions, as well as developing a game to map to these actions to create a realistic virtual environment, an intense and enjoyable gaming experience was created. Finally ExerSurvey was employed in order to present a unified validation approach to effectiveness of the exergame. The survey, as well as the activity data presented in this work will be made available for future use in order to enable an acceleration of research in mobile active exergaming applications based upon realistic sports environments and other similar ideas to help further gaming and healthy exercise, as well as address potentially devastating epidemics such as childhood obesity.

Movement	**Description**	Approx. Duration
*Back Heel	Pass backward	.5s
*Behind Foot Pass	Pass left behind left leg	.75s
*Square Pass	Pass left (standard)	.75s
*Through Pass	Lead pass forward and left	1s
*Flick Pass	Quick pass with outside of foot	.5s
*Chip/Lob	cross a ball forward and left	1.5s
*Fake Shot	Fake shot (only backswing)	.75s
*Full-Swing Shot	Full powered swinging strike	2s
*Medium Powered Shot	Strike ball with laces	1.5s
*Quick Shot	Toe poke quick release	.5s
*Curved Shot	Placed shot, using side of foot	1.5s
Cut Left	Running, step right cut left	2.5s
Cut Right	Running, step left cut right	2.5s
*Side Step	Step on ball, roll it to right	1.5s
Spin Move	360 degree spin	2.5s
Step on Ball	Stopping ball in place	.5s
*Step Over Ball	swing leg around ball	1s
*Run	Running in place	1s
*Sprint	Sprinting in place	.5s
Walk	Stepping forward	1s

Table 5.1: Full-list of soccer movements, those marked by * retained in the final system

Algorithm	**Subset of moves[MCL12]**	**Full set of moves**
PCA	89%, 89%, 89	77%, 77%, 77
Soccer [MCL12]	81%, 80%, 80	63%, 61%, 62
General Activity[RDM05]	66%, 41%, 51	40%, 19%, 26

Table 5.2: Precision, Recall, and F-Score respectively of Mobile Soccer Exergaming System using a subset of movements and full set versus a general-activity monitoring system [RDM05]

Q. 1	Q. 2	Q. 3	Q. 4	Q. 5	Q. 6
2.97 ± 1.16	3.50 ± 0.90	3.60 ± 1.04	3.93 ± 0.98	3.47 ± 0.94	2.90 ± 0.71

Q. 7	Q. 8	Q. 9	Q. 10	Q. 11	Q. 12
3.67 ± 0.55	4.13 ± 0.82	2.23 ± 1.19	2.40 ± 1.07	2.77 ± 0.86	2.83 ± 1.09

Q. 13	Q. 14	Q. 15	Q. 16	Q. 17	Q. 18
3.77 ± 0.77	N/A	3.27 ± 0.87	4.13 ± 0.73	3.13 ± 0.82	3.50 ± 0.82

Table 5.3: Average scores and standard deviations for the 18 question ExerSurvey given to all 30 participants

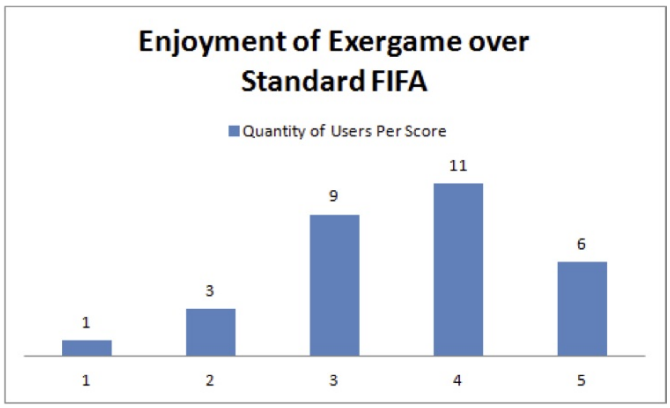

(a) Preference

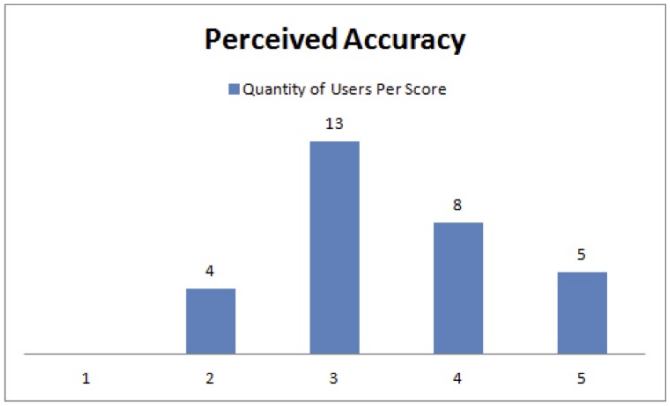

(b) Accuracy

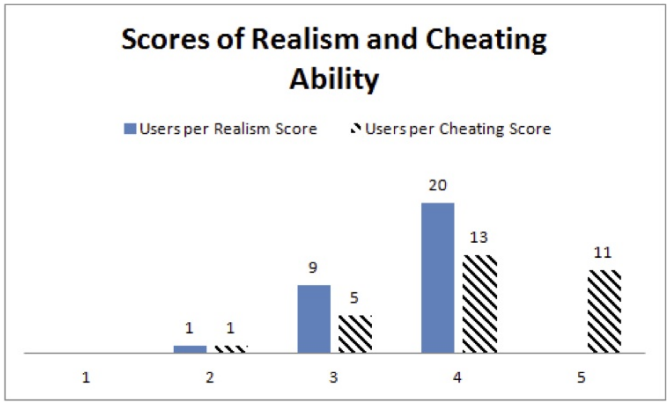

(c) Realism and Difficulty Cheating

Figure 5.4: Quantity of Users per Score Category in (a) preference (b) perceived accuracy and (c) realism of controls and ability to cheat that realism

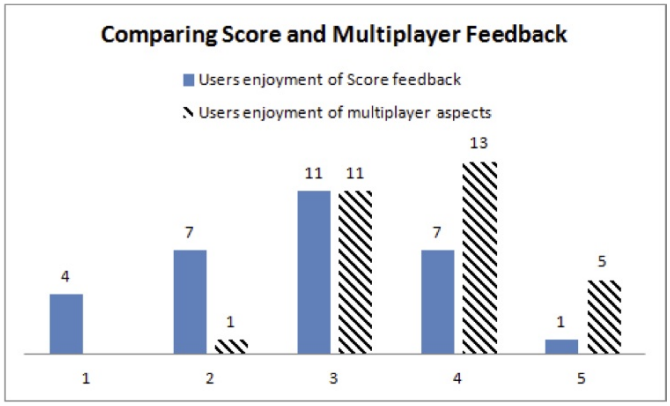

(a) Scoring and Multiplayer Aspects

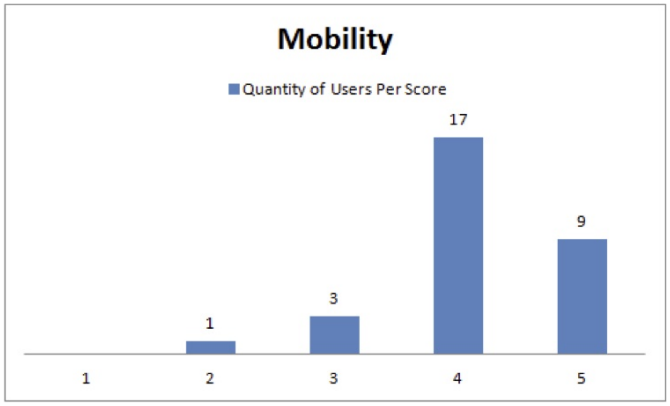

(b) Mobility of System

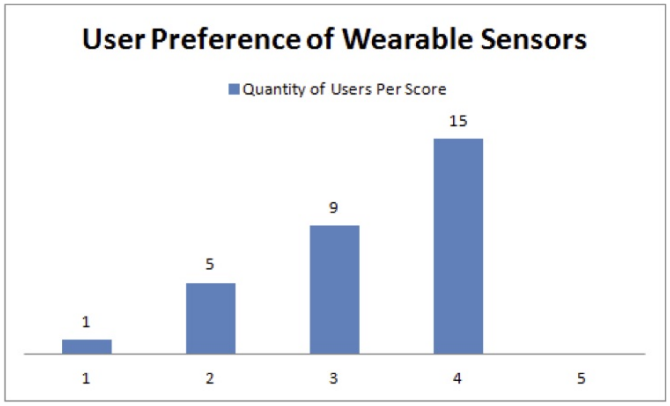

(c) Sensor System

Figure 5.5: Quantity of users per score for (a) perception of benefit to scoring and multiplayer information (b) mobility of the system and (c) perception of pervasive sensor system

CHAPTER 6

Clinical Trial and Regression Techniques

Chronic conditions affect nearly half of all individuals in the United States; 133 million Americans have at least one chronic illness [WA00], accounting for 70% to 80% of health care costs [BWG02]. Most patients with chronic conditions such as obesity, hypertension, diabetes, hyperlipidemia, heart failure, asthma, and depression are not treated adequately, and the burden of chronic illness is magnified by the fact that chronic conditions often occur as comorbidities. Obesity, for example, is becoming a cost and health epidemic in the world [WA10]. The ever-increasing trend has the potential to affect over half of the population of the United States by 2030 [**?**], potentially resulting in exploding medical costs. Indeed, work in [FKT12] estimates that, over the next two decades, there will be a 33% increase in obesity and 130% increase in severe obesity in the United States. Further, this trend if curbed to 2010 levels of obesity, has the potential to save almost $550 billion in medical expenditures over the next two decades [FKT12]. Engagement in physical activity has been shown to be effective in mitigating complications associated with many chronic diseases. Because of this, many approaches to measuring physical activity in adults and children have become popular. In particular, wireless health systems that use wearable motion sensors have been proposed to remotely and continuously measure physical activities [PTS05, BKW12].

The growth of body-wearable accelerometers has given rise to a number of techniques to monitor one's energy expenditure when performing general daily activity [PBW13, LGF12]. Accelerometer systems generally output information

that can calculate energy expenditure and the Metabolic Equivalent of Tasks (hereafter METs) in order to indicate to users their activity levels. METs are an approximation to the level of energy expenditure the metabolism achieves. While it is impossible to monitor the metabolic pathways of a given individual using accelerometers, the goal of accelerometer-based energy expenditure systems is to attempt to predict the energy expenditure achieved. METs are an example of that, where a given number represents the overall level of work and effor the metabolism achieves (e.g. 4 METs for running at 4 miles per hour, 7 METs for casual soccer). From this, an approximation to the energy expenditure of a given user can be achieved. Many approaches exist in determining this information, from calculating activity intensities [BKV97, MCL12, EPM08] to using proprietary counts and formulas from product manufacturers [CCB06b]. Fundamentally, as described in [KLH10], counts are specific to brands of accelerometers and, therefore, their methods cannot easily be adapted to one another. There are more recent methods which take their regressions and formulas from general daily activities and treadmill activities that simulate running [LKS11, KLM13].

The ability to attach low-cost sensors to the body to track movements has given rise to the field of exergaming [MCL12]. The usage of exergames, or active video games for health, to promote physical activity where there was once sedentary behavior [Dal09, SHD09, WJN10, APK11] has presented results in light-to-moderate physical activity [PLC11]. These games can affect the body composition of overweight children [MFM11], though how exergame systems output the actual health information can vary. As exergaming through the use of accelerometers has increased in prevalence in response to the worldwide health epidemic, so has the need to approximate energy expenditure from such systems. The calculation of energy expenditure for exergaming movements using wearable accelerometer sensors, however, has not been addressed in literature previously. Precise measurements for caloric expenditure in exergaming has been calculated in a number of stud-

ies [SHD09] that use invasive measurements of oxygen consumption (VO_2) to get precise measurements of energy expenditure [SHD09, Dal09, PLC11, MTM12] or heart rate [WJN10]. Thus, the focus of this paper is to propose an approach for measuring energy expenditure of exergame movements from wearable accelerometers.

This paper will propose a framework for context-aware MET calculation. This paper takes the approach in [CDH11, PBW13, KLH10, LGF12, SZC12] as a basis to create an acceleration approximation to the METs [AHH11] achieved during exergame activities in an active sports video game such as in [MCL12]. The proposed approach, however, can be applied to any MET calculation method in order to improve its accuracy and reliability. That is, this proposed framework aims to build upon quality metrics that are important in developing wireless health platforms. In Section 6.1, this work briefly discusses these quality metrics and describe how our MET calculation framework improves these metrics. The general theme of our quality enhancement method is context-awareness. This paper particularly focuses on two contextual factors, 'activity type' and 'sensor location', and attempt to incorporate such information through the data processing flow. It will present representations for each movement and sensor node to give more detailed future possibilities instead of finding only a general value for the overall usage so that future systems will not require invasive techniques to gather accurate results. This work will present MET values for the actions and overall game play for a soccer exergame. The soccer exergame is considered due to soccer's popularity as a worldwide sport, and soccer video games are also immensely popular. For example, Electronic Arts, a company that produces a yearly soccer video game, has found recent editions of the game quite successful. EA's FIFA 12, for example, was the fastest selling sports video game of its time [Ele11b], while EA's FIFA 14 had over 5 million users for its demonstration version before the full commercial product had even been released [Ele13]. Using oxygen consumption values

of users within a trial, a base linear regression formula will be derived for each movement. Finally, this work will compare these results to those found via table look up to indicate the necessity in analyzing specific data for each activity for future approximation systems, rather than applying other regressions and tabulated information on generalized activities.

6.1 Preliminaries and Related Work

Wireless health systems perform best by providing quantitative data for the desired goals of producing efficient and effective qualitative data. Desired systems must accurately measure desired features of a given system and provide metrics that can be used for various factors from quality of data processing to approximations based on applications such as caloric expenditure of a given wearable sensor system. Thus, the design of such a system must take several preliminary factors into consideration and then delve into the application specific roadblocks for each desired outcome.

6.1.1 Quality Metrics

A wireless health system is typically composed on a front-end sensing platform and back-end data analytics. It is then essential that the front-end sensing platform provides constructs that consider the quality of information driven from the system and ensure the reliability/validity of the outcomes to support the decision-making processes make it through the back-end framework. In general, a comprehensive quality framework must incorporate the following metrics 1) quality of data; 2) quality of information; and 3) quality of user, as explained later in this section. Our focus in this paper, however, is not on dealing with challenges arising from user-induced errors. In other word, dealing with quality of user is out of the scope of this paper. For the two other types of quality metrics, however, we build our

MET calculation framework such that contextual information about sensors and activities are used to enhance quality of information. We also use a preprocessing algorithm to deal with noise in the sensing data.

Quality of data: Sensors are not perfect. They may be miscalibrated or malfunctioning, and often encounter environmental interference that can result in noisy, imprecise data. The frequency of sampling and the latency associated with delivery of a sensor reading can also impact the utility of the reading. In addition, there is a spectrum of the quality of reading that is obtainable from sensors of the same type. These concerns are related to the quality of data; metrics for describing the quality of sensor data include accuracy, timeliness, confidence, throughput, and cost. Sensors must be calibrated and validated for functionality everyday using intelligent algorithms that do not require user intervention. Such data quality metrics can be directly incorporated into our proposed framework. In this paper, we are only concerned about the effect of noise on MET measurements. Thus, as will be discussed in Section 6.2.3, we combine the three accelerometer readings over a sliding window in order to compensate for noise.

Quality of information: Typically, raw data from simple sensors are interpreted and fused into higher level information that can be used by a user, health-science researchers or health-care providers to make decisions. This includes translation of raw sensor readings into movements along with their timing characteristics. The degree of utility of the derived high-level information for a particular purpose is captured by quality of information metrics. Quality of information metrics for physical activity monitoring will need to be determined and logged and validated in conjunction with gold standards. We take an approach that uses contextual data about the system to enhance quality of information and therefore final MET measurements. In particular, sensor locations (and their associated contributions to MET calculation and movement detection) and activity type are used to 1) develop a sensor weighing approach for MET calculation; 2) develop activity-

specific MET calculation approaches.

Quality of the user: Human error may result in improper placement of sensors required for a specific application. The quality framework must identify potential error caused by users and provide alerts for their correction. This information needs to be further logged to assist clinicians and researchers to identify user's non-compliance. Additionally, as age-associated differences in conscientiousness exist [RWB05] participants needs to be assessed using the conscientiousness component of the Revised NEO Personality Inventory [CM08] and this measure will be used to account for potential age-related discrepancies in the quality of data between younger vs. older cohorts.

6.1.2 Research on Exergaming

Work in [MCL12] presented an active exergaming application as a potential solution for childhood obesity. The authors present a soccer exergame that argues intensity values from velocity calculations guarantees a certain level of physical activity. Further, they cite [BKV97] for the method of calculating METs online, after using a regression from running on treadmills with well known MET values to present their caloric expenditure results to users. However, like many exergaming papers, such as [APK11, WJN10, SC11], the results presented do not focus on the exercise levels achieved by each activity and, instead, focus on primary goals such as cheating prevention [MCL12], range of motion [APK11], or effectiveness of exergames for long-term studies [SC11, WJN10, Dal09].

Only a few papers, such as [BM11] compare the energy expenditure of particular forms of exergames. The method in [MCL12], which is based on the well-cited IMA value calculated in [BKV97] to show the need for movement specific regressions is based upon general daily activity movements. [KLH10] shows that each set of movements and accelerometers has and needs their own regression formu-

las, in that the comparisons are unique due to accelerometer types, outputs, and movements calculated. [MCL12] uses general daily activity movements for regression, this paper will run regressions on the specific soccer movements with ground truth MET values, more accurate than regression on other movements based upon assumed MET values, similar to the MET calculations in [BM11], but with an appropriate accelerometer approximation. Further, this work will consider accelerometer placement in the location presented in [MCL12] for classification purposes as well as the hip and ankle, two common locations for activity monitoring [XZS12].

6.1.3 MET Calculation for Exergaming

Work in [AHH11] compiled a compendium on physical activity, which is used to compare against several activities of physical exercise, daily living, and sports. Indeed, this compendium is the source of many approximations to physical activity in monitoring papers. [RAO08] has put together a compendium of energy expenditure on youth, in particular. However, neither has analyzed detailed motions and METs for those necessary in exergaming systems. In covering a wide range of general daily activities, many approximations can be used, but, in order to have a more accurate representation of exergaming, this work will collect exergaming specific movements in order to supplement such materials for future work. In particular, a comparison will be drawn between the actions of the exergaming environment and those of the actual sport it is comparing against, in this case being soccer.

6.1.4 Regression Models for MET Approximation

Many devices [YH10] have been used and tested in several studies to predict the MET physiological variable using values from uni- and triaxial accelerometers.

[KLH10] discusses the use of multiple regression techniques to calculate MET values of common physical activities from accelerometer output. This work shows the necessity of calculating specific regressions for specific devices and activities. In fact, the work presents results showing approximations from the METs in [AHH11] were, indeed, inaccurate for over 80% of the activities measured. Further, the accelerometer counts ranged from 11 to 7490, a wildly large range. The R^2 value from the regression techniques developed reaches 0.65 in the best settings. As a result, work in this paper will not use accelerometer counts, but instead, raw acceleration values so that comparisons will be easier to draw for future works. Further, the regression techniques should result in comparable results if the work is considered to be accurate. Finally, work in [KLH10] resulted in authors from [AHH11] to update previous work with corrected formulas. This work will also show that such corrected formulas, while appropriate for general populations and activities, do not allow for great variability across users that are possible due to a number of physiological considerations.

Work in [CCB06a] discussed how there are more than 30 regression techniques that produce very different results. [HMB00] discussed the differences in energy expenditure from accelerometer data resulting from inconsistencies in the calibration process, making comparing results among studies difficult. Many systems compare results from devices based on non-universal metrics, such as counts, which are specific to one accelerometer. This work maps specific soccer motions using regression techniques that differ according to activity, using typical accelerometer outputs in units of gravity (based on acceleration as $\frac{m}{s^2}$) to establish metabolic equivalent (MET) equations for soccer exergaming activities.

Work in [AIH10] began identifying improvements by using activity-specific models for regression. By capturing a variety of activities on 24 subjects, from gym activities to daily living, they see an improvement of 15% in their estimates. They show that having multiple sensors on the body to accurately capture the

data improves models, but having multiple regression models and using not only actively captured data but "simulated days" helps improve results. This work will take this approach by using multiple sensors on the body, create activity-specific regressions, and will not only capture data on each individual movement type but also a "simulated game play" session at the end of each collection trial.

Recent works in [AXL13, CKH10] showed that more advanced regression models can provide even more accurate results of MET calculations. In particular, [CHB12] showed that using separate regressions for different classes of motions provided for more accurate results and lower mean errors. This work will adapt such methods for soccer-exergaming by developing regression models for each activity identified, showing that having MET calculation equations for each activity will result in more closely-related regression models than a general regression. While different types of divisions can be examined, this was chosen as it is closely related to the classification results already required in playing any such exergame.

Studies in [KLH10, LKS11] review evaluations of different accelerometers with counts derived from movement specific regressions. While counts will not be used, the movements specific regressions method will be applied to this work, with raw gravity values of accelerations instead of proprietary count values. In this work, the different groupings are set forth by the different movements recognized by any given system. This method, however, can be applied to any setting with any contextual information on the difference between classes that are being considered. If the classes are different, it is suggested that models be created for each class or each cluster of classes that are significantly different, and that this differentiation can result in stronger models across the dataset. Taken into consideration will be the placement of the sensors, the number of sensors, and the activity intensities in order to generate more accurate expenditure values for individual movements as well as establish an MET value for exergames.

Table 6.1: Data collection protocol

No.	Description
1	Sit for 3 minutes to achieve normal breathing with metabolic cart
2	Stand for 3 minutes to establish baseline rest
3	Run designated activity for 3 minutes to establish oxygen uptake for activity
4	Rest (Stand) for 3 minutes to establish baseline rest before next activity
5	Repeat

6.1.5 Extensions

The methods presented in the following section are intended to demonstrate the use of contextual information, knowledge about the given movements, in order to present a stronger model for predicting energy expenditure. Readers should consider the quality metrics presented earlier along with contextual information provided from knowledge and potentially other sensors in order to improve the development of models and regressions to other applications as necessary.

6.2 Methods

The trial run in this work, as with many initial wireless health applications, consisted of two separate phases. The first such phase is a data processing phase in which a collection protocol is set up to determine the feasibility of a given application and generate models for large scale usage. The second phase is the processing techniques used on that data to generate those models.

6.2.1 Clinical Setup

Work in this paper presents a method to approximate metabolic equivalents (METs) of various exergaming activities, an IRB approved study (UCLA IRB #12-000730). The approach relies on leveraging contextual information about the sensing platform in order to improve MET calculations based on regression models. The purpose of the study was to develop an approximation for the METs

Figure 6.1: Subject running trial with metabolic cart and accelerometers attached

produced when using exergame movements, in order to set up future studies analyzing body composition changes. Participants were given three accelerometers to wear, including two GCDC +/- 2g accelerometers worn on the hip and ankle [Con], and a +/- 5g Memsense IMU [Mem] worn on top of the foot to help simulate motion at contact with a soccer ball (and to correlate with work in [MCL12]). Users were then attached to a metabolic cart that examines the volume of oxygen taken into the lungs during activity, a key measurement in determining actual MET values. In fact, the oxygen uptake, presented as $VO_2(\frac{ml}{min})$ can result in METs given by

$$MET = \frac{VO_2}{f \times m} \tag{6.1}$$

where m denotes the mass of the user in kilograms, and f represents a factor that changes based upon the general fitness of the group analyzed, 3.5 here in the case of healthy, active adults. Six healthy male subjects between the ages of 22 and 31 were selected and ran the protocol shown in Table 6.1.

This allowed for testing of each activity, to be described in Section 6.2.2, and determine the oxygen uptake of each motion in order to obtain a ground truth

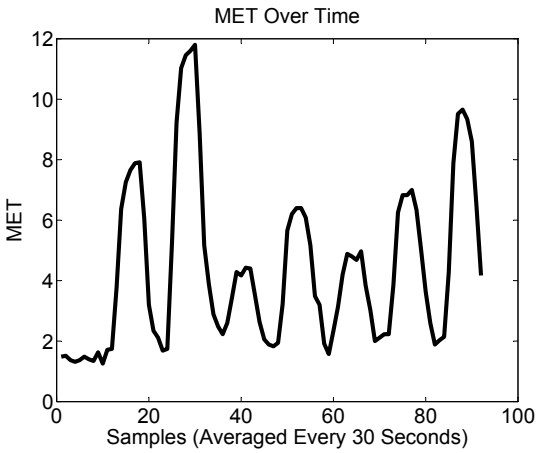

(a) METs

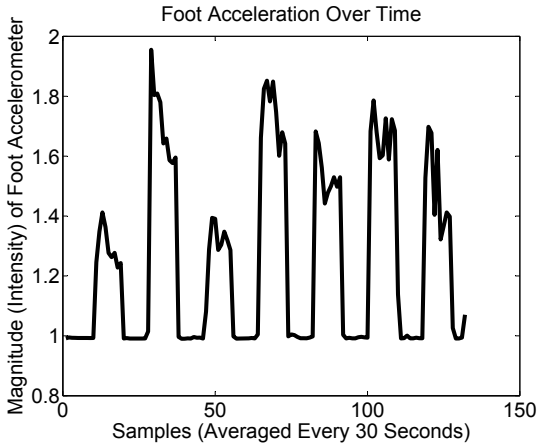

(b) Magnitude of Acceleration for Foot Sensor

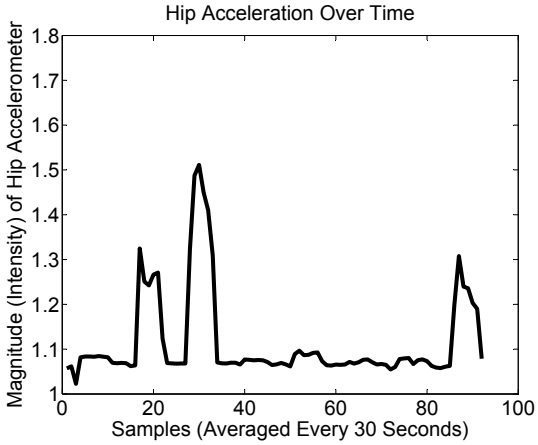

(c) Magnitude of Acceleration for Hip Sensor

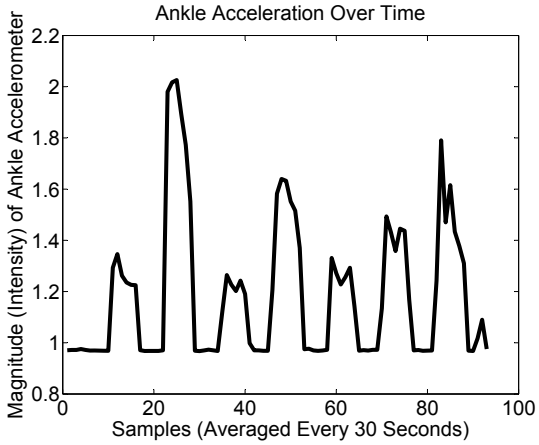

(d) Magnitude of Acceleration for Ankle Sensor

Figure 6.2: User 1 values for (a)METs and accelerations for (b) foot, (c) hip, and (d) ankle

MET value to obtain accurate caloric expenditure information for each exergaming activity. It is known that for constant load activities, a steady state is typically achieved by three minutes of exercise and to only use data after this point in analysis [CGS84]. Fig. 6.1 shows an image of a user running the designated protocol for data collection. It seems, as expected, the sensors closest to the greatest point of action might correlate most closely to the resultant METs. However, we notice that in some cases the intensities are similar, such as in Fig. 6.2b for the foot

Table 6.2: Collected Soccer Moves

No.	Move	Description
1	Run	Running in place
2	Sprint	Sprinting in place
3	Pass	Passing Ball Directly left
4	Chip	Chipping a ball up and to left
5	Medium Shot	Medium Powered Laces Shot
6	Full Powered Shot	Full Swinging Shot
7	Simulated Game	Simulated Exergame-play

Table 6.3: Simulated game movements

No.	Description
1	pass, pass, medium shot
2	pass, pass, strong shot
3	sprint for 5 seconds (defense)
4	pass, chip, shoot
5	running, fake shot, pass, strong shot
6	sprint for 5 seconds (defense)
7	sprint for 5 seconds

accelerometer, despite the METs being different. Thus a combination of results may produce the best value.

6.2.2 Exergaming Movements

From [MCL12], six soccer movements were selected for data collection. Those movements and their descriptions are shown in Table 6.2. Each movement was repeated for the full 3 minutes. Users would perform the motions at their desired intensities (showing variability in the intensities recorded, as expected) and at roughly the same pace (enough time for users to settle and repeat the action, approximately 3 seconds between each action). This gives the activity intensities if one were to repeat each soccer action, which happens in many games. Repeated

actions are more realistic in an exergame than a real soccer environment as most team-play video games change the focus to the player with the ball every time the ball is passed between players on screen. However, as it is not entirely realistic to simply pass for 3 minutes straight a simulated game play mode was created for the testing environment (kept the same to generate uniform results). This simulated game play ran as described in Table 6.3, with 5 second gaps between movements and running in place for the duration of the 3 minute trial, based off of an exergame like that of [MCL12].

This set of actions simulates movement of the soccer ball in a soccer environment including a series of running actions and sprinting actions that happen throughout game play to give a more realistic overall game play MET value. It is intended to simulate a series of offensive moves and defensive running activities that occur throughout normal game play.

6.2.3 Context-Aware MET Approximation

Due to the variability in any individual's breathing pattern, the VO_2 data was calculated in 30-second averages. As a result, the accelerometer data needed to be synchronized in the same format. Further, systems such as [BWV94] and [BKV97] use a variation of either the integrated absolute values or the magnitude of the accelerometer data. For this work, the magnitude of each accelerometer is considered in order to combine the x-axis, y-axis, and z-axis for an overall intensity calculation, as well as account for the effects of gravity by setting a new baseline value for inactivity. Thus, after each axis of the accelerometer is averaged over 30-second windows, the magnitude of the acceleration vector is calculated by

$$\| \mathbf{a} \| = \sqrt{{a_x}^2 + {a_y}^2 + {a_z}^2} \tag{6.2}$$

This value is collected for each accelerometer. Then the peaks of each intensity

point and each MET point were correlated and a regression analysis was run to determine the curve of best fit.

6.2.3.1 Binary Sensor Weighting Model

The idea behind this approach is that different sensors contribute to calculation of the MET values differently. Sensors that provide most information regarding movements of interest can be used for MET calculation. Thus, a binary selection of the sensors will be applied to find the best subset of sensors for a particular physical movement monitoring application. Several approaches were taken in testing the best combination of sensors for the most appropriate and accurate regressions. The first is a simple selection of sensors, in which a 2-D linear regression is run where

$$MET_{reg} = \alpha_0 + \alpha_1 \cdot s \tag{6.3}$$

where s is a potential combination of each sensor is given by

$$s = \sum_{i=1}^{n} c_i \cdot s_i \tag{6.4}$$

where i is the number of sensors available (e.g. 1 is hip, 2 is ankle, 3 is hip and so on), and each c_i is 0 or 1 whether it is used or not. This method was first presented in [MAL13]. Thus, the best regression may be selected from the most appropriate range of data. In particular, this method might be best used in relation with sensor selection techniques for other purposes, including classification [GGJ09] and power usage [GJ11] by determining how many sensors should be necessary for any given application. In our experimental results, however, we perform an exhaustive analysis and find the MET values for all combinations of the sensors used for data collection in this study.

6.2.3.2 Sensor Weighting and Activity-Specific Model

If all the sensors will be used then perhaps binary selection of each sensor would not produce the best results. A more complicated regression would allow for fractional constants. For this work, three nested for loops were written to range the constants c_i from 0 to 1, in this case in increments of 0.1. This included 0 and 1 so as to encompass the previous method's results as well, with all formulas saved so that the best results could be selected when the number of sensors are decided from any other application, or the best picked here. This sorting was based upon the R^2 value of the regression.

Once the appropriate weighting of each sensor was found, the comparison between MET and MET_{reg} can be analyzed further. As indicated in Section 6.1, analyzing each activity independently can provide stronger regression results, rather than developing a universal model for the all movements. An extra iteration of the method indicated here is run per activity, to develop individual regression models for each activity and the simulated game play independent of each other. Thus, any future exergaming system that has a classification system, will not only identify the movement performed, but the appropriate regression model necessary to calculate the most accurate approximation of caloric expenditure. As will be shown, this activity-specific regression technique provides a much stronger linear regression for each movement.

6.2.3.3 Optimal Sensor Weighting

The sensor weighting technique described in Section 6.2.3.2 is a heuristic approach. Instead of weighting the regression as such, a multi-dimension regression can be run to select the weights in a completely variable format such that an objective function (e.g., regression error) is minimized. In this case, the MET estimations are given by

$$MET_{reg} = \alpha_0 + \sum_{i=1}^{n} \alpha_i \cdot s_i \tag{6.5}$$

where the s_i are just as in the previous section. In this case, instead of looping through the c_i and setting the weights directly, the algorithm will select the weights through a multidimensional linear regression. This approach finds the best fit by minimizing the amount of mean square error. The method can be applied to both algorithms discussed previously. That is, the multidimensional regression can be used to optimally weight sensors either with or without integrating the 'activity type'. In our experimental results, we will demonstrate the accumulated improvements made by integrating the two contextual factors (e.g., sensor weighting and activity type) within the MET calculation model. A concern about such a system would be the accuracy of determining such an 'activity type'. If such a system is not accurate in determining the motion, then it is better to use a general approximation. The movements in this work are, however, accurately detected. The method described in [MCL12] reaches 81% precision and 80% recall on the desired movements, while improvements on such an algorithm approach 90% [MNL14]. Thus, it is safe to assume an 'activity type' specific regression is a safe model to develop, as each movement will be identified with high accuracy.

6.2.4 Cross Validation

In order to verify the strength of such regression models, one must measure its ability to predict appropriately the MET values being outputted. As a result, a leave-one-subject-out cross-validation is run to verify that the model presented on the given data set can predict the appropriate MET values. In particular, for each subject, the regression model is run on the subset of data training data. Then, the testing data from the left-out subject is run. The MET from the regression model is compared against the ground truth, and the results are averaged across

Table 6.4: R^2 values from Sensor Selection Regression Analysis

No.	Description	R^2
1	MET vs. Foot	0.2431
2	MET vs. Ankle	0.5662
3	MET vs. Hip	0.2342
4	MET vs. Foot + Ankle	0.4655
5	MET vs. Foot + Hip	0.3355
6	**MET vs. Hip + Ankle**	**0.7147**
7	MET vs. Foot + Hip + Ankle	0.5472

all users for each movement as:

$$MAE_{move} = \frac{1}{n}\sum_{i=1}^{n} |MET_i - MET_{i_{reg}}| \tag{6.6}$$

where n is the number of subjects in the data set.

6.3 Experimental Results

This section covers the three processing techniques designed for this protocol. Beginning with the general sensor location problem, this section covers the progression to the generalized multidimensional technique for activity type regressions.

6.3.1 Binary Sensor Weighting Results

Fig. 6.2a shows the METs as calculated from VO_2 data; associated accelerometer magnitudes for one of the users of the trial are shown in Fig. 6.2b, Fig. 6.2c, and Fig. 6.2d respectively. Table 6.3 shows the results of the regression run on the analysis. At each movement point the peaks were detected after the three minute mark and used for the polyfit regression run in Matlab. A combination of the

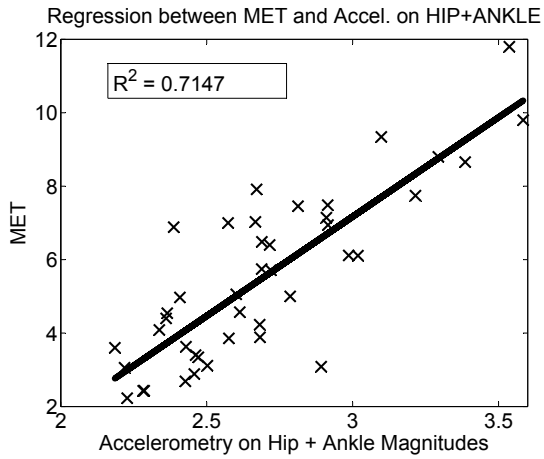

Figure 6.3: Regression run on data from Hip accelerometer and Ankle accelerometer at peaks

hip accelerometer and ankle accelerometer seem to do better than using the foot, like is used in [MCL12]. It seems there is perhaps too much activity at the top of the foot, or rather, perhaps the peaks themselves should not be used. As can be seen in Fig. 6.2b, the average intensity value over a period seems to differ from the peaks; however, this analysis is left for future work, as it does not correlate in time with the oxygen consumption, and therefore requires further analysis. The best fit line produces the following model:

$$MET_{reg} = 5.3 \times (\| \mathbf{hip} \| + \| \mathbf{ankle} \|) - 8.6 \tag{6.7}$$

where in this case, the magnitude of each accelerometer is summed together. The sum, or an average, would result in the same regression. A sum is taken in order to calculate the intensity at a given point in time. This follows from the plot in Fig. 6.3. As can be seen from the plot, there is significant variability from user to user, calculating based off of simple METs from a table such as done in [MCL12] to derive MET formulas will not provide accurate representations unless those tabled values consider a wide enough population. Regression analysis must be run on a large number of subjects with varying levels of intensities and body

composition in order to do better; however, finding the exergaming specific METs can improve approximations for those not wishing to run a clinical study. As such, a more detailed online calculation of caloric expenditures can be run when knowing the MET values for each activity as has become clear here.

6.3.2 Activity-Specific Results

As suggested, regressions based on context information can provide stronger results. In this case, knowing the activity and running variable weighting on each of the activities results in significantly stronger results. In Table 6.5, we show several key factors. The first is that, some of the individual regressions shows stronger results than found in Table 6.4. Second, notice the generalized sensor-selection method used in the previous section. While results seemed strong in the general case, notice how weak the results are, in particular for certain movements like the medium strength shot. This sensor location and selection method lists the best possible combination in each movement type, as depicted in Table 6.6. The results are not simply the hip and ankle sensors as discussed in the generalized case, but different sensors for different movements. Thus, for exergaming movements similar to other works, activity specific regressions perform better generally than overall regressions, even with problem movements such as the medium shot. When empirically determining the weighting, all the results improve generally. When developing an exergame, one can choose to use an overall regression based upon the simulated gameplay only, or can create formulas that are chosen based upon any classification result given. For each movement, as expected, the multidimensional regression produces the best results. In the example presented in this work, the multidimensional regression has four parameters, one being the constant, the other three being scaling factors on each of the sensors. The best parameters for each movement are listed in Table 6.7. Finally, the average R^2 value is calculated for each movement, including and removing the medium shot as it appears to be

Table 6.5: R^2 values of best regression for each activity

Movement →	Run	Sprint	Pass	Chip	Med. Shot	Full Shot	Sim-Game	Avg. w/out Med. Shot	Avg.
Sensor-Selection	0.81	0.85	0.81	0.56	0.10	0.66	0.71	0.73	0.64
Variable-Weighting	0.86	0.89	0.85	0.66	0.74	0.68	0.76	0.78	0.77
Multidimensional Regression	0.98	0.95	0.99	0.75	0.90	0.94	0.78	0.90	0.89

Table 6.6: Mean Absolute Difference for each Movement using Activity-Specific Models of Regression

Run	Sprint	Pass	Chip	Med. Shot	Full Shot	Sim. Game	Avg.
2.336	2.550	2.021	2.061	2.107	2.085	2.460	2.231

Table 6.7: Sensor Location Choices for each Activity Type in Sensor Selection Regression

No.	**Movement**	**Locations**
1	Run	Ankle
2	Sprint	Hip + Ankle + Foot
3	Pass	Ankle + Foot
4	Chip	Hip + Ankle
5	Med. Shot	Hip + Ankle
6	Full Shot	Foot
7	Sim. Game	Hip + Ankle

a problem movement. Such movements should be investigated further, needing more data for more valuable models.

6.3.3 Cross Validation Results

The results of the leave-one-subject-out cross-validation are shown in Table 6.6. The mean absolute difference shown results in an error of about 2 METs, which would still put the general intensity levels in the correct ranges. This shows that the method and R^2 values show a model that can accurately predict the intensity of a new individual. The widest range of error comes in the effort put forth while running. This is likely due to the different physical conditions of the subjects, the speed with which they ran, and the strain this put on the body. While there is obvious room for improvement, as will be discussed further in section 8.4.

Table 6.8: Values for each parameter in the multidimensional regression

No.	**Movement**	α_0	α_{hip}	α_{ankle}	α_{foot}
1	Run	-35.5	52.6	10.5	-30.6
2	Sprint	2.83	18.5	-10.6	-9.37
3	Pass	105.6	99.0	-69.7	-139.5
4	Chip	14.5	15.7	-7.99	-21.0
5	Med. Shot	-10.0	-10.7	7.38	17.4
6	Full Shot	-7.23	-4.91	9.20	5.40
7	Sim. Game	5.69	-3.17	7.83	-7.61

6.3.4 Caloric Expenditure

The purpose of such MET calculations is to ultimately calculate the energy expenditure through caloric expenditure. The MET is an approximation of the metabolic expenditure of the body. Further, using a method from [MHW11], the caloric expenditure can be extracted from this information using the following equation:

$$Calories = \frac{k \times MET \times m}{200} \times t \tag{6.8}$$

where k is the same factor used in the MET predictions, m is the mass in kilograms, t the time in minutes, and 200 a scaling factor. Thus, to ultimately prove the validity of such a system, the caloric expenditure of the trial is shown in Table 6.11. This table shows the actual caloric expenditure achieved by the model for each given user over the course of the trial, summed over each activity (three minutes per activity).

Table 6.9: Comparing True Exergaming MET values with Ainsworth

	Ht (cm)	Wt (kg)	Age (yrs)	Run	Sprint	Pass	Chip	Med Shot	FP Shot	Sim-Game	Ains (Light)	Ains (Intense)
1	170	76	28	7.91	11.80	4.40	6.40	4.97	7.0	9.66	7.57	10.80
2	187	82	29	4.20	7.34	2.29	6.49	2.97	5.0	7.29	7.47	10.68
3	174	68	29	4.20	8.66	5.06	4.94	4.11	7.83	9.82	7.16	10.23
4	183	79	26	3.80	7.03	2.60	3.97	3.11	6.11	7.49	7.38	10.54
5	174	70	31	4.54	6.51	3.49	4.63	3.17	4.23	5.74	7.36	10.44
6	175	66	22	6.89	8.80	2.80	3.63	4.14	5.94	7.60	6.83	9.75

Table 6.10: Average METs for each activity

No.	Activity	AVG ± STD
1	Run	5.26 ± 1.70
2	Sprint	8.36 ± 1.92
3	Pass	3.44 ± 1.09
4	Chip	5.01 ± 1.20
5	Med Shot	3.75 ± 0.79
6	FP Shot	6.02 ± 1.30
7	Sim-Game	7.93 ± 1.55

Table 6.11: Caloric expenditure achieved by each user during the trial

User	**Calories**
1	208
2	153
3	159
4	141
5	118
6	138

6.4 Discussion and Future Work

While the regression can indicate a more accurate way of calculating METs in an online fashion while participating in exergaming activity, it may also be interesting to see a general MET value for each activity, including a comparison to what [AHH11] uses as the corrected formulas for METs per person. Since there is great variability among individuals from height and weight to age, the corrected formula is supposed to indicate the appropriate MET for that individual. As can be seen in Table 6.9, the final two columns show what soccer (casual and intense) would be with the corrected models for each of the individuals involved

in the study. As can be seen, there is still little variability. However, looking at the MET value of each individual for each of the actions shows great variability across the user base, a reason for needing large populations for future regressions, but also for the regressions themselves, as the basic table approximation can vary for specific actions like these of soccer exergaming, showing need for specific values for exergaming. Exergaming systems should take the approach of [MCL12] to ensure only those reaching a certain level are accepted as appropriate activities, though even then, variability will exist due to intensity desired as well as physical fitness of user. Table 6.10 shows the average MET and standard deviations for all the movements and the simulated game. It seems the simulated game play can reach that of soccer, a promising result for exergaming research. Further, having an MET for each movement can allow for better realism, using such an MET calculation as a cheating prevention cut-off instead along with other techniques to ensure realism and activity. Finally, it is obvious that a general level of activity can be guaranteed but that specific caloric expenditure approximations may need more user information than simply accelerometer intensities. Further movement data is necessary to better validate models of specific movements, such as the medium powered shot discussed in the previous section. This is shown with the mean absolute error of the model predicting the METs in cross-validation. While the activity levels are in the general correct ranges, over a long period of time an error of 1 or 2 METs might result in too large an error. It is still a more accurate model but room is clearly left for improvement.

This work presents a baseline approach to calculating the METs of a soccer exergame ranging from its movements to a simulated game play calculation. These values and the regression formula will be used as a baseline for an extended study on the overall values reached actually playing particular exergaming systems. Further, instead of signal processing simply on the peaks, perhaps an average across the climb, peak and descent of each activity can be taken. Finally, when a more

accurate determination of METs achieved during exergaming is concluded upon, such a system must be re-incorporated into an exergaming system to give accurate long-term caloric expenditure calculations for users of these exergaming systems, in particular due to the heavy importance placed on sensor location for classification techniques as the primary requirement for many of such systems. Variability must be better modeled into such systems. When conducted in a laboratory setting perhaps the body strains more by wearing the oxygen equipment. Variability needs to account for more users, different factors on the actual MET value, and account for users becoming more efficient over time. Further, this variability can be compared over population ranges. The data in this study is taken from healthy young adult subjects. Separate models can be created for overweight/obese adults, and for healthy and overweight/obese children as well, in order to create an even more general framework for MET approximations across all population types. Finally, the sensor selection approach indicated here can be further analyzed, such as in [MPA14], in order to reduce the computational complexity and power optimization. As was seen here, the power optimization can be improved by roughly one-third to two-thirds depending on the movements monitored and sensors desired.

The MET and caloric expenditure information presented in this work is applicable to exergaming and other physical activity monitoring tools with the use of accelerometers. This information should be provided to a user in order to better represent their physical activity information. Further, this data can be transmitted to any user's clinical professional. Doctors can use this information to assess the length of time playing video games, the activity levels achieved playing those games, and using this information to better assess the physical activity levels achieved with real data instead of simple questions asked in periodic checkups. This is applicable to treating obesity of many levels in children and adults. In particular, this method is flexible, can handle lots of variation, and thus is appli-

cable to a large population of healthy and unhealthy subjects to give Doctors the greatest power in addressing obesity as a result of physical inactivity.

6.5 Conclusion

This work developed a procedure and a regression technique to determine the METs achieved when participating in soccer exergaming. Several sensor locations were tested, as well as results compared with the individual locations and the fusion of multiple locations. Further, by using context information, stronger correlations can be determined when the activity information is given. Each individual movement regression results in a stronger model for approximation than any of the generalized formulas. As exergaming research expands to target the growing childhood obesity epidemic, caloric expenditure results must be verified to ensure an appropriate intervention is possible and measurable to a certain degree of accuracy. This work produces an oxygen consumption data set for exergaming activities and produces METs of each particular action, instead of general use values. Instead of using table values to approximate METs and create a regression from this, this work used actual volume of oxygen uptake to determine an accurate representation of the METs found when participating in many movements in an exergaming specific fashion instead of those from general daily activities. Finally, this paper also concludes that soccer exergaming can reach an MET value of 8 even across variable subjects, which is roughly the same as the predicted value for actual light/casual intensity soccer.

CHAPTER 7

Advanced Regression for Modeling Clinical Results

Sedentary behavior is a root cause of several chronic conditions affecting health of adult and children in the United States and worldwide[WA00][BRL12]. Such physical inactivity often leads to overweight and obese populations [FKT12] that result in chronic conditions such as cardiovascular disease or diabetes. Cardiovascular disease and diabetes both present a significant economic burden in the United States as well as health, with cardiovascular disease accounting for an estimated $670 billion in health care costs in 2010 [HTK11] and estimates suggesting diabetes may approach $900 billion in health care costs by 2015 [Ass13]. As a result, many solutions using wireless wearable sensors have been presented to monitor activity [EPM08][BKW12].

One such solution is the use of these body-worn sensors, such as accelerometers, to enable the actions of an individual to control a video game[WJN10]. Exergaming, or the mix of physical activity to control an electronic game, has been shown to increase physical activity and promote a potential solution for obesity [Dal09] as well as other activities, such as stroke rehabilitation[APK11]. Such exergames, such as the ones presented in [MCL12], show a framework for potentially active video game play, but the energy expenditure of such systems needs to be verified. Some work has been done to determine the energy expenditure and metabolic equivalent of task (MET) [AHH11] of exergames to show that they can burn calories and be a useful intervention over time [BM11][MTM12][SC11]. These systems

do promote activity and thus, it is better to have accelerometer approximations to this energy expenditure so that the systems can provide useful information found in works that show it provides real exercise [SHD09][PLC11][MFM11].

In order to provide useful information while exergaming, accelerometer approximations to the energy expenditure need to be created. Since it is difficult to measure the oxygen consumption while playing, such as in [MAL13], approximations need to be developed that correlate closely to the caloric expenditure values. Work such as [LKS11][KLH10][CDH11] develop advanced models of mapping accelerometer values to energy expenditure. This work applies such systems to create an advanced model for energy expenditure while playing a soccer exergame, which often involves unique movements derived from reality but adapted for gameplay. It will measure the appropriate output in terms of METs, which can then be translated to caloric expenditure directly based on each user. By using advanced, non-linear models of regression, namely a support vector regression, it will provide stronger mappings between exergaming movements, the accelerometer readings for those inputs, and the METs generated. It will be shown, through a leave-one-subject-out cross-validation, that the robustness of such a model for exergaming movements by measuring the mean absolute error (by calculating the mean absolute difference) in each approximation is greatly improved over standard linear regression techniques.

7.1 Related Works

7.1.1 Common Regressions

Work in [SWC05] showed an advanced, non-linear regression method for walking energy expenditure approximation. By using a support vector regression, instead of a standard linear regression, [SWC05] shows reduced mean square errors for walking energy expenditure estimations. The results support that using advanced

regression techniques can provide stronger results. This work will adapt such a methodology to the data collected and set of movements to a soccer exergame, to expand upon results beyond walking to many types of leg movements. Similarly, work in [SZC12] uses an artificial neural network, which also models a non-linear regression method, to estimate METs of activities with low root-mean-squared error by combining it with the activity detection. This work will use the advancement of using the individual identified activity in order to develop movement-specific models that approximate the energy expenditure.

7.1.2 METs for Exergames

Work in [MAL13] was one of the preliminary investigations into the MET values associated with exergaming movements. A trial conducted on six people showed that specific regressions are needed for each type of activity, because the regular table look up of energy expenditure does not allow for a vast amount of variability. Work in [MAL13] developed MET values for a soccer exergame, using a simple linear regression, showing simulated gameplay manufactured an average of 7 METs. This work extends off of [MAL13]. Namely, more advanced regression techniques are designed here to allow for robustness across subjects. Work in [MAL13] runs a simple regression and does not use the data in cross-validation in order to measure error across subjects. This work will continue the work of developing accelerometer approximations for specific activities, in this case soccer exergame movements, but develops a more robust model to allow variability across subjects. First, this work will adapt the method shown in [MAL13] to work on individual activities, and then will further apply a support vector regression in order to produce strong results in cross-validation.

7.2 Method

This section describes the data collection procedure and trial conducted for MET approximations. This study was a UCLA Institutional Review Board(IRB) approved study # 12-000730, where 14 healthy adults, from ages 18-35, were selected to run through the same protocol defined in [MAL13]. This trial consists of six movements and a simulated gameplay phase to relate the data with actually playing a game and the mix of movements necessary. The movements are running in place, sprinting in place, passing a ball left, chipping a ball left, medium powered shot, and a strong powered shot. Each movement was repeated consistently for three minutes to measure the constant load of each movement, as necessary to determine the oxygen consumption levels of a user. Three minutes of rest between each set of moves allowed for the user to return to a state of rest. The simulated gameplay, in order to compare directly, was used from [MAL13] where the movements are mixed along with running actions to mimic the game described in [MCL12]. User's wore a mask connected to a metabolic cart in order to measure the volume of oxygen consumed during the movements. This data was averaged over 30 second windows in order to filter the individual breaths that might skew the results. The motion data was captured by wearing accelerometers strapped to the body in three locations. The accelerometers were Shimmer3[shi] wireless IMUs with a +/- 6g accelerometer sampled at 50 Hz. Those locations are the hip, on the right side, the ankle, on the outside of right leg, and the foot, on top of the right foot, as shown in Fig. 7.1.

7.2.1 Data Processing

Data that was collected involved the volume of oxygen consumed during three minutes of constant load activity. From this information, listed as $VO_2(\frac{ml}{min})$, MET can be calculated as

Figure 7.1: User wearing sensors on the hip, ankle, and foot for collection trial

$$MET = \frac{VO_2}{k \times m} \tag{7.1}$$

where k is a factor that scales based upon the physical condition of the user (3.5 in the case of the trial in this work), and m is the mass in kilograms of the user. Since the data is collected from the IMUs at 50 Hz, the data needs to be

Activity	AVG ± STD
Run	6.08 ± 3.32
Sprint	9.27 ± 3.14
Pass	4.39 ± 2.18
Chip	5.77 ± 2.58
Med Shot	5.13 ± 3.34
FP Shot	6.76 ± 2.61
Sim-Game	8.24 ± 2.09

Table 7.1: A high standard deviation of METs for each activity

averaged. Similarly to [MAL13], the magnitude of acceleration is calculated, then averaged over 30 second windows. The peaks as well as flat rest data are selected and used for the regression.

7.2.2 Cross Validation

While the R^2 correlation coefficient of a regression reports the fit of the regression to the data, this does not necessarily account for variability across the users, even those shown in [MAL13], greater variability is found over the 14 subjects collected here. A leave-one-subject-out cross-validation is run to test the generated models' variability across each test subject. Table 7.1 shows greater variability than reported in [MAL13], where in [MAL13] most standard deviations were reported to be under 2, even close to 1 in many cases, here they range between 2 to 3 $METs$ that, if reported incorrectly, is a wide range for physical activity (perhaps the difference between low and moderate physical activity). For each test subject, the training data is split between each activity. Each of these activities have their own model. The test data is then run on the model to see what determined predicted MET value is produced, labeled MET_{reg}. Then the mean absolute different (MAD) is calculated for each activity and reported to determine the best

Activity	Run	Sprint	Pass	Chip	Med. Shot	Strong Shot	Simulated Gameplay
Sensor Selection	2.99	3.06	1.97	2.54	3.12	2.29	1.83
eps-SVR (RBF)	1.37	1.55	0.82	1.09	1.65	1.08	0.58
nu-SVR (RBF)	1.872	1.909	0.740	1.099	2.026	1.154	0.548

Table 7.2: The Mean Absolute Difference for each activity in cross-validation for the Sensor Selection method and the Support Vector Regression

results. The MAD is calculated as

$$MAD = \frac{1}{n}\sum_{i=1}^{n} |MET - MET_{reg}| \tag{7.2}$$

where n is the number of subjects, MET the ground truth measured value and MET_{reg} the output of the regression model. A model's robustness will be determined by a low MAD (also reported as mean absolute error in some works).

7.2.2.1 Sensor Selection Method

The sensor selection method, first reported in [MAL13], was implemented using the polyfit function in Matlab. For each activity, the seven different training configurations had each model created. The configurations were using only the hip, only the ankle, or only the foot sensor data, or the combination there of. Then the R^2 of each model is evaluated on the training data alone to pick the best reported model as they choose. This model is then used to calculated the absolute difference of predicted MET.

7.2.2.2 Support Vector Regression

Using LibSVM [CL01], a support vector regression is created, using an epsilon-svm and a radial basis function kernel. Similarly to the Sensor Selection method each activity model has its reported MET_{reg} value outputted and the absolute difference stored for the overall mean absolute difference.

7.3 Results

As seen in Table 7.2, the robustness of each method is shown. The mean absolute difference of each movement type is shown, validating the need for movement, specific models. Certain movements, including the simulated gameplay, have much better mean absolute difference values than other associated moves. Further, the movement models can be separated easily by using a classifier to find a move type first. The Support Vector Regression performs better in every movement, and is particularly strong for the simulated gameplay. This would lead one to believe a robust, advanced regression model, can produce highly accurate energy expenditure values for such an exergame.

7.4 Future Work

There is higher variability still presented by the increase in data. Thus, more data needs to be collected, and across different body types to account for various states of well-being from fully healthy and active to completely sedentary and obese. Further, models that allow for more variability across types should be compared. Finally, the variability between such exergaming movements and the energy expenditure values, should be compared over longer periods of actual gameplay.

7.5 Conclusions

Exergaming, as a solution for obesity has shown the need for advanced energy expenditure models for the specific type of motions necessary for such games. By calculating the energy expenditure through the usage of oxygen consumed by a user, this work developed a model and approximation method for estimating the metabolic equivalent of task (MET) of each activity in a soccer exergame through

the usage of a support vector regression. This higher dimensional regression tool, in cross validation, had a mean absolute error (mean absolute difference) of less than 2 for every movement, and in some cases under 1. This low error is a full MET, in some cases 2 METs, better than similar linear regression techniques, providing a significantly more accurate accelerometer approximation to energy expenditure for exergaming movements.

CHAPTER 8

Segmenting and Movement Specific Models

Body-wearable sensors for personal health monitoring have become an important tool in solving health problems. Chronic illness, which affects 133 million Americans [WA00] resulting in a majority of the health care costs [BWG02], often stems from physical inactivity [BRL12]. Indeed, cardiovascular disease and diabetes, the two most common chronic diseases resulting from inactivity are the source of an increasing economic burden on the United States. For instance, cardiovascular disease is projected to cost up to $818 billion in direct medical costs and $276 billion in indirect (loss of productivity) costs by 2030[HTK11], while diabetes, estimated at $245 billion in 2012, will increase due to the estimated prevalence doubling by 2050 [Ass13]. As a result, there is a dire need of a solution that addresses this problem, and wireless health systems, used to monitor the physical activity of wearers [PTS05, BKW12], are now looking to wrist-worn platforms for tracking various forms of exercises. Aerobic and resistance training exercises have both been shown to help prevent and address these diseases [BNT13]. In some cases, resistance training alone can be effective [PFB00].

Wearable sensor systems provide a large number of monitoring applications for energy expenditure. They can range from using only a single sensor[CLA13] for general energy to networks of sensors for detailed applications, like swimming actions[DAC13]. Monitoring energy expenditure while performing general daily activity [PBW13, LGF12] is the most common such application. Calculating activity intensity, and approximation to metabolic equivalent of tasks allows for

estimation of actual energy expenditure [BKV97, CCB06b]. A review of such wearable accelerometers and activity energy expenditures shows popularity of such systems in promoting physical activity[BBS14, LKS11, KLM13]. For each existing platform, many application specific solutions exist as well, from design of gameplay [Dal09, MNL14] to light-to-moderate physical activity monitoring[PLC11].

With the popularity of Fitbit, and the emergence of the smartwatch platform from companies like Sony, LG, and Samsung, wrist-worn sensor devices for exercise monitoring will become an increasingly important tool in personal health monitoring. In particular, exercise routines and repetitions can be counted in order to track a workout routine as well as determine the energy expenditure of individual movements. Indeed, mobile fitness coaches [KMH13] for counting repetitions of exercises [CCC07, CJK13] have become a growing topic of interest, including the selection of sensors and locations for tracking activities [PEA07]. Mobile fitness coaching has covered the range of topics from quality of performing such sports actions [MRD12][VBG13] to detection of the specific sports activity [LFH13].

This paper introduces a framework for platform creation (e.g., accelerometer only system vs. accelerometer and gyroscope) and machine learning of some activities, which can be especially useful in the emerging market of smartwatches. By identifying the most informative activity-specific features, a system can be optimized to reduce the computational load as well as the power consumption through appropriate sensor selection. In particular, the decision to add an accelerometer and a gyroscope, such as the Samsung Galaxy Gear [Sam], or accelerometer only in Fitbit and Sony platforms [Son] gives rise to the questions of what sensory devices are necessary in a single smartwatch platform to make it a successful exercise tracking device for software platforms, such as those available on the market by companies such as Focus[Foc] (who provided the desired workout routine for this work) as well as similar platforms in research fields[SBV11]. This work will attempt to analyze the performance difference for five workout exercises using a

Samsung Galaxy Gear smartwatch platform in order to count repetitions for use in a personal gym coach application, by leveraging the contextual information of a given workout routine to classify movements and use said classification as a repetition counter. It will analyze the difference between the accelerometer, the gyroscope, and the combination there of, as well as optimization techniques to reduce the computations necessary to accurately count the repetitions of such exercises as a guideline to future applications using wrist-worn devices.

8.1 Related Works

While exercise routines and detection are a popular field of research, ranging from quality to energy expenditure, several related works aim to detect particular exercises and count their repetitions. In [SBV11], a mobile phone platform is used to aggregate data from two custom accelerometer devices operating at 100 Hz. These sensors, on the arm and the leg, as well as a heart rate sensor on the chest, show the ability to count repetitions of exercises and calories burned from the increased heart rate through the use of a six-dimensional Gaussian distribution for each movement class. They show an accurate system from 71% to 100% for 16 various activities using said system and counting the repetitions of an activity through a peak-detection method, which may or may not be a robust method when used across different populations. This work will attempt to approach a similar problem from that of [SBV11] by using the classifiers to identify not the general motion patterns but instead use the classification to accurately count the motions detected. Further, this work will identify the key features in classifying free-weight and body-weight exercises.

In [PHK13], a smartphone is attached to the users arm for unconstrained exercises (those capable of being performed in any environment) and placed on the weights of a machine in a constrained (gym) environment. They use a method

based upon dynamic time warping to identify the activity, and possible repetitions, then count off of this information. In particular, they indicate that the DTW is too time-consuming to be performed real-time, and thus, pre-filter data by using peaks of similar height and a threshold window to only consider given acceleration windows. The similarity method appears robust in their results resulting in perfect precision and high recall (above 93% in all cases) in their unconstrained environment. This paper, similarly, deals with an unconstrained environment and movements monitored on the wrist (5 in their case) and attempts to find an even less-computationally intensive procedure as DTW to classify its movements. Further, no gyroscope exists in their platform, but will be analyzed here similarly to [PEA07]. This work attempts to classify each individual movement repetition instead of the entire data set, similar to what was achieved in [PHK13].

In [CCC07], two different classifiers are used in determining a workout exercise then counting the repetitions of said exercise. Using a single accelerometer on the back of the hand and one on the hip, they are able to accurately identify most work out routines with most exercise accuracies in the $90 - 100\%$ range using a Naive Bayes classifier or a Hidden Markov Model. Further, after identifying the movement, they use either a peak detection algorithm, or the state transitions of their HMM in order to count the actual repetitions and have an error rate under 10% for most exercises and under 20% for all. They test their system in two different cross-validation schemes, with user-specific models at different weights and in a leave-one-subject-out cross validation for robustness, showing similar results. This work attempts to build upon results of [CCC07], by assuming first that the motion may be identified correctly, either through their method, or through knowledge of the workout routine to identify the activity, and use a classifier to improve upon the counting error rates. Further, they eliminate the gyroscope, stating cost issues. This work will use a gyroscope to identify accuracies so that, if cost is not a concern, results can be evaluated to determine

Figure 8.1: Device worn on wrist for exercise with dumbbell

if it is necessary to include all available sensors in a smartwatch platform. Finally, instead of using the entire data set to separate the differences between motion classes, this work looks at individual repetitions and attempts to classify each.

In [XPK13], a large-scale system is created for annotating data in a large database system for use in ground-truth labeling. This works shows the need for finding an automated process to simplify ground-truth labeling and annotation for data systems in wireless health applications. This work will examine a similar issue, in determining the appropriate movement annotations and duration in an automated manner.

8.2 Methods

Many workout routines start and end the same way with a defined set of exercises. This work aims to leverage this information to build classification and repetition counting into a single algorithm to quantify these exercises. This effort will focus

Table 8.1: Movements selected and whether they use a dumbbell or not

Movement	**Uses Dumbbell**
Bicep Curls	yes
Crunches	no
Jumping Jacks	no
Push Ups	no
Shoulder Lateral Raises	yes

Table 8.2: List of Calculated Features per Axis

Feature	**Count**
Amplitude	6
Median	6
Mean	6
Maximum	6
Minimum	6
Peak-to-Peak	6
Variance	6
St. Dev.	6
Root Mean Square (RMS)	6
Skewness	6
Derivative Mean	6
Derivative St. Dev.	6
Derivative Variance	6
Derivative RMS	6
Axis Correlations	15

on building binary classifiers for each specific repetition of each movement, classifying individual actions instead of the repetitive pattern of the general exercise. Using a smartwatch platform that consists of an accelerometer and a gyroscope [Sam] operating at 50 Hz, shown heuristically to be +/- 2g and +/- 200 degrees per second respectively, worn on the left wrist. Fig. 8.1 shows a user wearing the watch during a workout routine while wearing the watch. A classification algorithm is run to determine accurately the workout routine movements listed in Table 8.1. When looking at the watch face (e.g., when looking at the watch

to read the current time), the x-axis points forward to the top of the watch, the y-axis points to the right, and the z-axis faces straight up from the face of the watch. While some of the routines could be done with or without dumbbells, the exercises selected were suggested by Focus[Foc] as a core set of initial exercises, using a 10 pound dumbbell. The end goal of the training and testing sets is to then develop an algorithm that is capable of recognizing these activities while running on the same watch.

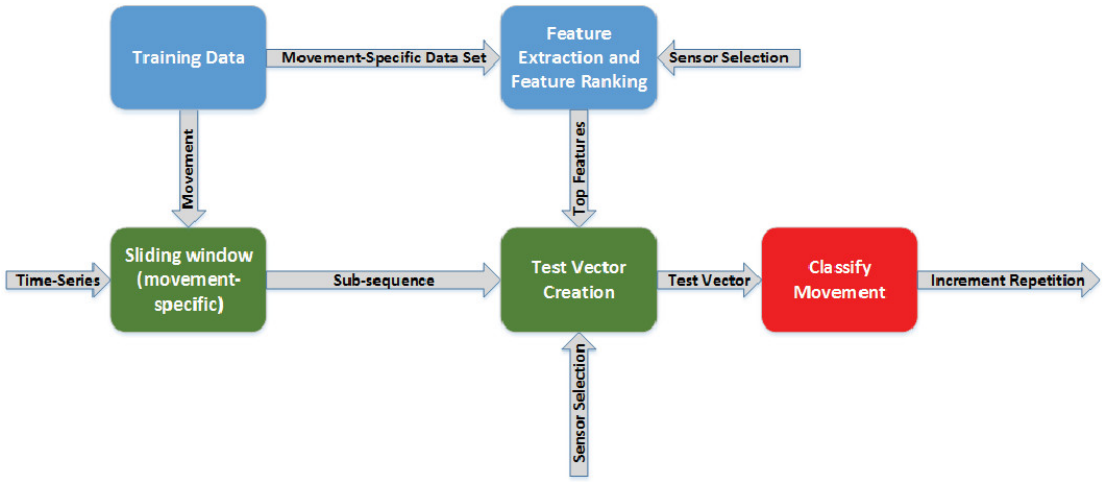

Figure 8.2: Feature Ranking and Classification Method for Identifying Repetitions

8.2.1 Data Set

Data was collected on 12 participants, male and female, ranging from age 23 to age 38 and thus a lower weight was selected to account for all users. Each user was asked to perform each given activity ten times. The data was collected with the watch on the left hand (although the algorithms can be adapted to either hand). Finally the users were also asked to simply perform no movement, random movement, and various arm positions in order to develop a no-movement class for an application. Random windows from this signal were selected for the training

sets.

8.2.2 Segmentation

When the data is collected, work must be done to find the appropriate segmentation, sizing of movements. This can be done through several methods. The most tedious is to manually annotate the data for a supervised learning algorithm. This method can be exhausting with large data sets. Automatic clustering is an option as well, although unclear from feature values and complex movements if such movements are unique repetitions or smaller components of larger pieces. Finally, some of the detailed motions considered in this dissertation are often short burst, singular movements. However, when collecting a data set, it is possible to take advantage of the collection strategy to create a repeatable environment by having the users perform each task multiple times and separating those files. Using contextual knowledge, it then becomes possible to identify not only the number of repetitions but also the appropriate length for each movement. This is done using the signal autocorrelation, modified to work for each individual axis. The autocorrelation of a discrete, real signal s_t with itself is represented as

$$autocorr(\tau, s_t) = \sum_t s_t * s_{t-\tau} \tag{8.1}$$

where s_t is the signal, $s_{t-\tau}$ is the signal at lag time τ. The two signals will be considered a repetition if $autocorr(\tau)$ is at the same value as $autocorr(0)$ (in other words the signal with itself). However, in a multichannel signal such as one presented by having accelerometer and gyroscopes working together, not every axis will represent a repetition. As a result, a maximum window size can be determined by the length of the training example, with the expected number of repetitions for a given move type known by those recording the data. An algorithm 2 can be run to identify repetitions on a given channel, the ideal move size for that

channel, and then decide among the channels the appropriate size. The algorithm can be run as is, or, if contextual knowledge is available (e.g. bicep curls work along one axis) then a reduction can be taken by modifying the channels to search in $getChannels$. Finally, if a certain percentage of given channels detect motion the $getWin$ function determines the window size. This can be done by taking the mean window size, the majority window size, or a sorted order of each channel's importance in a given move type, for example. Once this is set, a given window size can be determined that optimally captures each repetition of each movement, which can then be used in the classification algorithms to follow.

Algorithm 2: Algorithm for an Appropriate Window Size for a Movement

Data: $D(t, C), maxWin, expRep, moveType$ where $D(t, C)$ is data with length t and C channels

Result: $idealWin$, ideal win size for given $moveType$

```
begin
    WinSizes ⟵ ∅;
    for i ∈ range(1, maxWin) do
        for each c in getChannels(moveType, D) do
            autoc_i ← autocorr(i, D(t, c));
            peaks ← IdentifyPeaks(autoc_i, moveType);
            if size(peaks) == expRep then
                WinSizes ← Distance(peaks);
    idealWin ⟵ getWin(WinSizes, moveType);
```

8.2.3 Training

8.2.3.1 Preprocessing

Once the data was collected, the appropriate window size for each movement needs to be determined. This was done by calculating the average window size

spit out by the $getWin$ function. For each move $m \in M$, where M is the set of all movements, each person $p \in P$, where P is the set of all people, has a move size determined by:

$$w_{m_p} = \frac{1}{n} \sum_{i=1}^{n} s_i \tag{8.2}$$

where s_i is each individual move sample defined by the annotated start and end points. This average is then calculated and the movement average is determined by:

$$w_m = \frac{1}{|P|} \sum_{j=1}^{|P|} w_{m_p} \tag{8.3}$$

This average window size is then used to alter the end point of each movement annotation determined by the autocorrelation to give the general movement size of each exercise. Since the workout routine is known, the context can be leveraged to build binary classifiers. As a user goes through a weight training program, the sets are defined. As a result, every model can be built to identify if a given movement window is a push up or is not a push up. Using this context information a series of features are extracted for each move.

8.2.3.2 Feature Extraction

The total list of features shown in Table 8.2 are extracted for each training sample. These features listed are calculated on each axis and the total number of the features is shown in the 2nd column. Thus, 6 features equates to the mean being calculated on each axis of gyroscope followed by each axis of the accelerometer. Once this is extracted for each sample for each person and each move, the number of features used in the testing must be reduced. Clearly the 99 features calculated would over-fit the data. Also, for computational performance it would be better to only need to calculate a subset of the features. Four testing configurations were created to evaluate the performance of the accelerometer, the gyroscope and the best axis. These configurations create 4 training sets per activity. The first is

using features from only the accelerometer, the second only the gyroscope, the third is the combination of features from accelerometer and gyroscope, and finally the fourth is identifying the single best axis from the third set, and using only those features. Since there are significantly larger quantity of negative examples than positive for each movement, the negative samples are randomly removed to balance the training set size. Finally, then, the training and testing sets are created to be used with feature ranking and selection, then classification. The training and classification flow is shown in Fig. 8.2. The features, once extracted, are normalized. In order to make the correlation features more informative, the signals are normalized for the extraction of those features.

8.2.4 Testing

8.2.4.1 Cross-Validation

In order to test the model, the Weka [HFH09] platform was used in order to feature rank, feature select, and cross-validate the results. The top six features were used in order to ensure that the accuracy information will not be as a result of over-fitting the training data. Six features were chosen as the minimal subset that achieves what was deemed acceptable performance. Using a correlation ranker, the top six features were used in each set and the results were past through a 10-fold cross validation to test robustness to variability. Several classification schemes can be tested in Weka, and in this case random forests, decision trees, SVM, and Naive Bayes classifiers were compared.

8.2.4.2 Testing in a Real Setting

For the test data, a sliding window of points, equal to the average window size calculated for each movement is chosen. Given a time series signal $T = t_0, t_1, ...$ where each index is a sample at 50Hz, a subsequence t_i is then selected as follows:

Table 8.3: Average Accuracy and Area Under the (ROC) Curve (AUC) of Each Movement in Cross-Validation for Each of the Test Settings

Movement	**Accel. Only**	**Gyro. Only**	**Accel+Gyro**	**Best Single Axis**
Bicep Curls	92%, 0.97	87%, 0.91	92%, 0.98	87%, 0.94
Crunches	98%, 0.99	86%, 0.92	98%, 0.99	93%, 0.94
Jumping Jacks	88%, 0.94	77%, 0.86	89%, 0.95	84%, 0.94
Push Ups	96%, 0.99	87%, 0.93	96%, 0.99	95%, 0.99
Shoulder Lateral Raises	88%, 0.95	90%, 0.94	90%, 0.96	90%, 0.94

Table 8.4: Selected Features of Each Best Axis for Each Movement

Movement	**Axis**	**Features**
Bicep Curls	a_z	minimum, median, mean, amplitude, maximum, root mean square
Crunches	a_y	median, mean, maximum, minimum, variance, root mean square
Jumping Jacks	a_x	minimum, root mean square, median, st. dev., mean, peak-to-peak dist.
Push Ups	a_z	median, mean, maximum, minimum, variance, root mean square
Shoulder Lateral Raises	g_y	root mean square, st. dev., minimum, peak-to-peak dist., variance, maximum

$$t_i = (T_i, T_{i+w}) \tag{8.4}$$

where w is the indicated move window as designated from the training set, assuming enough points remain in the time-series. This subsequence is then used to extract features, resulting in a feature test-vector, as in:

$$F(t_i) = \{f_i | 0 < i \leq MAX_f, each f_i \in F_s\} \tag{8.5}$$

where MAX_f indicates the largest number of features collected, then each f_i is from the sorted order of features as ranked by the feature ranking algorithm, denoted F_s, such that f_0 is the highest ranked feature, f_1 the next and so on. A sliding window is then shifted over the testing data testing each activity. Every time an activity is classified as an individual sample, then it is counted as a repetition, the window cleared and the time-series sub-sequence jumps past this point and continues (in order to avoid re-classifying the same repetition when the window is slid only one point). Otherwise, the overlap for the sliding window is every point. For every yes classification result a counter is incremented to count repetitions. Thus, high classification accuracy of individual movement repetitions will equate with low error rate in counting.

8.3 Results

8.3.0.3 Accuracy and AUC

In order to evaluate the results the correlation tool to rank the features in the Weka environment was used, followed by a random forest classifier due to its highest classification accuracy. The top six features were picked for each movement. As seen in Table 8.3, all testing configurations provide strong results. This table shows the classification accuracy as well as the area under the curve (AUC) of

the Receiver Operating Characteristics (ROC) of each movement in each of the four test settings. The gyroscope seems to add some strength in a few of the movements and thus, for a system that wants accurate counting, is necessary in order to have the highest classification accuracy. The jumping jacks, in particular, have curves that occasionally saturate the 2g accelerometer, thus, most likely exhibits the lowest classification accuracy. Thus, by using context information to create binary classifiers, a high classification rate for each example movement is found, thus indicating a low error rate on counting for each of these movements in a trainer application.

8.3.0.4 Selecting the Best Axis

In order to reduce the computational load, a trade off between the best single axis and the full data set can be compared. Two examples of movements and their ROC curves are shown in Fig. 8.3a and Fig. 8.3b, showing a combination of accelerometer only, gyroscope only, accelerometer and gyroscope, as well as the best selected individual axis. Note, in these examples, a combination of accelerometer and gyroscope provide the strongest result. The best individual axis and their features for each movement are listed in Table 8.4, listed in order of strongest first. While the features are similar, a binary classifier allows for the selection of different features for each movement. Thus, contextual information can optimize a system to use the ideal features and reduce computational load. Further, by only attributing a specific axis, less data can be computed or transmitted, and further, uniaxial sensors can be chosen in place of triaxial if so desired. For certain movements, this procedure yields accurate results that allow for the reduction of the computational load. Fig. 8.4 shows one such movement where the accelerometer only and the accelerometer plus gyroscope configurations use the same best features as the single individual axis of the accelerometer. This allows for the same ROC curve and same AUC as a result.

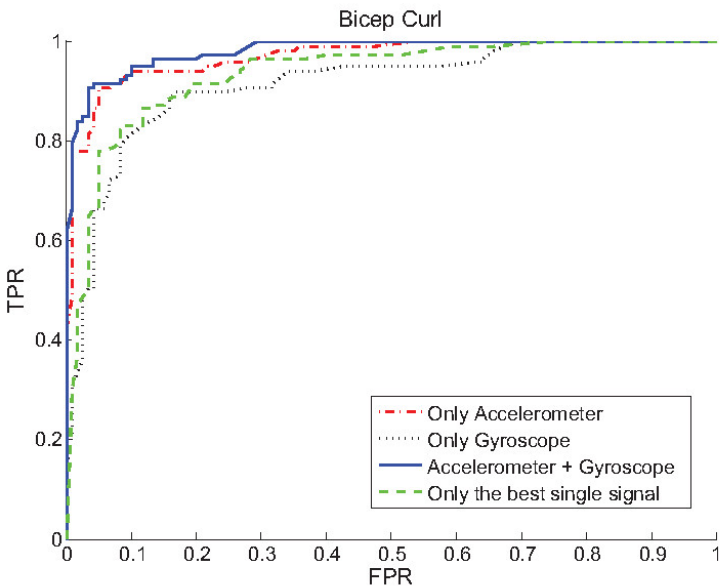

(a) ROC Curves for Bicep Curls

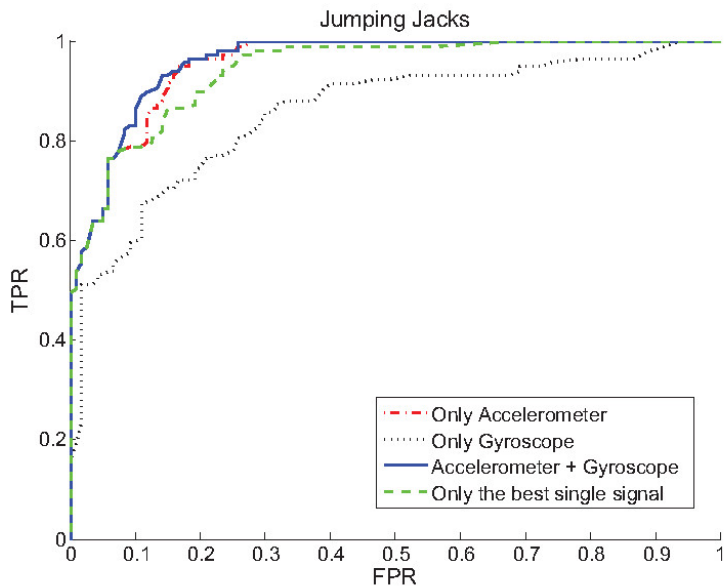

(b) ROC Curves for Jumping Jacks

Figure 8.3: ROC Curves for (a) Bicep Curls (b) Jumping Jacks and (c) Push Ups showing Accelerometer Only, Gyroscope Only, Accelerometer + Gyroscope, and Best Axis Classification Methods

8.4 Future Work

The system shown here provides ample opportunity to further test and evaluate workout routines. Further feature extraction techniques can be used to determine stronger best-axis information to reduce the computational load, or to ensure a minimal subset of considered sensor data, as well as applying contextual information or calibration steps to create a user-centric platform. Once this is developed, further detailed motions can be analyzed in a similar manner, such as those needed for rehabilitative exercises. As the workout routines are increased this may be more important and so more exercises should be considered, as well as the real-time responsiveness and user-experience considered. Further, the algorithm itself should be re-implemented on to the watch and a user-experience trial performed to see if users appreciate the accuracy and speed in which the repetitions are counted, and the actual delay values calculated. Such systems should also be tested for robustness in terms of varying weights, in particular considering the trade-offs between movement-specific models and weight-class specific models and how this contextual information may be passed into the system. Finally, the addition and fusion of other sensors, such as heart rate, may provide valuable differentiating characteristics and should be further investigated in the realm of exercise-sport applications.

8.5 Discussion

Smartwatches show a new realm of activity sensors for personal health monitoring. Some have accelerometers while others also include a gyroscope. As wrist-worn wearables increase in popularity, the investigation of the sensor platforms needed and computational considerations for the response time of such counting algorithms must be considered, ranging in applications from general monitoring to real-time exercise repetition. This work presents the feature extraction and se-

lection necessary to not only identify but at the same time count repetitions of free-weight and body-weight exercises by leveraging the context of the workout routine in order to develop strong classifiers. While the accelerometer alone is fairly consistent in classifying the individual repetitions of the motions, the gyroscope does improve the classification accuracy by having an average area under the curve of .974, up from .968 using only the accelerometer. Adding a gyroscope with the accelerometer increased the average AUC from .968 to .974, increasing the accuracy of specific movements as much as 2%. The best single axis method achieves an AUC of .95 showing that a reduced method can provide the necessary accuracy to accurately count repetitions in a lower computational power setting. The segmenting approach here for annotating of data signals runs an $O(n^2)$ algorithm for data processing. It should be investigated in further systems whether such a system is feasible in comparison to systems such as those in [XPK13].

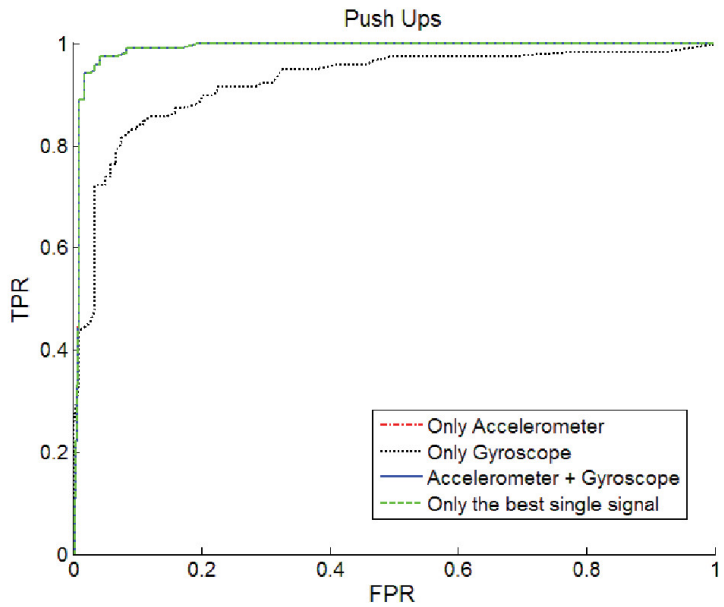

Figure 8.4: ROC Curves for push ups. Note that three configurations share the same curve

Acknowledgment

The authors would like to thank the team at Focus, Cavan Canavan and Grant Hughes, for their input on ideal initial exercises and the workout routine.

CHAPTER 9

Dynamic Optimization for User-Centric Exergaming

Accurate physical activity monitoring of human motions has importance in many applications. These applications include medical monitoring systems [PSM07], personal training [OF06], and personal gaming [BCF10]. The latter, personal gaming, is of particular interest, since it spans the realm of different forms of human activity monitoring. Such monitoring can be on general movements [RDM05], general pose-based activity monitoring for gaming [SGB09a], or exergaming, which is activity monitoring for gaming with targeted health applications [GPH09]. This latter application realm is a target solution for childhood and adult obesity, a growing epidemic [WBL08a] and economic burden [FKT12]. Indeed, work in [FKT12] estimates that, over the next two decades, there will be a 33% increase in obesity and 130% increase in severe obesity in the United States and proposes methods to save almost $550 billion in medical expenditures over that time. Exergaming has potential to help both the health epidemic as well as relieve some of the health care economic burden. As exergaming expands, its application realm has grown, increasing the need for more immersive and responsive exergaming applications. User adherence to a game drops with noticeable latency[CHL06]. Latency, in fact, can alter the user experience in a negative fashion in many application realms, for example in user performance in gaming consistency and experience [WAH10]. Ultimately, this shows the importance of analyzing each aspect of any interactive system and the importance of the human element in that loop. Many activity

monitoring systems work on long and repetitive actions to find general patterns, thus not tailored to a particular user [MHK06].

This paper presents a system and optimization approach for playing a realistic soccer exergame, based on a system developed in [MNL14]. That game consists of a sensor system for user input, a classification method for fine-grain activities, and a visual interface on a mobile tablet for gameplay and feedback. This paper dynamically adjusts the algorithm presented in [MNL14] through the use of a derivative free optimization approach in order to define a unique exergaming experience for each individual user.

9.1 Related Works

9.1.1 General Daily Living Activity Monitoring

In [RDM05], hereafter known as RDML, a method is presented that uses a combination of the mean of accelerations, the standard deviations, the energy or power expenditure, and a correlation between the channels of an accelerometer worn on the body for general activity monitoring of standing, walking, running, climbing up stairs, climbing down stairs, sit-ups, vacuuming, and brushing teeth. Features are extracted from windows of size 256, which is five seconds of repetitive, cyclical movement data. Fine-grain movements that fit in much smaller windows could be problematic in such a scenario. Ultimately, these calculations do not properly classify fine-grain movements[MNL14], but are still presented as applicable to the dynamic optimization approach presented here.

9.1.2 Exergaming

There exist many vision-based exergaming approaches. The Microsoft Kinect camera system and SDK for human motion monitoring has proven effective [RKH11]

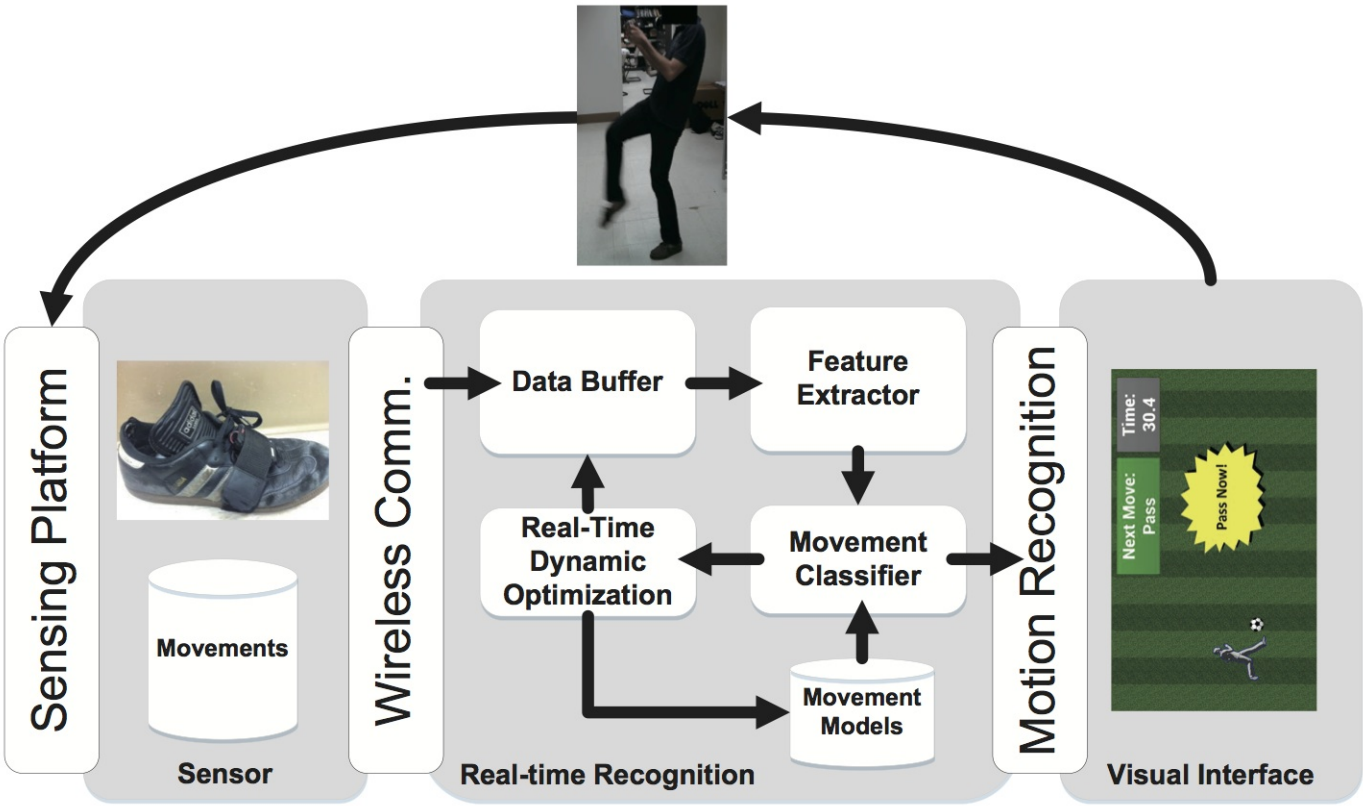

Figure 9.1: System Architecture of User Playing Soccer Exergame

at monitoring activity. Skeletal gestures are combined principal component analysis (PCA) [RKH11] to classify movements. The goal of this work is to develop a general system that can be used in a mobile setting that approaches the classification of many movements and delay in a similar fashion to those presented in [RKH11] that use specialized hardware for processing.

9.1.3 Derivative Free Optimization

Dynamic optimization problems encompass the range of optimization problems where the constraints themselves often require complex solutions or simulations [RAG11], and indeed, might be learned or sampled instead of known. Often, work involves evaluating different solvers for smooth or noisy constraint functions [MW09] or interpolation and approximation [KCK11]. This paper takes the latter approach, by sampling the function for the constraint instead of knowing it outright. [CLL09] uses a case based dynamic window to improve results and this work will take a similar approach for dynamically adjusting the window size, not for rule refinement, but for real-time responsiveness.

9.2 System Architecture

9.2.1 Exergaming Overview

This section describes the full system architecture presented in Figure 9.1. This figure shows the user playing the exergame. It starts with a motion from the user, follows to the sensing platform and movements required. From here, data is wireless transmitted to a computing device (e.g., a tablet computer). This computational device contains the recognition engine as well as dynamic optimization tools. When movements are detected they are passed to the game and the information is returned to the user via a visual interface. This loop needs to happen

Table 9.1: Movements Captured

No.	Move	Description
1	Back Heel	Kick Backward
2	Behind Foot Pass	Pass to left behind left leg
3	Chip/Lob	Lift a ball up and forward diagonally
4	Fake Shot	Fake a kick
5	Flick Pass	Flick a pass outside of foot to right
6	Full Swing Shot	Full powered shot
7	Laces Shot	Mid powered shot
8	Quick Shot	Low powered shot
9	Curved Shot	Placed shot, using side of foot
10	Through Pass	Diagonal pass
11	Pass	Pass direcly left
12	Step Over Move	Swing foot around ball in circle
13	Side Step	Step on ball, roll it to the right
14	Run	Step
15	Sprint	Step 2x speed of run

in a responsive fashion for an appropriate, interactive user experience.

9.2.2 Soccer Exergaming

Work in [MNL14] develops a mobile, interactive gaming system for classifying what this paper denotes as fine-grain movements. The system developed in this paper incorporates an optimization approach to such a system, in order to develop a more responsive, unique exergaming experience. The classification engine in that game is based on twenty four users collecting the movements listed in Table 9.1, all referenced from the right foot wearing a three-axis accelerometer and three-axis gyroscope sampled at 100 Hz. A nearest-neighbor classifier based upon reconstructions from a principal component analysis are used to find the appropriate movements, based upon a 330 point sliding window.

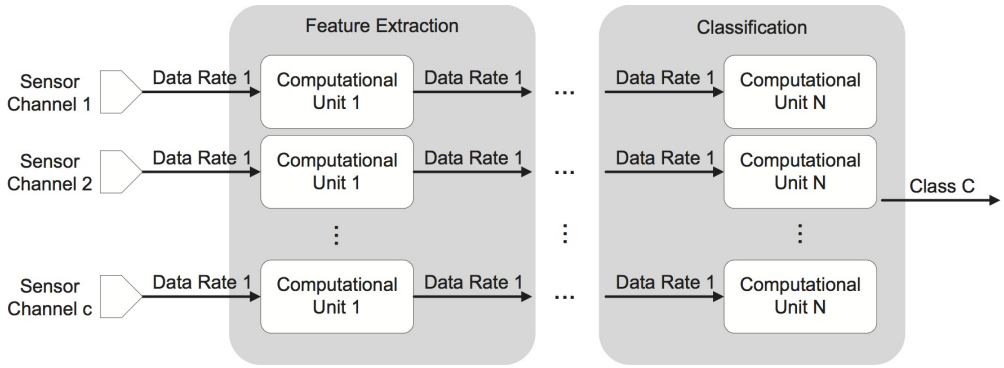

Figure 9.2: Sensor System and Computational Units

9.3 Dynamic Optimization

9.3.1 Problem Formulation

At its basis, the trade off between classification accuracy and latency can be modeled as an optimization problem. This work will use the Micro F1 score to measure overall accuracy, as it accounts for the precision and recall of a system. Using accuracy alone could bias a solution since a data set this large might result in high accuracy even if it classifies 0 movements because the number of true negatives would have been high. The F measure is calculated as follows:

$$F = 2 \times \frac{P \times R}{P + R} \tag{9.1}$$

where P is the precision and R is the recall rate of the learning model, and the factor of 2 puts the end result back in the familiar 0 to 1 range of precision, recall, and accuracy. Implicitly, F can be defined as $F(\delta)$ where δ is a delay factor for the system. This is because the more data present the better the system can classify; however, the more data used before a classification is made, then the longer the feature extraction steps take before classification. If one considers a system layout, as in Figure 9.2, where each sensor channel (from

possibly different sensors) provides data at a given rate to different computational modules, an eventual classification is made when considering the combination of the data provided to each computational unit and the delay with which its computation proceeds. These values are, however, very application dependent and need to be learned. In fact, delay can be considered as the following:

$$\delta = \max_{c}\{\sum_{i}(\alpha_i \times r_i \times k_i) - \beta\} \tag{9.2}$$

where r_i is the data rate to a given module, k_i is the computational delay for that operation, such as feature extraction, α_i is a scaling factor for a given module's importance to the overall delay, β is the real-time length of the motion (e.g. two seconds to shoot a soccer ball, delay only matters after the motion is completed at the point the ball should leave the foot) and the delay is the maximum such delay over all the sensor channels necessary to calculate the class C. From this, F as a function of δ becomes more clear and this presents the optimization problem of:

$$\min \delta \tag{9.3}$$

subject to

$$F(\delta) \geq \tau \tag{9.4}$$

where τ can be some predefined accuracy threshold. The difficulty comes in determining both the $F(\delta)$ and τ because the former is learned, and the latter is application specific. Figure 9.3 shows how the F measure changes as the window size changes. In particular, window size is one of the features that represents overall delay, and it seems here that the window size can range from about 150 to 330 samples with similar accuracy results. If a move takes 2 seconds (or 200 samples) then the most accurate point also represents the highest delay, a trade off the user might not want, when the 0 delay point has reasonable accuracy.

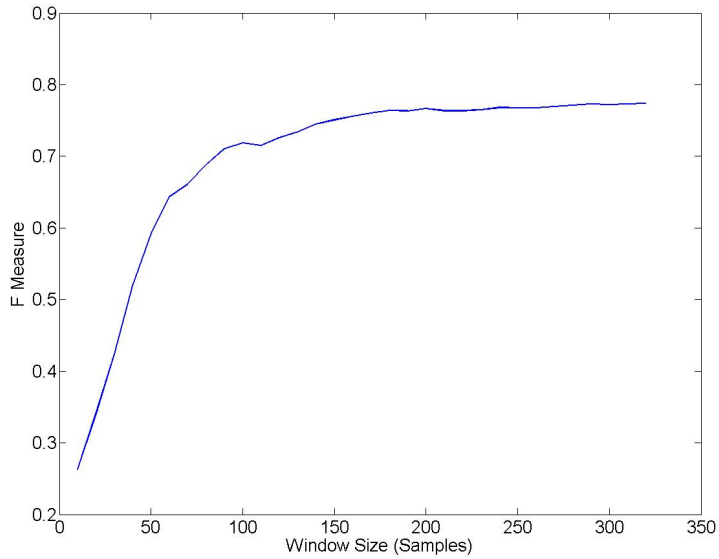

Figure 9.3: Sample F Measure as scaled with Window Size

9.3.2 Dynamic Adjustment

As shown in Figure 9.4, different curves result from leaving specific users out of the training in a cross-validation setting. In other words, the way in which individual user might perform particular actions either conforms to the model or shows their necessity in the training set for a more robust model for other users. The right hand side of the constraint then becomes $\tau(\delta)$, which is also sampled and learned. From this, a dynamic solution must be derived. A method to alter the delay must be chosen that is adaptable while in use. In particular, a dynamic window-size algorithm is chosen, where the constraint optimization problem tells us the left hand side wants to shrink the window size as much as possible while the right hand size is attempting to resist this shrinking. The learned model has a maximum classification accuracy at the full window size of 330 points. The dynamic windowing compares the classification result at the smaller sizes with the full size that is most accurate on the largest number of subjects. If a series of

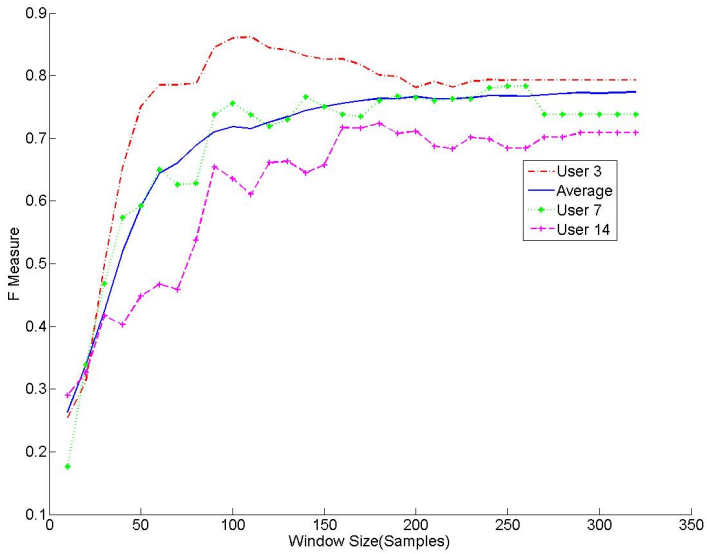

Figure 9.4: F Measures as scaled with window size for average and for different users as test subjects

matches are found, where reduced window sized classification matches full window sized classification, the window size is shrunk. If a series of mismatches are found, the window size is increased again. In both cases, the threshold to allow for re-sizing can be adjusted per application use. In this case, since the latency is of greater concern, a series of only five matches in a row will be enough to shrink the window size (a number deemed an aggressive shrinking factor that still clearly represents a pattern of correct results). This aggressive shrinking must then be robust enough to regrow the window if it shrank in error (e.g., the user only performed one action correctly ten times and it was a small window sized action). Thus, five misses (based on the F measure as a value clearly higher than a simple misclassification or two) in a row will set a lower boundary on the possible window size and the window is grown, half way back to the previous set. As a result, a form of binary search results in eventually settling on an operating window size that is user specific. If the accuracy is of greater concern, the system can be

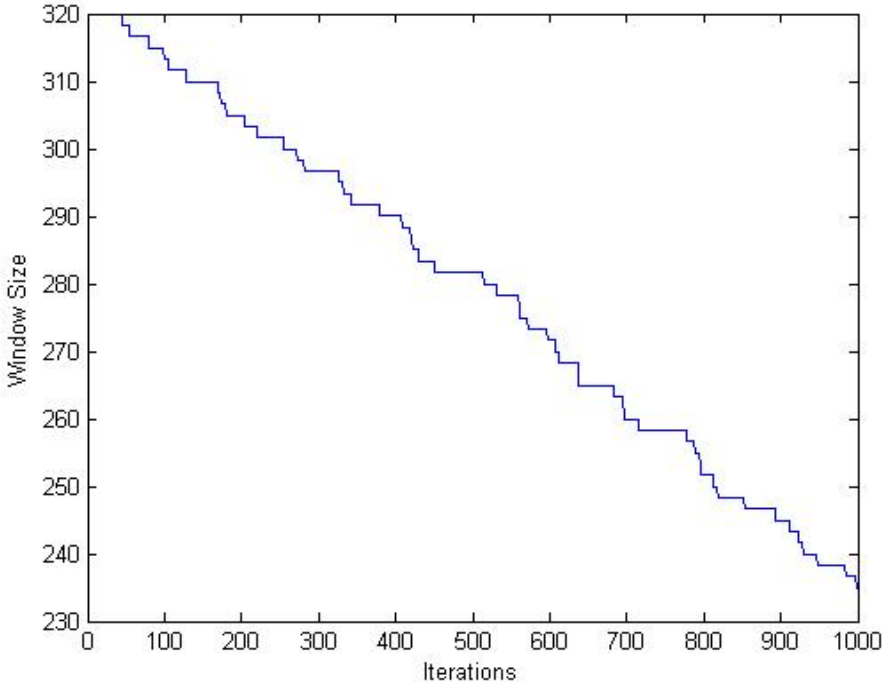

Figure 9.5: Dynamic windowing of RDML algorithm

altered to shrink less aggressively and re-grow more aggressively, depending on the application and the learned model of $F(\delta)$.

9.4 Results

A leave-one-subject-out cross-validation was used in order to interpret the results of this interactive system, and was applied to an algorithm for fine-grain motions as well as one designed for general daily living to show the adaptability of such a method.

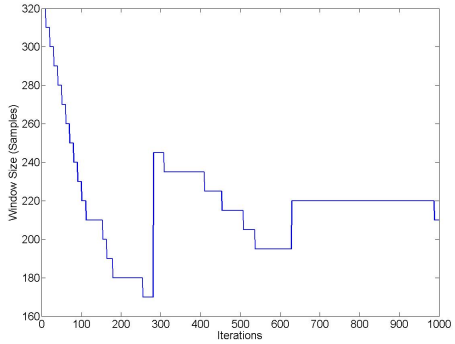

(a) A user with large window size

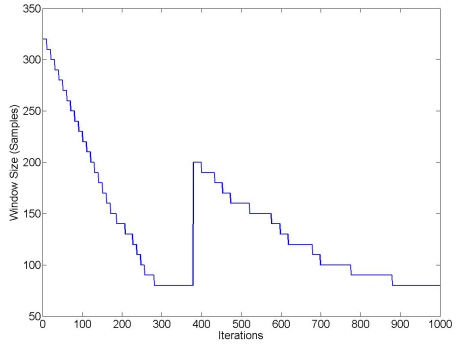

(b) A user with a smaller window

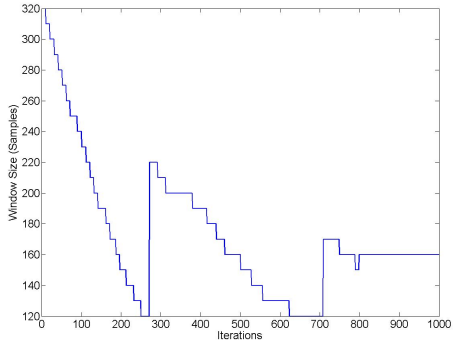

(c) A user with an average window size

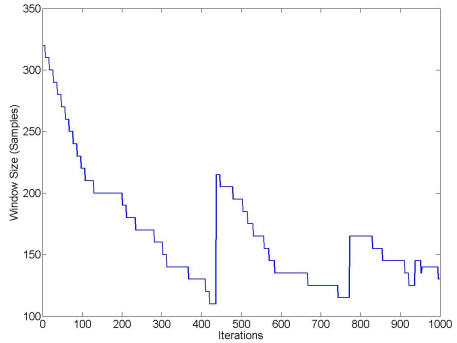

(d) Another user with an average window size

Figure 9.6: Dynamic windowing of Soccer Exergame for (a) a below average user (b) an above average user (c) an average user with expected window size and (d) another user with an expected window size

9.4.1 RDML

The dynamic windowing algorithm was run a data set of moves from RDML to show a slow shrinking result in a system where accuracy is considered more important than delay, needed a greater number of correct matches. This shows the system is adaptable to general daily living and can be used to convert such a method to a real-time one. Figure 9.5 shows that this algorithm can slowly shrink to what is even less than the 256 point window size that is chosen by RDML.

9.4.2 Soccer Exergaming

Figure 9.6 shows four such runs of four users then the iterations run to dynamically adjust the window sizes (where iterations are classification queries, including those that result in no movement, as defined in [MNL14]). This shows each individual user has a different threshold, as explained previously. Since this value is learned and sampled instead of defined, each user receives a tailored experience with a particular-sized window fit for that individual's performance with the learned algorithm. Note that Figure 9.6a settles on a much larger size than might have been picked simply looking at Figure 9.3.1, while Figure 9.6b is a much smaller window size. Figures 9.6c and 9.6d show an expected window size of 160 points. However, all show different stabilizing points for the algorithm and that in some cases this algorithm picks a much smaller window size than would be expected when looking back on Figure 9.3 and seeing the optimal window size versus F measurement.

9.5 Future Work

The system described in this paper shows a soccer exergame developed on a tablet using a wearable sensor on a user's foot. The game has an algorithm tailored to

fine-grain activities that dynamically adjusts the window size. Further, the formulaic representation of the problem allows for adjustments to other parameters, including data rate, sampling rate, feature extraction and selection, and/or power considerations to fine tune the human interaction.

9.6 Conclusion

This paper presents an interactive exergame system with a method for optimizing the detection of fine-grain physical activity of the human body in a real-time environment for each individual user. It defines fine-grain movements to be those that are often quick and singular segment movements, instead of the standard cyclical, repetitive movements of previous human physical activity monitoring systems. The derivative-free optimization metrics can be adjusted on a per-application basis to help users design an appropriate system with the trade off of delay versus accuracy in mind. This gives each user a unique experience in a generalized exergame.

CHAPTER 10

User-Specific Models and Mobile Exergaming

10.1 Desired Outcomes

With clinical outcomes and user-centric optimization developed, a comprehensive exergame that achieves a user-specific experience is desired. The previous chapter covered a dynamic solution to optimizing delay versus accuracy on the trained models. However, more work can be done to train the models in a user specific way. This chapter covers such an approach, examining the ability to create more accurate models for users as they play a desired exergame. Consider again the algorithm 1 from the introduction. As referenced in the previous chapter, work can be done to create a more user-centric modeling to enhance the user experience. However, while modifying the model type used and feature sizes optimizes the results by selecting a particular model in a given context, it does not actually modify the modeling itself. This chapter will cover an approach to improve a model for any given user of a system as well as improve classification of the fine-grain movements presented earlier, specifically the line involving the $updateModel$ function. This chapter presents SoccAR. An Augmented Reality Soccer exergame based upon the Temple Run platform to incorporate a dynamic mobility game with detailed sports-type actions.

10.2 Background and Related Works

With a worldwide trend of inactivity, with 15% of men and 20% of women surveyed in 51 one countries being at risk for chronic diseases [GOS08], monitoring of physical activity is of utmost importance. Physical inactivity causes not only health concerns, as discussed in chapter 1, but as shown, the inactivity can causes many diseases as well as increase healthcare costs for each individual[GBS04]. Exergames and other physical activity monitoring systems have been developed to track the users and identify their activity levels with different degrees of accuracy and application [ALM10a][AXL13][KLH10][LL13]. However, generalized models may not always be the most effective. With exergaming, the user has the benefit of knowing whether movements are classified as correct or incorrect based upon the context of the game. As such, miss-classifications can be dealt with by updating models for user-specific modeling.

10.2.1 Incremental Learning

Work in [CP01] introduces an incremental support vector machine. These systems apply updates to a defined model at a given period of time. [LTB12] uses this method to update human recognition in video applications to improve accuracy. Further, [HB11] and [SKM06] use contextual information and a method called active learning to update the models based upon contextual information. [SVS08] uses contextual information for co-training and re-training. All of these methods use context information to create a subset of labeled data in order to reduce training time. This work builds upon these ideas, not necessarily concerned with reducing the training time, but using active selections of training data to improve the accuracy of a given system by analyzing the data in an on-line fashion and re-training any model when determined necessary.

10.2.2 Sports Exergames

Exergaming involves incorporating health information and healthy activity into gameplay with the hope that the entertainment factor of gaming coupled with the health information can make better use of the time spent playing games. Many exergames exist that address particular aspects of personal health. Work in [ALM10a], for example, presents a platform for developing an exergame to assist in stroke rehabilitation using accelerometers, representing a class of exergames with a specific goal and set of actions in mind. Work by [PHL12] develops an entertainment-based exergame with multiplayer aspects to encourage competition. While successful, ultimately the game is an adaptation, to standard exercises, namely, using a hula-hoop, jump rope, or riding a stationary bike. This stationary game presents the possibility for active exergaming but does not address the mobility or correlate to a gaming environment well. Similarly, work by [GLN12] demonstrates a motion-based game for exercise and entertainment of adults, but uses a Microsoft Kinect which, while allowing freedom in a particular space, does not allow for a game that can be played anywhere and across any distance. Finally, a set of mobile games, such as those in [Mac12][KMK11] allow for gaming in any environment by using a mobile computing device (such as a smartphone) as the controller. So while these games allow for freedom in environment, they do not address a wide range of possible motions and exercises.

One such work, in particular, addresses the concerns of the related games in development to build a solid foundation for exergame demonstrations. Work in [MNL14] (chapter 5) describes the steps necessary to building an exergame that results in healthy and enjoyable gameplay with accurate motion classification. A tablet-based game is built from a list of soccer movements collected to develop an energetic and fast-paced game. A game-specific movement classification algorithm is created based upon a principal component analysis. This algorithm classifies movements shown in the previous chapters to produce exercise-levels of inten-

sity by guaranteeing a certain level of metabolic output. This paper introduces a framework for creating successful mobile, wearable exergames by outlining a general procedure of data collection and game built together; however, the game ultimately designed in that paper uses only a single sensor and a tablet computer that must be held by the user throughout gameplay. Further, the actions, while intense, are essentially completed in place. As a result, this work is built off of the principals presented in [MNL14] in that movements are collected for a fast, intense exergame based upon realistic actions (in this case some soccer movements as well as some general obstacle course movements) with a wearable system that does not restrict the user's limbs or movement ranges, but instead allows for the full body to be involved in the action and for the user to run around in any environment the game is being played in.

This chapter introduces soccAR, an exergame derived from soccer movements, Temple Run style dashing, and a head-mounted eyewear display to allow for full range of motion and mobility. By wearing a translucent display, users are able to focus on the game while seeing their surrounding environment. This, coupled with the four wearable sensors, allows for a demonstration of a fully mobile exergame where the sensors can be embedded and the eyewear display connected to a mobile computing device like a phone or tablet. The current demonstration runs on a laptop, worn in a backpack, with communication with four accelerometer and gyroscope sensors attached to the wrists and the feet. This work focuses on the design of this game and the leave-one-subject-out cross-validation run to validate the strength of the fine-grain classifier.

10.3 System

This section describes the development of the soccAR platform. It will cover the data collection platform and game design for use. The game has two distinct

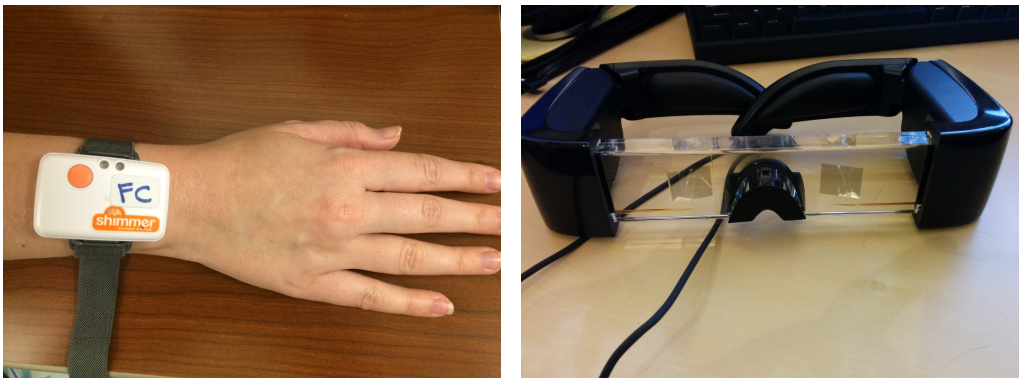

(a) Shimmer3 Sensors (b) Head-Mounted Display

Figure 10.1: Shimmer sensor (left) and Espon Moverio glasses (right)

components that work together to make SoccAR a successful mobile exergame with intense activities from a soccer environment. The first is a sensor platform to input actions from the user in any environment, the second is the computing device and display for the user to interact with the game and gaming environment, these are shown in Figure 10.1. In order to determine the movements and actions of each individual limb, four Shimmer [shi] inertial measurement units (IMU) were used. Each streams data at 50 Hz and presents three axis accelerometer ($\pm 6g$) and gyroscope ($\pm 250\frac{deg}{s}$) data as well as comes with code to estimate the orientation of the sensor, presented as four quaternions. One sensor is on each wrist, one is on the top of each shoe (to capture the best possible data for soccer-type movements like kicking a ball). The data is streamed, via bluetooth, to a computing device.

10.3.1 Head-Worn Display

Ultimately, what can make a game successful as a mobile game is the ability to display to the user the gameplay environment while leaving the hands and legs free (as opposed to, for example, using a hand held tablet). To enable this, head-worn displays are needed. However, these displays must also let the user see the actual environment, so that they do not run into actual physical objects. As a

result, the game is implemented on a laptop and displayed on the Epson Moverio BT-100[eps]. The user is presented with the game in the small displays in the middle of the lenses.

10.3.2 Data Collection

Figure 10.2 shows a user playing soccAR. SoccAR is an augmented reality Temple Run-style soccer exergame Data was collected from 16 individual volunteers (5 female, 11 male). They were asked to wear four Shimmer [shi] wireless IMU platforms, two attached to each wrist, two on top of each foot (similar to the placement of the sensor in the previous chapters). With this sensor placement a truly mobile range of motion data was collected for 26 different fine-grain movements as well as some various other random motion to help build a "no movement class" if desired. For each move, ten repetitions were recorded. Since a temple-run game is being adapted for the game in this chapter, the motions include soccer motions from the previous chapters as well as newer activities to make an enjoyable mobile exergame. Table 10.1 shows the list of moves both sports related and general running motions.

10.3.3 Game Design

Figure 10.2 shows a user playing the soccAR game and a screenshot of this game. The game is an adaptation of a Temple-Run game called Ruin Run based off of the Unity platform. It allows the game designer to use the Unity3D game engine to make a run/dash style obstacle course (similar to the idea presented in [MNL14] but with visual cues). Users can walk, run or stand still to perform certain actions. The game presents the user in a third person view, like in Figure 10.2b. This screenshot presents the user with a ball that needs to be kicked out of the way. This adventure game will require the user to run around any environment dodging,

Soccer	**Running**	**Misc.**
Shoot	Shoot (while running)	Open Door
Square Pass	Run	Unsheathe Sword
Through Pass	Jump (in place)	Draw Bow
Chip	Jump (while running)	Shoot Bow
Flick Pass	Jab (while running)	Jab
Placed Shot	Sword Slash (while running)	Sword Slash
Cut Left	Cut Left (while running)	Slide
Cut Right	Cut Right (while running)	Walk
Spin	Spin (while running)	

Table 10.1: Fine-Grain Movements for soccAR

punching, and kicking at obstacles. When presented with particular challenges, like needing to open a door or shoot an arrow, the game zooms into first person to allow the user to perform specific actions. The duration of the game can be set before the user plays, allowing for a short game for a quick break or a long game for serious exercise. These moves, among others, have been shown to generate moderate levels of intensity [MAL13]. Given how long it takes to avoid obstacles, and how many calories are burned by the user, a complete score is presented at the end of the game to give the user an indication of how well they played. The calories burned factor in to the final score in a simple linear combination. The movements and their order are randomly presented to the users to give them different experiences. For each move, a timer is started to see how long it takes the user to react to the presented object. This delay is inverted and scaled, and added to the calories burned in the following formula:

$$Score = \alpha \times \frac{1}{\delta} + \beta \times Cal \tag{10.1}$$

where δ represents the total delay in performing all actions, Cal the calories

(a) User playing soccAR (b) Game Screenshot

Figure 10.2: User playing soccAR (left), kicking a ball. Note the shimmer sensors on each wrist and each foot. Screenshot of the game (right)

burned as computed by the MET information of movements as well as the user's height and weight. α and β can be adjusted to give better scores to gameplay factors or to exercise factors as desired.

10.4 Methods

The recognition algorithm is the heart of soccAR. The algorithm for detecting the fine-grain motions listed here involves two steps. The first is a new recognition algorithm that performs stronger than previous chapters, detailed singular movements that are recognized with high confidence. The second is updating and optimizing this algorithm for creating a stronger user-centric model.

10.4.1 Training

Data collected were done so with an annotation process that marked the beginning and end of each movement. As the data was recorded, the individual in charge of running the data collection marked the beginning and end of each movement. Since there is some delay built in, the midpoint annotation is drawn from this start and end information. If the data is represented as a time series $D(t)$, where time t starts at 0 and goes until the length of D, windows are built around this as such:

$$move(midpoint, w) = (D(t - \lfloor \frac{w}{2} \rfloor), D(t + \lfloor \frac{w}{2} \rfloor)) \tag{10.2}$$

where w is the desired window size and $move(midpoint)$ is the windowed move around the desired midpoint. Before extracting these movement windows from the training set, the data is filtered with a low-pass filter via a moving average. This moving average was found, heuristically, to be one-fourth the average window size. The average window size is one second, or 100 points. The reason this length was determined, for movements that range between $1 - 2$ seconds, and a window size was selected to encompass all the movements. Once the data is filtered, the movements are supplied to a support vector machine classifier for training. The SVM used in this work is LibSVM [CL01] used in Matlab. The RBF Kernel is used with default values for the other parameters. Finally, four sensors are used at the same time, thus, each move window is the following information from each sensor:

$$SensorData_{loc}(t) = \langle G_x(t), G_y(t), G_z(t), A_x(t), A_y(t), A_z(t), \tag{10.3}$$

$$Quat_1(t), Quat_2(t), Quat_3(t), Quat_4(t), \|\mathbf{A}\| \rangle \tag{10.4}$$

where loc is the sensor identification on the body, the four $Quat$ signals are the four

Feature	Description
Minimum	Minimum value attained over window
Maximum	Maximum value attained over window
Sum	Sum of all values attained over window
Mean	Average of all values attained over window
Std. Dev	Standard Deviation of all values attained over window
Skewness	Measure of asymmetry in signal attained over window
Kurtosis	Measure of peakedness in a signal attained over window
Energy	Measure of frequency domain in a signal attained over window

Table 10.2: Features extracted for fine-grain activity classification and their descriptions

quaternions derived from the accelerometer and gyroscope readings as designed by [shi] using a standard attitude and heading reference system (AHRS), and finally, the magnitude of the acceleration is calculated. If time length t data is collected, each component of the vector above is a length t time-series signal. Thus, the full data set at any point can be indicated by:

$$MoveData(t) = \langle SensorData_{LA}, SensorData_{RA}, SensorData_{LL}, SensorData_{RL} \rangle \tag{10.5}$$

where LA is the left arm, RA the right arm, LL the left leg, and RL the right leg respectively.

10.4.1.1 Feature Extraction

As discussed in chapters 2 and 5, the recognition algorithms for general activities of daily living do poorly at recognizing fine-grain motions [MNL14]. As required, in order to make such a system work, the appropriate set of features need to be

extracted. In particular, a combination of features that generally represent the motions must be matched with features that indicate the structure. For each channel above (44 in total) a set of features is extracted. The gyroscope gives approximations of rotational movement, accelerometer the linear movements, and the quaternions on the orientation of the sensors, giving the ability to predict the position in space of each sensor. From this, the range of values possible can represent not only the direction of movement but the amplitude. The list of features that represent these are shown in Table 10.2. In particular, the energy of a signal is calculated as in [RDM05][BI04], and other similar work, where

$$fft(move) = \{x_i | i = 1, \cdots, k\} \tag{10.6}$$

$$Energy(move) = \frac{\sum_{i=1}^{k} |x_i|^2}{k} \tag{10.7}$$

the fft of a move window results in k fft components, and the energy then follows from the components. Calculating all of these features, however, will result in overfitting the data sets in training. In particular, [PJF12] suggests that, for a linear SVM, the number of data to feature ration should be $10 : 1$. However, this assumes each data point is an independent subject in a binary classification. In this case, the data set consists of 16 individuals and 26 movements. The number of samples, then would be 416 (in this case the 10 copies of each move are not considered as they inflate the data set size with polluted data, which should be avoided). As a result, no more than 42 features should be considered (and 50 are considered in this work for a multiclass classifier with a non-linear kernel). For SVM all the features selected must then be normalized over the feature space by subtracting the mean and dividing by the standard deviation.

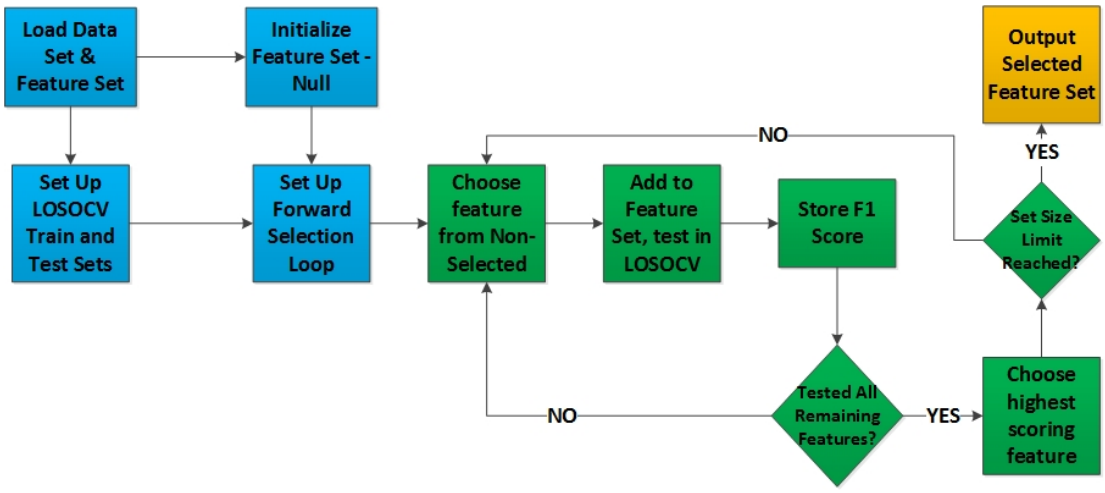

Figure 10.3: Forward selection wrapper algorithm flowchart for feature selection

10.4.1.2 Feature Selection

When all of the features from Table 10.2 were calculated over each of the channels of the data set, 352 features resulted for each sample window. In order to select these features, a selection algorithm must be run. This algorithm, however, must not consider the entire data set at once, otherwise, it will be polluted when testing the algorithm in cross validation by considering the features relations to the entire data set, not just the training set. Instead, a wrapper method will be employed to test the features. While a backward elimination approach is the best to find the optimal subset of features, the duration and complexity of such an algorithm with 352 features precludes this from being run in this instance. Instead, the forward selection approach is employed as an appropriate approximation. The flowchart for this algorithm is shown in Figure 10.3. The algorithm can use a classifier of choice, in this case, the same SVM with RBF kernel as will be used in testing the classification algorithm in order to determine the important features for the given classification scheme. Finally, the leave-one-subject-out cross-validation (LOSOCV) is used to avoid pollution of the training set for the testing set. As a result, in the classification algorithm, when a cross validation scheme is used,

this LOSOCV will be of the subset training set (in other words nesting LOSOCV valuations, feature ranking each of the subsets).

10.4.2 Classification

Once the appropriate features are selected and a model created, live testing of data is run for gameplay mechanisms. Each incoming point from the four sensors is combined together to one vector. This point is put into a sliding window of size 100 points as determined earlier, filtered with a moving average filter of 25 points. Once the features are extracted from this movement buffer, it is supplied to the classifier in order to determine the movement. The biggest issue with using such a system is to deal with all the extra points in time when no movement is happening. Most of soccAR requires walking or running. By contextually knowing when a movement needs to occur, any false positives for other movements during these periods can be ignored to improve the perceived accuracy. Further, when the movements are supposed to occur, two different approaches can be taken. The first is to create a no-movement class where all negative examples are trained. Unfortunately, to keep the training sets equally weighted, it is unlikely to capture all dissimilar movement. Another approach is to use the probability estimate outputted by LibSVM and find an appropriate threshold for requirement on a given movement classification. For soccAR, the probability estimates are used and determined heuristically for each movement type. Finally, similarly to the previous chapter, a user-centric approach can be determined by attempting to reduce the number of features extracted. In a completely powerful computing environment, extracting a large set of features can be done in parallel and no computational delay will follow. For any tablet-like device that would be ideal for such a mobile exergame, reducing the number of features computed can not only reduce delay, but potentially improve the user accuracy, similarly to adjusting the window size in the previous chapter.

10.4.3 Re-Training

Algorithm 3: Algorithm for training and using a recognition algorithm in an Exergaming environment

Data: $D(t, C)$ where $D(t, C)$ is a data point at time t with C channels,
$expMove$ is the expected move

Result: $class$, classification label for move performed (if any)

```
begin
    Data ⟵ loadDataSet(); model ⟵ trainModel(Data);
    Buffer ⟵ ∅;
    while Playing Game do
        Add D(t,C) to Buffer; FeatureSet ← extractFeatures(Buffer);
        class ← classifyMove(FeatureSet);
        if expMove ≠ class then
            missMove ← missMove ⋃ class, expMove;
            if shouldUpdate(missMove) then
                Add missMove to Data with correct labels;
                model ← updateModel(model, missMove, Data);
    saveUser(model, Data)
```

The contextual information, regarding when movements need to happen, can be used not only to ignore miss-classifications but to also improve any classifying environment. By recalling algorithm 1, shown here again from the introduction, this section concerns itself with the *shouldUpdate* method and its results. Essentially, unlike many general activity monitoring systems, in an exergaming environment, the ground truth movement can be estimated by keeping track of what move is expected versus what class label is determined by LibSVM. In order to create a more robust algorithm for any user playing the game repeatedly, the model trained can be updated with miss-classifications determined. For each movement that is miss-classified, it can be stored by the label it should have had, as deter-

mined via the context of the gameplay. If a certain number of miss-classifications is reached, this data, with the appropriate label, can be added to the algorithm and the model retrained. While most general activity monitoring systems aim at developing models that are robust for a wide population, the goal here is to eventually build a user-centric model.

10.5 Results

In order to evaluate the algorithm for the system developed, a leave-one-subject-out cross-validation is run. For each movement, the precision and recall can be used to determine the F-score, calculated as:

$$F = 2 \times \frac{P * R}{P + R} \tag{10.8}$$

where P is the precision, R the recall, and 2 a scaling factor to place the results back in a familiar $[0, 1]$ interval. The F-Score must then be weighted over the movements to created a weighted F score for each person (in this case, the same number of movements results in an equal weighting). The average F score is then calculated over all the individuals to determine the robustness of a given model. In order to test not only the accuracy but the improvements possible through the re-training, the testing set is split in half. Any Data Set is split into *Training Set* and *Testing Set*. The *Testing Set* is then split in half to create the *Test Set* and the *Increment Set*. The goal of the *Increment Set* is to combine with the trained model to bias it to the given user and test results. In order to compare the results to the original model, only the *Test Set* is used in both situations.

Features 1-10	11-20	21-30	31-40	41-50
Avg(RA,Mag)	Energy(RL,AccY)	Energy(LA,AccZ)	Std(LA,AccY)	Sum(LL,AccY)
Std(LL,GyroY)	Avg(LL,AccY)	Energy(RA,GyroZ)	Std(LA,AccX)	Std(LA,GyroZ)
Max(RA,AccY)	Energy(RL,GyroY)	Min(LL,Mag)	Energy(LL,AccX)	Min(LL,AccY)
Max(RA,Mag)	Min(LA,AccX)	Energy(LL,GyroX)	Energy(LL,AccY)	Std(LL,GyroZ)
Energy(LA,AccX)	Std(RA,GyroZ)	Skew(RL,Mag)	Skew(RA,Mag)	Std(LL,AccY)
Energy(RA,AccX)	Std(LA,GyroY)	Std(RA,Mag)	Std(LL,AccX)	Std(RL,AccY)
Energy(LL,GyroY)	Energy(LA,GyroZ)	Std(RA,AccX)	Sum(RA,Mag)	Std(RL,GyroY)
Energy(RA,Mag)	Std(LA,Mag)	Energy(RL,Mag)	Std(LL,AccZ)	Min(RA,GyroY)
Max(LA,AccY)	Std(RA,AccY)	Std(RL,Mag)	Sum(LL,Mag)	Std(RA,GyroY)
Kurt(RA,Mag)	Min(RA,Mag)	Energy(LL,GyroZ)	Max(RA,GyroZ)	Std(LL,GyroX)

Table 10.3: Top 50 selected features for fine-grain classifier by forward selection and wrapper method

10.5.1 Feature Ranking

Table 10.3 shows the top 50 features selected by the algorithm for training the classifier. Note that many of these features are structurally important for different movements, showing different limbs moving or not (for example if the right leg is stationary, such as in drawing and shooting an arrow, versus a pass where the right leg is moving). Notice the spread of features across all limbs, showing the necessity for wearing all the sensors available and combining the general windowed results with structure specific results. While work has been conducted to determine the sensor location, it is interesting to note that the quaternions appear to be of no use. As a result, perhaps only the relative adjustments of motions are required to determine the different movements.

10.5.2 Cross Validation

A leave-one-subject-out cross-validation (LOSOCV) was run on the SVM implementation presented here, as well as compared to the method from chapter 5 and the general monitoring scheme from that chapter, known as RDML[RDM05]. Note that the best f-Score comes with 25 features selected, of .72, although more

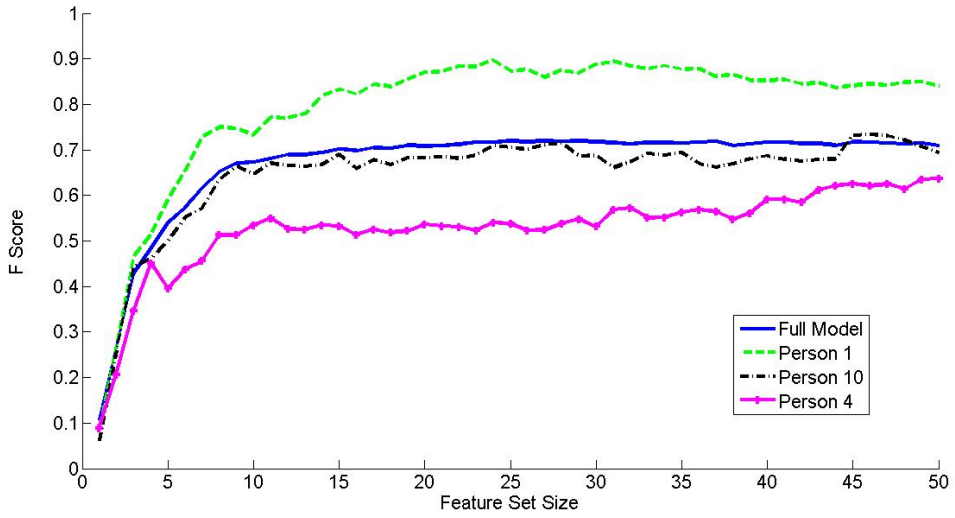

Figure 10.4: Results of the LOSOCV for entire set, and individual results for user 1, 10, and 4

features can help in certain scenarios. Consider Figure 10.4 that shows not only the full set tested and averaged, but three representative subjects individual cross validation runs. Notice, similar to the previous chapter, that an ideal number of features can be found for person 1 and person 10, although person 4 will not see any improvement in such dynamic optimization, and may, in fact, see an increase to all 50 features if accuracy is more desirable than the delay trade off.

Figure 10.5 plots the results of the LOSOCV validation when tested on the test subset, and shows the results of the re-training of the model and re-testing. Notice the significant and expected boost in results as a result of re-training the model with the user-specific data. The expected biasing/polluting of the training set performs as desired increasing the reliability of the algorithm. Table 10.4 shows the results of the best F-measure and its incremental improvement. Note we present the F-Score, a combination of the precision and recall of the system, to provide more representative accuracy information. The simple accuracy measure for the Fine-Grain PCA measure, for example, is over 98%, simply due to the

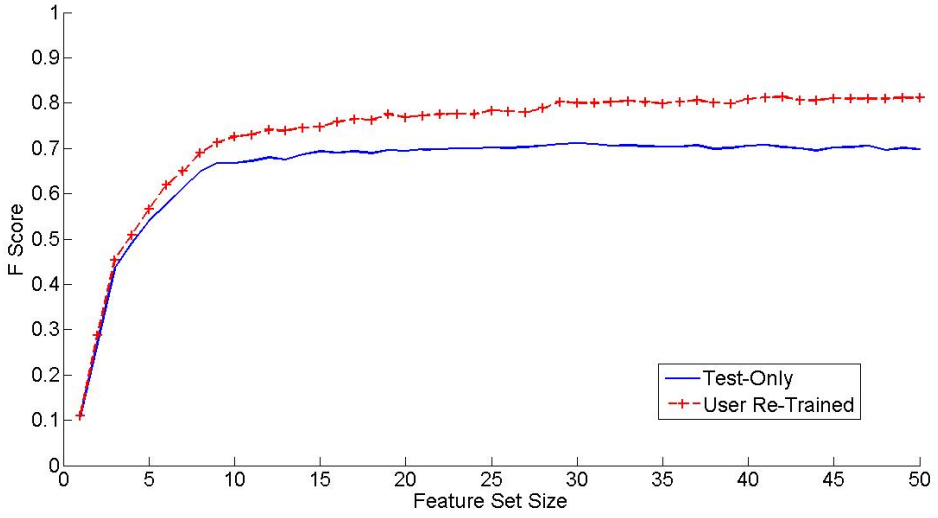

Figure 10.5: Results of LOSOCV on the test set and the re-trained LOSOCV

correct number of true negatives in such a multiclass environment. As expected, from chapter 5, the Fine-Grain PCA measure degrades again as the movement set increases, though stays fairly robust. Further, it shows that such a re-training improves multiple methods, not simply the Fine-Grain SVM presented here. With the test-subset selected, the best Fine-Grain SVM results come with 30 features. The best re-trained SVM comes from 42 features at .814, but for consistencies sake, the same number of features are presented to match the same number of eigenvectors compared in the PCA method (10, as was selected in chapter 5). More eigenvectors result in an improved result for the fine-grain PCA method, showing equally strong results close to .8 but using 35 eigenvectors, which is computationally expensive. Ultimately, while similar, the Fine-Grain SVM shows a computationally improved testing procedure that benefits more strongly from a re-training than its PCA counterparts. Indeed, even person 4, listed above, saw percentage change improvements in classification f-score from a minimum of 7.7% to a maximum of 38.3% and a max f-score of .82.

Method	F-Score	Re-Trained F-Score	Percentage Improvement
Fine-Grain SVM	.72	.81	12.5%
Fine-Grain PCA	.70	.78	11.4%
RDML	$< .05$	$< .05$	0%

Table 10.4: Results of cross validation, re-trained cross validation, and percentage improvement for Fine-Grain SVM method, chapter 5 method, and general activity of daily living method

10.6 Discussion

This chapter introduced soccAR, a fully-mobile head-mounted exergame that allows for information from multiple sensors to accurately recognize a large number of movements. Further, by using an SVM, which is relatively quick in classifying information, the improvement over previous chapters will result in a quicker classification despite using more information than the Fine-Grain PCA method would allow. What this work has shown is the ability to improve any user's information specifically by re-training any given model during the game play. Exergaming gives a unique ability to determine the correct label for miss-classifications as they occur, eventually providing a data set for re-training models and improving performance. Further, by determining a set of structure specific features to include with general features, a method to use appropriate feature extraction and classification techniques has been found to perform as desired, versus the original general activity monitoring algorithms that fail to classify this large set of fine-grain activity. What is left, then, is to compare further methods, including leveraging contextual information about the movements in the classification step, not simply in the feature and re-train steps.

CHAPTER 11

Multiple Model Recognition of Fine-Grain Motions

11.1 Desired Outcomes

This chapter develops upon the work of soccAR presented in the previous chapter. The development of user-specific models has been covered in the previous chapters, and a reasonably strong algorithm presented, based upon an SVM with a radial basis function (RBF) kernel built with good accuracy and a small number of selected features. Contextual information presented yielded possibilities for improving a classification model, introducing the idea of building user-specific exergaming models. This chapter will take this a step further, leveraging the contextual information to the classification step itself, not simply the targeted retraining. By taking knowledge of the system, multiple models can be developed to improve the accuracy of such a system greatly.

11.2 Background and Related Works

Many activity recognition problems concern themselves with similar issues, namely, setting up a model and then proving this model is robust and general, either in a testing environment or in cross-validation[HB11][QMX10][RDM05][BI04][HTW04]. However, it is possible to leverage specific information and context to improve classification results [XCC13]. Indeed, often multiple models or hierarchical classifiers

can improve classification results, from using structural information [HB11], conditional information [SKM06], or knowledge information to create ontologies to improve classification [CNW12]. These systems all attempt to relate specific information about the given environment or classification goals to improve the results and behavior of such systems, from the speed to the accuracy.

Several works leverage this information in particular format to improve their results. In [XSW11], authors define a framework for contextual information resulting in the selection of different models. It uses this contextual information to determine the labels necessary to classify and the systems in which the classification can occur. Further they define a data acquisition platform to allow for context, using a webserver for the interface between client and server information for real-time classification. This work builds upon a similar idea, using contextual information to create and select specific models for classification. Because the defined realm and contextual information is known, the system can build this context into the models, not needing an interface and data acquisition platforms. Specifically, information within the data itself is used to determine the appropriate division of work.

Work in [SWY09], a hierarchical model is built off of data driven contextual information to find certain daily living actions out of video, such as answering the phone, sitting and standing. Their work demonstrates that using contextual information from features, specifically in regards to SIFT image features in spatial and temporal relation, three different layers are created and significant classification improvement shown. This work will attempt to use similar principals with regards to time series data to improve classification results of fine-grain motions. While the features will not have spacial or temporal relations, the knowledge of movement types will be used to split the motions into hierarchical modeling situations in similar fashion.

Work in [DPB09] builds hierarchical hidden markov models in order to better

identify activities of daily living by using multiple models. Multiple models allow for their work to improve classification results of activities of daily living, particularly in an environment in which these movements might overlap. In this chapter, overlapping motions are not a concern, but the idea of using this information to reducing computation is one that is adapted here. While the work in [DPB09] does reduce computation and help set up a semi-supervised setting, the system is not set up to adapt in a real-time setting, and thus, is not applicable to the results desired in this chapter.

11.3 Methods

In order to use knowledge of movement information to build a multiple model approach to classification, the same data set and movement information from the previous section is considered. When considering the general exergaming algorithm, the $trainModel$ function is now fed contextual information instead of simply viewing the data as a single large multiclass solution. As shown in chapter 8, the ability to reduce complexity of a model by reducing the number of labels can help a classification algorithm such as an SVM. In this case the desired goal is to use feature information to group similar movements together and generate models for each of these groups. In terms of the movements, previous chapters, such as chapter 5, showed that by increasing the number of classes, the algorithm performances tend to degrade. Further, those number of classes tend to have an effect when they are most similar in terms of movements (for example, two different types of passes). Thus, the method presented will use this information to create subclusters of movements that are similar and the weka platform will be used to test the results of the LOSOCV[HFH09]. Figure 11.1 shows the general flow chart of applying expert knowledge to the clustering algorithm as desired.

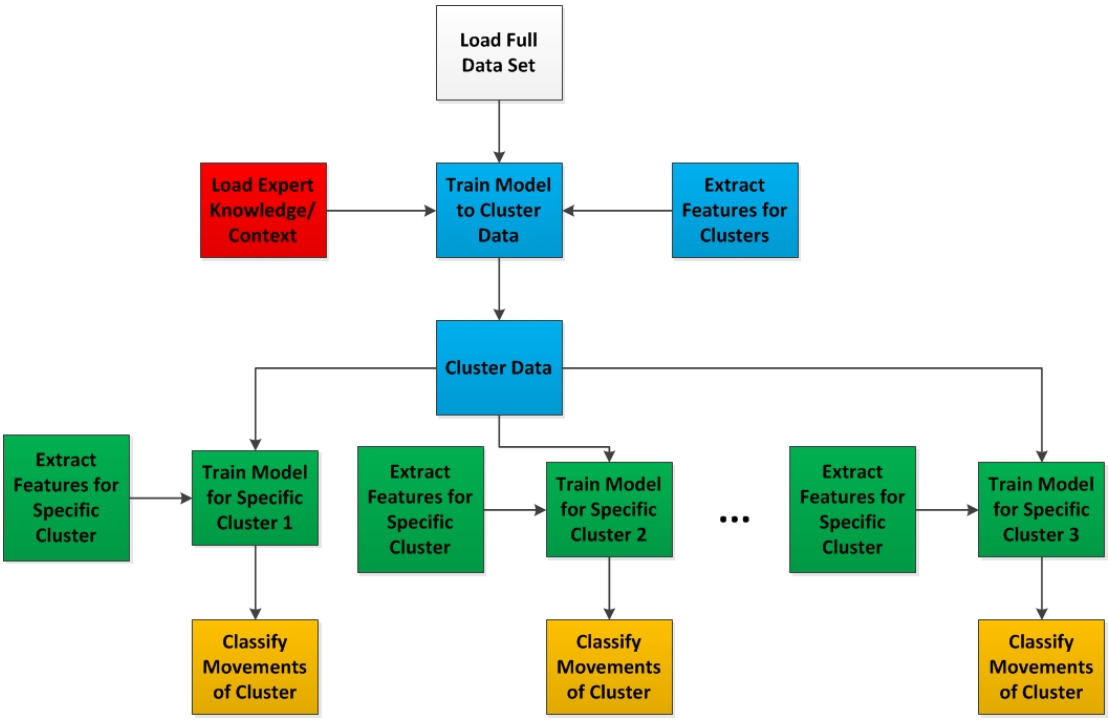

Figure 11.1: Training with expert knowledge and clustering

11.3.1 Model Generation

Table 11.1 shows the list of moves in their appropriate clusters. As soccAR contains some soccer moves and some other moves adapted for obstacles in its Temple Run-style game, contextual knowledge of the movements yielded these fairly distinct clusters of movements. Cluster 1 consists of most moves for soccer type actions. Cluster 3 results in movements that occur while currently running, including identifying running. Cluster 2 results in the remainder of movements in a miscellaneous grouping. These also happen to correspond to different expected readings from sensors. For example, most of the soccer movements are primarily foot movements, most of the miscellaneous mostly arm movements, while the other running cluster is going to be a noisy cluster of motions that consist of a fairly consistent linear velocity while performing them. Thus, four models will need to be generated. One model differentiates between the clusters, and the remaining

Cluster 1 - Soccer	**Cluster 2 - Misc.**	**Cluster 3 - Running**
Shoot	Open Door	Run
Square Pass	Draw Arrow	Cut Left (while running)
Through Pass	Shoot Arrow	Cut Right (while running)
Chip	Jab	Jab (while running)
Flick Pass	Unsheathe Sword	Jump (while running)
Placed Shot	Sword Slash	Shoot (while running)
Cut Left	Walk	Spin (while running)
Cut Right		Sword Slash (while running)
Jump		
Spin		
Slide		

Table 11.1: Fine-Grain Movements for soccAR

three models classify moves in each cluster, with significantly fewer labels than 26 in the previous chapter. This top level clustering could be done in an unsupervised method and that is left for future work. Contextual knowledge here lends itself to an easy supervised nature of classification. This work will define a new metric for measuring the accuracy of such classifiers as well. For any multiclass classification, the F score will be used, where:

$$F = 2 \times \frac{P \times R}{P + R} \tag{11.1}$$

where P is the precision and R the recall, as in previous chapters. Most multiple model systems [SKM06][DPB09][SWY09] still present classification results as an overall metric. Given a multiple model scenario this often makes the most sense for analyzing how well a given system performs. For this work, the classification accuracy of each model will be analyzed and a metric created to calculate the weighted accuracy of multiple models. The reason this is done is, in the context

of a gaming environment, as discussed in previous chapters, sometimes the performance and speed are more important than the accuracy. Take, for example, a situation in which a particular cluster contains two different passes to the left. Perhaps, in the context of a game, it is not as important to differentiate between the two passes as it is to identify that either pass has occurred. In that setting, the accuracy of the system to identify that cluster (the higher level model) is more important than the accuracy of the lower model. As such, different accuracy metrics must be calculated. For a given level and model, the accuracy, in terms of F-score, for the level is calculated as:

$$F_l = \frac{\sum_{i=1}^{n} \alpha_i * \|C_i\| * F_{l_i}}{\sum_{i=1}^{n} \alpha_i * \|C_i\|} \tag{11.2}$$

where F_l is the F-score for a certain level of multiple models (1 being the top, 2 the next and so forth), F_{l_i} is the F-score for the particular cluster/model in this level (e.g., F score for the soccer cluster in the table above), α_i is the weighted importance of this cluster in range $[0, 1]$, $\|C_i\|$ is the weighted size of the cluster. If, for example, all the α_i are 1 then the F score for the level will be the weighted average of the F score for each model in this level of the classifier. If the F_l values for each level of a multiple-model scenario are considered, then the overall F score will be calculated as:

$$F_{mm} = \frac{\sum_{i=1}^{L} \beta_i}{\sum_{i=1}^{L} \frac{\beta_i}{F_i}} \tag{11.3}$$

where the F_{mm} represents the F-score for the multiple model scheme, β_i is the importance of each model level in range $[0, 1]$ and F_i is the F-Score for the given level of the model. In other words, a weighted average is calculated for the F-score of a given level of models in this multiple model scheme, since even a single outlier should affect the result greatly, where as the weighted harmonic mean is chosen for the comparison across levels since a given level may not be crucial to the operation

of the exergaming recognition system. This work considers this top level F-Score more representative of the multiple model scheme, in particular, where all weights are 1 this scheme represents the same accuracy as an overall system. Further, by developing a single metric, it becomes easier to compare results across the models in a situation where the designer of such an exergame may not care about the distinction in a given model, such as the distinction between two types of passes, in an exergaming environment, if those movements are deemed similar enough not to effect user experience.

11.3.2 Feature Extraction

With multiple models, it becomes necessary to extract multiple features. While the features extracted for a single model might be robust enough for an entire system, multiple models should have their own features extracted and selected. This is because a multiple model solution represents, essentially, multiple independent classification problems. As a result, the same feature extraction technique presented in the previous chapter should be applied to the top level clustering scheme as well as each of the three clusters selected here. Similarly to the calculation of the F measure above for the entire selection, the number of features in use must be calculated over the entire model. To be useful in the sense of the optimization techniques of previous chapters, an overall feature number needs to be calculated. Otherwise, a sorted order of the primary reductions can be determined and used for each cluster and each level respectively. Similarly to above, the weighted feature value can be calculated, although the value of this possible computation will be left for the discussion section below.

11.3.3 Puk Kernel and Weka SMO

For modeling this problem, the Weka SMO implementation of Support Vector Machines (SVM) are implemented [HFH09], called through Matlab. For this, the complexity parameter is increased to 100 and the Puk kernel is selected, with default values for its potential robustness. In particular, the Puk kernel is a universal kernel based on the Pearson VII function. This kernel is applied here to activity monitoring, as it has been shown to be a robust, universal kernel for support vector regression [UMB06], discriminating protein types [ZG13] and on gesture recognition [DCR14]. The kernel is based on the Pearson VII function for curve fitting

$$f(x) = \frac{H}{\left[1 + \left(\frac{2(x-x_0)\sqrt{2^{\frac{1}{\omega}}-1}}{\sigma}\right)^2\right]^{\omega}} \tag{11.4}$$

where H is the peak high at the center x_0, x is the independent variable being fit, σ is the half-width of the curve, ω is the tailing factor for the peak [UMB06]. This curve fitting is particularly powerful because, by varying ω it can fit different distributions, such as Gaussian and Lorentzian [UMB06]. From this, the kernel function for an SVM can be defined as a kernel of two vectors by

$$K(x,y) = \frac{1}{\left[1 + \left(\frac{2*\sqrt{\|x-y\|}\sqrt{2^{\frac{1}{\omega}}-1}}{\sigma}\right)^2\right]^{\omega}} \tag{11.5}$$

where a variable and its midpoint of the curve is replaced by the Euclidean distance of two vectors, and the height is simply set to 1. Work in [UMB06] guarantees that this function matches the requirements for a kernel (symmetric matrix and real valued symmetrical and satisfies Mercer's conditions).

Feature	Description
Minimum	Minimum value attained over window
Maximum	Maximum value attained over window
Sum	Sum of all values attained over window
Mean	Average of all values attained over window
Std. Dev	Standard Deviation of all values attained over window
Skewness	Measure of asymmetry in signal attained over window
Kurtosis	Measure of peakedness in a signal attained over window
Energy	Measure of frequency domain in a signal attained over window

Table 11.2: Features extracted for fine-grain activity classification and their descriptions

11.4 Results

For the purposes of this experiment, the data is stored slightly differently than the previous chapter. Each sensor provides data in the following format:

$$SensorData_{loc}(t) = \langle G_x(t), G_y(t), G_z(t), A_x(t), A_y(t), A_z(t), \tag{11.6}$$

$$Quat_1(t), Quat_2(t), Quat_3(t), Quat_4(t)\rangle \tag{11.7}$$

where the magnitude of the acceleration vector is not calculated here, instead is provided at the end as in:

$$MoveData(t) = \langle SensorData_{LA}(t), SensorData_{RA}(t), SensorData_{LL}(t) \tag{11.8}$$

$$SensorData_{RL}(t), Mag_{LA}(t), Mag_{RA}(t), Mag_{LL}(t), MagRL(t)\rangle \tag{11.9}$$

where Mag is the acceleration vector for the associated sensor type.

Model Type	**Top 50 Features**
Single Model	f_5,f_6,f_{14},f_{15},f_{21},f_{22},f_{23},f_{24},f_{25},f_{26},f_{35},f_{44},f_{53},f_{55},f_{56},f_{57},f_{61} f_{62},f_{64},f_{65},f_{72},f_{74},f_{75},f_{76},f_{78},f_{79},f_{81},f_{84},f_{85},f_{86},f_{93},f_{95},f_{96} f_{101},f_{102},f_{103},f_{104},f_{105},f_{106},f_{107},f_{112},f_{115},f_{124},f_{125},f_{126},f_{130} f_{131},f_{133},f_{135},f_{136}
Top-Level Model	f_{12},f_{14},f_{45},f_{51},f_{63},f_{74},f_{78},f_{93},f_{95},f_{122},f_{124},f_{125},f_{138},f_{146},f_{157} f_{185},f_{188},f_{195},f_{202},f_{211},f_{212},f_{232},f_{235},f_{241},f_{243},f_{244},f_{246},f_{251} f_{254},f_{256},f_{258},f_{282},f_{292},f_{295},f_{313},f_{315},f_{329},f_{330},f_{331},f_{332},f_{336} f_{340},f_{348},f_{350},f_{335},f_{351},f_{312},f_{205},f_{236},f_{206}
Soccer Model	f_{15},f_{16},f_{21},f_{22},f_{31},f_{44},f_{46},f_{49},f_{53},f_{55},f_{56},f_{64},f_{65},f_{72},f_{76},f_{79},f_{85} f_{86},f_{87},f_{89},f_{95},f_{101},f_{102},f_{104},f_{105},f_{112},f_{114},f_{115},f_{124},f_{126} f_{130},f_{131},f_{132},f_{133},f_{136},f_{142},f_{154},f_{155},f_{156},f_{180},f_{181},f_{182} f_{203},f_{204},f_{223},f_{231},f_{232},f_{233},f_{256},f_{261}
Misc. Model	f_2,f_4,f_5,f_6,f_{11},f_{12},f_{15},f_{21},f_{22},f_{23},f_{24},f_{25},f_{29},f_{35},f_{42},f_{43},f_{45},f_{48} f_{57},f_{60},f_{61},f_{62},f_{65},f_{66},f_{71},f_{75},f_{81},f_{85},f_{86},f_{93},f_{95},f_{101} f_{102},f_{103},f_{104},f_{105},f_{106},f_{107},f_{124},f_{126},f_{127},f_{130},f_{133},f_{134} f_{139},f_{141},f_{142},f_{143},f_{145},f_{146}
Run Model	f_{44},f_{49},f_{63},f_{91},f_{96},f_{106},f_{131},f_{134},f_{135},f_{143},f_{145},f_{151},f_{152},f_{181},f_{183} f_{202},f_{212},f_{232},f_{301},f_{306},f_{318},f_{324},f_{342},f_{338},f_{326},f_{85},f_{22},f_{32},f_{193} f_{191},f_{112},f_{102},f_{125},f_{350},f_{263}, f_{273},f_{156},f_{126},f_{55},f_{124},f_{325},f_{346},f_{65} f_{105},f_{115},f_{123},f_{94}, f_5,f_{345},f_{341}

Table 11.3: Top 50 Selected Features for Each Model Type

11.4.0.1 Features Extracted

From this, the same set of features as the previous chapter are extracted and numbered, as in Table 11.2. Thus, feature f_1 is the minimum value of the X axis of the *Gyro* of the Left Arm. f_2 is the minimum value of the Y axis of the *Gyro*. The minimum features for the left arm are the first 10 features (3-axis gyro, 3 axis accelerometer, 4 channel quaternions). Then the next 10 features are the maximum, and so on. After all the feature types are extracted on the left arm, the right arm begins. Features $1-80$ are the left arm, $81-160$ the right arm, $161-240$ the left leg, $240-320$ the right leg, $320-328$ the magnitude of the left arm, $329-336$ the magnitude of the right arm, $335-344$ the magnitude of the left leg, and $344-352$ the magnitude of the right leg. The top 50 features for each cluster and model type are listed in Table 11.3. Notice while many features are

repeated each model has its own set of selected features. Similarly to the previous chapter, all the features are normalized over the feature space for use in an SVM classifier.

11.4.1 Cross Validation

Model											
Features− >	1	5	10	15	20	25	30	35	40	45	50
Single Model	0.129	0.557	0.671	0.697	0.714	0.720	0.722	0.723	0.719	0.718	0.707
Top Model	0.877	0.980	0.983	0.984	0.987	0.987	0.987	0.987	0.986	0.984	0.979
Soccer Model	0.239	0.561	0.639	0.681	0.690	0.686	0.682	0.688	0.688	0.690	0.673
Misc. Model	0.604	0.976	0.989	0.993	0.994	0.994	0.994	0.993	0.99	0.991	0.986
Run Model	0.305	0.587	0.648	0.672	0.686	0.698	0.7	0.691	0.685	0.689	0.675
Multiple Model	0.490	0.790	0.831	0.850	0.857	0.858	0.857	0.857	0.855	0.856	0.845

Table 11.4: F-Scores for the single multiclass classifier, each model of the multiple model method, and the overall multiple model method for RBF Kernel in SVM

A leave-one-subject-out cross-validation was run in order to test the robustness of such an algorithm. The results of the system in comparison between RBF (LibSVM in Weka [CL01]) and Puk kernels are also compared. Table 11.4 shows the results of running the multiple model scheme in the RBF Kernel. For this application, since the second level of modeling still contains different motions, the importance of each model is considered paramount. As a result, all of the α and β parameters are set to 1. From this, it is clear that the RBF model can differentiate between the three types of movements with high accuracy. While certain motion classification in each of the clusters might be lower, this method allows for high accuracy of many of the miscellaneous movements, as well as an overall performance gain over the single-model classifier from the previous chapter.

Table 11.5 shows the results of the same experiment with the Puk kernel. The results show several important factors. The first is that the single model classification scheme is significantly improved, even over the re-trained version of the previous chapter. Further, the soccer model and the run model both see sig-

Model											
Features− >	1	5	10	15	20	25	30	35	40	45	50
Single Model	0.144	0.485	0.694	0.748	0.755	0.782	0.847	0.879	0.902	0.911	0.913
Top Model	0.595	0.843	0.963	0.981	0.994	0.996	0.995	0.996	0.996	0.996	0.995
Soccer Model	0.186	0.441	0.694	0.749	0.902	0.93	0.945	0.941	0.956	0.961	0.962
Misc. Model	0.447	0.858	0.941	0.963	0.972	0.972	0.985	0.992	0.996	0.997	0.993
Run Model	0.176	0.506	0.618	0.685	0.796	0.829	0.872	0.877	0.874	0.881	0.878
Multiple Model	0.343	0.668	0.828	0.868	0.937	0.951	0.963	0.964	0.968	0.970	0.969

Table 11.5: F-Scores for the single multiclass classifier, each model of the multiple model method, and the overall multiple model method for Puk Kernel in SVM

nificant improvements in classification. The overall multiple model classification scheme not only reaches a high level of accuracy, with an F-Score reaching as high as .97, but already outperforms the best versions of previous models with only 10 features. If computational complexity is a larger concern than accuracy, then this multiple model scheme is more applicable to finding optimally fast solutions for real-time classifications as well as finding highly accurate systems. Finally, Figure 11.2 shows the individual plots for each examination type in response to the number of features used.

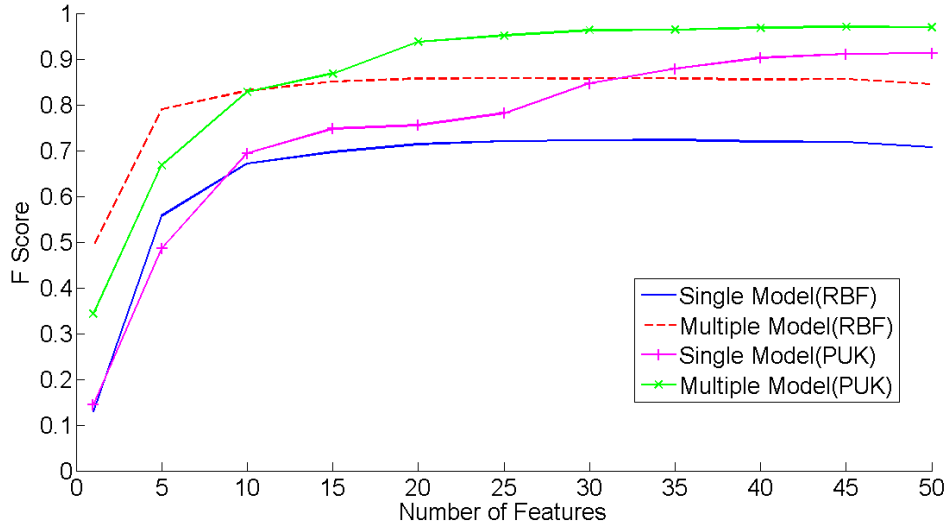

Figure 11.2: F-scores per feature of the single and multiple model SVM RBF and PUK kernels.

11.5 Discussion

This chapter covered the use of contextual, expert knowledge to differentiate between the movement types in a large, multiclass environment. Using this information, a multiple model-based supervised classifier is created. This classification scheme shows improvement over the standard single model, showing that such a large, multiclass environment needs to have such a system to be as robust and generalized. The F-score for the multiple model approach reaches as high as .97 in the defined F-measure produced in this chapter. Further, the Puk kernel is shown to be adaptable to human activity just as well as it had to other types of classification work. It is possible to extend such work to using different kinds of models at each level, be it SVMs or hidden markov models. Further, multiple levels can be further investigated, but two was chosen in this work for its applicability to soccAR in terms of recognizing movements within a desired point in time. Determining the average feature number requires further work. When comparing the real-time responsiveness of such an algorithm, it must be determined what is the best approach to determining the appropriate number of features. For example, the miscellaneous model does not need as many features selected as the soccer or run models do. However, what is considered the overall number of features used? Is it the average number of features? The maximum? The intersection? What if features are extracted in parallel while the first level model is being run? Depending upon the implementation of the real-time aspects of the work, this metric will require different representation and, as such, is left for future investigation. Finally, generalizing the contextual information leveraged to create the separate models can be performed in an unsupervised manner when memory usage is of less concern.

CHAPTER 12

Conclusion

This dissertation presented a complex end-to-end system for detailed, realistic-motion activity recognition. Such a complex system attempts to address the sedentary behavior and nature of the population and the chronic conditions and economic burdens these create. This inactivity leads to diseases and reduced life-expectancy [LSL12]. This global problem requires several different solutions working together to address clinical problems [CSL10]. Many systems exist to track the activity levels of users, particularly in movements often classified as activities of daily living, that are cyclical and repetitive [BI04]. Motivation and entertainment [MEV11] are factors that can help promote a healthier life-style. In that light, exergaming is presented. This dissertation investigates the creation of a truly mobile, active exergaming system and outlines the framework and methodology required to make a successful system. Work in the literature has shown that the successful creation of an exergame requires several factors, including interactive motions and gameplay[MVG12]. As such, it was shown, here, that the following is necessary to create a successful such game:

- Appropriate sensors for recording the appropriate motions (chapters 3,4, and 5)
- Developing a game to match the motions required (chapters 5, 9)
- Collecting the data necessary to build such a system (chapters 5, 9)

- Validating the energy expenditure levels and guaranteeing the activity levels (preventing cheating) (chapters 4, 6, 7)
- Developing an algorithm to identify these detailed, fine-grain motions (chapters 5, 8, 9, and 10)
- Improving the realistic nature of such classification schemes by creating user-centric components (chapters 8, 9, and 10)

By achieving these pieces, a working, end-to-end system named soccAR was created and its components and features demonstrated. This complex system can:

- Recognize detailed, fine-grain, realistic and specific motions
- Develop user-centric, multiple models for such recognition systems
- Analyze the trade offs between accuracy, computation complexity, and real-time responsiveness

The system in this work started with a soccer exergame due to the popularity of such games [Ele13] as an opportunity to take a global problem and solve it with an activity popular around the world. Various sensor components were compared, from using accelerometers to track activity to pressure sensors to help prevent the cheating of such activities, as well as correct the errors that accelerometer systems have. Those systems, however, cannot address the detailed tracking needs of fine-grain, realistic motions. These motions are singular, short burst motions of the human body, often in complex fashions that build upon another, as opposed to the standard cyclical, repetitive patterns monitored by standard activity monitoring algorithms. Algorithms that preserve the structure and detail of these motions when tracking them are required, and one such algorithm is used in the initial mobile exergame, presented in chapter 5. This algorithm achieved an almost 80% precision and recall, robust to the increase size of such a data set

in comparison with a well-cited activity monitoring algorithm that degrades as the multiclass problem becomes larger and larger. The initial framework for a successful exergame is laid out and the components verified, starting first with a user experience trial that indicates such a game will be successfully adopted by the users surveyed.

After the initial algorithm and system were created, improving the modeling of caloric expenditure and energy expenditure became necessary. This dissertation showed a selection of appropriate sensors to create an appropriate linear regression for modeling the energy expenditure. This was expanded upon and developed into a non-linear, advanced regression technique to show the ability to predict the metabolic equivalents within 1 MET of error. This regression was made possible through the ability to use contextual information to produce realistic monitoring of energy expenditure. However, the previous algorithm produces enough error that this system is not ideal. As such, the next chapter showed the ability to automatically detect repetition, identify the movement sizes of detailed motions, and then extract movement-specific features to create movement-specific models, even reducing the sensors used to save on battery power and other computational complexities by reducing the data necessary to classify a movement. By properly classifying a given movement, the appropriate movement model can be used to monitor the energy expenditure of that movement.

The final step of designing such an end-to-end system was to create a truly mobile, user-specific game. Leveraging the contextual information used for energy expenditure, a derivative-free optimization approach was developed to model the dynamic optimization of an exergaming system by finding the ideal point in a user-specific response to a given model. This approach was extended to the features selected and used in a given system, as well as using the context of the game to properly identify miss-classifications and re-train algorithms for a user input. By increasing the number of sensors and increasing the number of movements,

the best way of maintaining a high level of accuracy is to re-train the algorithm for each individual, improving the accuracy F-score by at least 10%. Finally, the knowledge and contextual information was applied to create advanced, multiple models to push this accuracy of fine-grain motions to an F-score as high as .97, a truly accurate method for identifying realistic, quick, fine-grain motions in a large, multiclass environment. This system requires all four limbs and accelerometers and gyroscopes to identify the motions. In order to do this, the game was developed to be displayed on head-mounted eye-wear, thus freeing up all the limbs for motion in any environment for any motion. This complex system combines all the given components into one working exergame with motions and gameplay developed together for a fun, fast, intense, interactive exergame to help promote increased physical activity.

12.1 Future Directions

12.1.1 Clinical Intervention

Such a system is thus developed and ready for future directions in research. The first such system to outline the trade off between accurate performance and real-time responsiveness, the long term adoption of such a game and its clinical outcomes must be investigated in an intervention-based clinical environment. Long term comparisons of playing such exergames versus playing standard games and the effect of body-composition need to be researched to determine a comparison to the shorter-term results presented here. Further, the modification of such games to be a more direct comparison may be necessary and, thus, required.

12.1.2 Multiple Models for Other Activity Recognition

Better modeling the contextual information and using it to develop other methods for other applications are extensions that should be investigated. Comparing the expert level of a user to determine appropriate models for each person can assist in giving more detailed information, not only in identifying the motions being performed, but being able to give instructional feedback for those motions. If extended out of gaming and into a sports environment, such a system can monitor individuals and groups, and model not only their abilities but instruct and adapt as necessary the actions prompted, identified, and classified. Further, as such data sets grow, unsupervised techniques should be investigated in identifying features for multiple models, labeling movements, perhaps in a semi-supervised state, and predicting motions that are likely to follow. As data sets grow in size, solutions that scale with size should be investigated and adapted to allow for varying sets of sensors and data applied to such motion and application systems. Finally, a comparison of further techniques and their abilities to monitor each of these components should be used in comparison to this work as a baseline result.

12.1.3 Sensors

The sensors used for such work, as well as their placement, should be allowed to be dynamic. As show in previous work, and discussed, sensor misplacement can be corrected already through calibration techniques [CXC][WWC13]. Dynamically adjusting such calculations to account for transnational movement in sensors during motion should be considered. Further, creating future non-invasive sensors embedded within things like the clothing worn by an individual should allow for various placements and orientations of such sensors. Future research should be conducted into the viability of such worn sensors, adjusting the granularity in which they need to monitor given sets of activities depending upon the situa-

tion the user is in, and adjusting the data on the fly to account for such motion monitoring.

References

[AC97] J.K. Aggarwal and Q. Cai. "Human motion analysis: A review." In *Nonrigid and Articulated Motion Workshop, 1997. Proceedings., IEEE*, pp. 90–102. IEEE, 1997.

[Act10] Activision Blizzard Inc. "Call of Duty(R): Modern Warfare(R) 2 Surpasses $1 Billion in Retail Sales Worldwide.", 2010.

[AHH11] Barbara E Ainsworth, William L Haskell, Stephen D Herrmann, Nathanael Meckes, David R Bassett, Catrine Tudor-Locke, Jennifer L Greer, Jesse Vezina, Melicia C Whitt-Glover, and Arthur S Leon. "2011 compendium of physical activities: a second update of codes and MET values." *Medicine and science in sports and exercise*, **43**(8):1575–1581, 2011.

[AHW00] B.E. Ainsworth, W.L. Haskell, M.C. Whitt, M.L. Irwin, A.M. Swartz, S.J. Strath, W.L. O'BRIEN, D.R. Bassett, K.H. Schmitz, P.O. Emplaincourt, et al. "Compendium of physical activities: an update of activity codes and MET intensities." *Medicine & Science in Sports & Exercise*, **32**(9):S498, 2000.

[AIH10] Fahd Albinali, Stephen Intille, William Haskell, and Mary Rosenberger. "Using wearable activity type detection to improve physical activity energy expenditure estimation." In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pp. 311–320. ACM, 2010.

[ALM10a] G. Alankus, A. Lazar, M. May, and C. Kelleher. "Towards customizable games for stroke rehabilitation." In *Proceedings of the 28th international conference on Human factors in computing systems*, pp. 2113–2122. ACM, 2010.

[ALM10b] Gazihan Alankus, Amanda Lazar, Matt May, and Caitlin Kelleher. "Towards customizable games for stroke rehabilitation." In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pp. 2113–2122, New York, NY, USA, 2010. ACM.

[APK11] Gazihan Alankus, Rachel Proffitt, Caitlin Kelleher, and Jack Engsberg. "Stroke therapy through motion-based games: a case study." *ACM Transactions on Accessible Computing (TACCESS)*, **4**(1):3, 2011.

[Ass13] American Diabetes Association et al. "Economic costs of diabetes in the U.S. in 2012. Diabetes Care 2013; 36: 1033–1046." *Diabetes Care*, **36**(6):1797, 2013.

[AXL13] Nabil Alshurafa, Wenyao Xu, Jason J Liu, Ming-Chun Huang, Bobak Mortazavi, Majid Sarrafzadeh, and Christian Roberts. "Robust human intensity-varying activity recognition using Stochastic Approximation in wearable sensors." In *Body Sensor Networks (BSN), 2013 IEEE International Conference on*, pp. 1–6. IEEE, 2013.

[BBS14] Andreas Bulling, Ulf Blanke, and Bernt Schiele. "A tutorial on human activity recognition using body-worn inertial sensors." *ACM Computing Surveys (CSUR)*, **46**(3):33, 2014.

[BCD02] C. BenAbdelkader, R. Cutler, and L. Davis. "Motion-based recognition of people in eigengait space." In *Automatic Face and Gesture Recognition, 2002. Proceedings. Fifth IEEE International Conference on*, pp. 267–272. IEEE, 2002.

[BCF10] S. Berkovsky, M. Coombe, J. Freyne, D. Bhandari, and N. Baghaei. "Physical activity motivating games: virtual rewards for real activity." In *Proceedings of the 28th international conference on Human factors in computing systems*, pp. 243–252. ACM, 2010.

[BI04] L. Bao and S. Intille. "Activity recognition from user-annotated acceleration data." *Pervasive Computing*, pp. 1–17, 2004.

[BI12] Regina Bernhaupt and Katherine Isbister. "Games and entertainment community SIG: shaping the future." In *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts*, pp. 1173–1176. ACM, 2012.

[BIM08] Regina Bernhaupt, Wijand Ijsselsteijn, Florian'Floyd' Mueller, Manfred Tscheligi, and Dennis Wixon. "Evaluating user experiences in games." In *CHI'08 extended abstracts on Human factors in computing systems*, pp. 3905–3908. ACM, 2008.

[BKV97] C.V.C. Bouten, K.T.M. Koekkoek, M. Verduin, R. Kodde, and J.D. Janssen. "A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity." *Biomedical Engineering, IEEE Transactions on*, **44**(3):136–147, 1997.

[BKW12] M Buman, J Kurka, E Winkler, P Gardiner, E Hekler, G Healy, N Owen, C Baldwin, and B Ainsworth. "Estimated replacement effects of accelerometer-derived physical activity and self-reported sleep duration on chronic disease biomarkers." *Journal of Science and Medicine in Sport*, **15**:S76, 2012.

[BM11] Bruce W Bailey and Kyle McInnis. "Energy cost of exergaming: a comparison of the energy cost of 6 forms of exergaming." *Archives of pediatrics & adolescent medicine*, **165**(7):597, 2011.

[BNT13] Elisabetta Bacchi, Carlo Negri, Giovanni Targher, Niccolò Faccioli, Massimo Lanza, Giacomo Zoppini, Elisabetta Zanolin, Federico Schena, Enzo Bonora, and Paolo Moghetti. "Both resistance training and aerobic training reduce hepatic fat content in type 2 diabetic subjects with nonalcoholic fatty liver disease (the RAED2 randomized trial)." *Hepatology*, **58**(4):1287–1295, 2013.

[BRL12] Frank W Booth, Christian K Roberts, and Matthew J Laye. "Lack of exercise is a major cause of chronic diseases." *Comprehensive Physiology*, 2012.

[Bro06] B. Brown. "Playing to win: Video games and the fight against obesity." *Journal of the American Dietetic Association*, **106**(2):188–189, 2006.

[BWG02] Thomas Bodenheimer, Edward H Wagner, and Kevin Grumbach. "Improving primary care for patients with chronic illness." *JAMA: the journal of the American Medical Association*, **288**(15):1909–1914, 2002.

[BWV94] Carlijn V Bouten, Klaas R Westerterp, Maarten Verduin, Jan D Janssen, et al. "Assessment of energy expenditure for physical activity using a triaxial accelerometer." *Medicine and science in sports and exercise*, **26**(12):1516–1523, 1994.

[CBC10] D. Curone, G.M. Bertolotti, A. Cristiani, E.L. Secco, and G. Magenes. "A real-time and self-calibrating algorithm based on triaxial accelerometer signals for the detection of human posture and activity." *Information Technology in Biomedicine, IEEE Transactions on*, **14**(4):1098–1105, 2010.

[CCB06a] S. E. Crouter, J. R. Churilla, and D. R. Bassett. "Estimating energy expenditure using accelerometers." *Eur. J. Appl. Physiol.*, **98**(6):601–612, Dec 2006.

[CCB06b] Scott E Crouter, Kurt G Clowers, and David R Bassett. "A novel method for using accelerometer data to predict energy expenditure." *Journal of applied physiology*, **100**(4):1324–1331, 2006.

[CCC07] Keng-Hao Chang, Mike Y Chen, and John Canny. "Tracking free-weight exercises." In *UbiComp 2007: Ubiquitous Computing*, pp. 19–37. Springer, 2007.

[CDH11] Scott E Crouter, Diane M DellaValle, Magdalene Horton, Jere D Haas, Edward A Frongillo, and David R Bassett. "Validity of the Actical for estimating free-living physical activity." *European journal of applied physiology*, **111**(7):1381–1389, 2011.

[CGS84] VICTOR A Convertino, DANIELLE J Goldwater, and HAROLD Sandler. "VO2 kinetics of constant-load exercise following bed-rest-induced deconditioning." *Journal of Applied Physiology*, **57**(5):1545–1550, 1984.

[CHB12] Scott E Crouter, Magdalene Horton, and David R Bassett Jr. "Use of a 2-Regression Model for Estimating Energy Expenditure in Children." *Medicine and science in sports and exercise*, **44**(6):1177, 2012.

[CHL06] Kuan-Ta Chen, Polly Huang, and Chin-Laung Lei. "How sensitive are online gamers to network quality?" *Communications of the ACM*, **49**(11):34–38, 2006.

[CJK13] Kyong Sik Choi, Yong Soo Joo, and Sang-Kyun Kim. "Automatic exercise counter for outdoor exercise equipment." In *Consumer Electronics (ICCE), 2013 IEEE International Conference on*, pp. 436–437. IEEE, 2013.

[CKH10] Scott E Crouter, Erin Kuffel, Jere D Haas, Edward A Frongillo, and David R Bassett Jr. "A refined 2-regression model for the ActiGraph accelerometer." *Medicine and science in sports and exercise*, **42**(5):1029, 2010.

[CL01] Chih-Chung Chang and Chih-Jen Lin. *LIBSVM: a library for support vector machines*, 2001.

[CLA07] E. Cordera, AM Lopez, D. Alvarez, RC Gonzalez, and JC Alvarez. "Physical activity classification from body center of gravity accelerations." In *Intelligent Environments, 2007. IE 07. 3rd IET International Conference on*, pp. 477–480. IET, 2007.

[CLA13] Shanshan Chen, John Lach, Oliver Amft, Marco Altini, and Julien Penders. "Unsupervised activity clustering to estimate energy expenditure with a single body sensor." In *Body Sensor Networks (BSN), 2013 IEEE International Conference on*, pp. 1–6, May 2013.

[CLL09] Pei-Chann Chang, Chen-Hao Liu, Jun-Lin Lin, Chin-Yuan Fan, and Celeste SP Ng. "A neural network with a case based dynamic window for stock trading prediction." *Expert Systems with Applications*, **36**(3):6889–6898, 2009.

[CM08] Paul T Costa and Robert R McCrae. "The Revised NEO Personality Inventory (NEO-PI-R)." *The SAGE handbook of personality theory and assessment*, **2**:179–198, 2008.

[CNW12] Liming Chen, Chris D Nugent, and Hui Wang. "A knowledge-driven approach to activity recognition in smart homes." *Knowledge and Data Engineering, IEEE Transactions on*, **24**(6):961–974, 2012.

[Con] Gulf Coast Data Concepts. "Gulf Coast Data Concepts USB Accelerometer Model X8M-3mini User Manual." `http://www.gcdataconcepts.com\/GCDC_X8M-3mini_User_Manual.pdf`. [Online; accessed 27-January-2013].

[CP01] Gert Cauwenberghs and Tomaso Poggio. "Incremental and decremental support vector machine learning." *Advances in neural information processing systems*, pp. 409–415, 2001.

[CSL10] Michele Cecchini, Franco Sassi, Jeremy A Lauer, Yong Y Lee, Veronica Guajardo-Barron, and Daniel Chisholm. "Tackling of unhealthy diets, physical inactivity, and obesity: health effects and cost-effectiveness." *The Lancet*, **376**(9754):1775–1784, 2010.

[CXC] Chieh Chien, James Yan Xu, Hua-I Chang, Xiaoxu Wu, and Greg J Pottie. "Model construction for human motion classification using inertial sensors." In *IEEE Workshop on Information Theory and Applications*.

[DAC13] Farzin Dadashi, Arash Arami, Florent Crettenand, Gregoire P. Millet, John Komar, Ludovic Seifert, and Kamiar Aminian. "A Hidden Markov Model of the breaststroke swimming temporal phases using wearable inertial measurement units." In *Body Sensor Networks (BSN), 2013 IEEE International Conference on*, pp. 1–6, May 2013.

[Dal09] A.J. Daley. "Can exergaming contribute to improving physical activity levels and health outcomes in children?" *Pediatrics*, **124**(2):763–771, 2009.

[DCR14] A Doswald, F Carrino, and F Ringeval. "Advanced Processing of sEMG Signals for User Independent Gesture Recognition." In *XIII Mediterranean Conference on Medical and Biological Engineering and Computing 2013*, pp. 758–761. Springer, 2014.

[DDD12] E. De Decker, M. De Craemer, I. De Bourdeaudhuij, K. Wijndaele, K. Duvinage, B. Koletzko, E. Grammatikaki, V. Iotova, N. Usheva, JM Fernández-Alvira, et al. "Influencing factors of screen time in preschool children: an exploration of parents' perceptions through focus groups in six European countries." *Obesity Reviews*, **13**:75–84, 2012.

[DPB09] Thi Duong, Dinh Phung, Hung Bui, and Svetha Venkatesh. "Efficient duration and hierarchical modeling for human activity recognition." *Artificial Intelligence*, **173**(7):830–856, 2009.

[Ele09] Electronic Arts Inc. "FIFA 10 Continues Record Sales Pace; FIFA 10 Ultimate Team Announced.", 2009. http://investor.ea.com/ releasedetail.cfm?ReleaseID=427414.

[Ele11a] Electronic Arts Inc. "EA SPORTS NHL 12 Shatters Records in First Week.", 2011. http://ir.ea.com/releasedetail.cfm?ReleaseID=606914.

[Ele11b] Electronic Arts Inc. "FIFA 12 Scores with Unprecedented Opening Sales and Record Online Play.", 2011. http://investor.ea.com/releasedetail.cfm?ReleaseID=611075.

[Ele13] Electronic Arts Inc. "EA SPORTS FIFA 14 Available Now in North America.", 2013. http://investor.ea.com/releasedetail.cfm?ReleaseID=792984.

[EPM08] M. Ermes, J. Parkka, J. Mantyjarvi, and I. Korhonen. "Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions." *Information Technology in Biomedicine, IEEE Transactions on*, **12**(1):20–26, 2008.

[eps] "Epson Moverio BT-100." `https://www.epson.com/cgi-bin/Store/jsp/Moverio/Home.do`. [Online; last-accessed 31-March-2014].

[FCK12] K.M. Flegal, M.D. Carroll, B.K. Kit, and C.L. Ogden. "Prevalence of Obesity and Trends in the Distribution of Body Mass Index Among US Adults, 1999-2010." *JAMA: The Journal of the American Medical Association*, **307**(5):491–497, 2012.

[FKT12] Eric A Finkelstein, Olga A Khavjou, Hope Thompson, Justin G Trogdon, Liping Pan, Bettylou Sherry, and William Dietz. "Obesity and severe obesity forecasts through 2030." *American journal of preventive medicine*, **42**(6):563–570, 2012.

[Foc] Focus. "Focus Digital Trainr.".

[FSF99] F. Foerster, M. Smeja, and J. Fahrenberg. "Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring." *Computers in Human Behavior*, **15**(5):571–583, 1999.

[GBS04] Nancy A Garrett, Michelle Brasure, Kathryn H Schmitz, Monica M Schultz, and Michael R Huber. "Physical inactivity: direct cost to a health plan." *American journal of preventive medicine*, **27**(4):304–309, 2004.

[GC09] C. Garcia Wylie and P. Coulton. "Persuasive mobile health applications." *Electronic Healthcare*, pp. 90–97, 2009.

[GGJ09] H. Ghasemzadeh, E. Guenterberg, and R. Jafari. "Energy-Efficient Information-Driven Coverage for Physical Movement Monitoring in Body Sensor Networks." *Selected Areas in Communications, IEEE Journal on*, **27**(1):58–69, January 2009.

[GJ11] H. Ghasemzadeh and R. Jafari. "Physical Movement Monitoring Using Body Sensor Networks: A Phonological Approach to Construct Spatial Decision Trees." *Industrial Informatics, IEEE Transactions on*, **7**(1):66–77, Feb 2011.

[GKG10] Andrea Grimes, Vasudhara Kantroo, and Rebecca E Grinter. "Let's play!: mobile health games for adults." In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pp. 241–250. ACM, 2010.

[GLN12] Kathrin Gerling, Ian Livingston, Lennart Nacke, and Regan Mandryk. "Full-body motion-based game interaction for older adults." In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, pp. 1873–1882. ACM, 2012.

[GOO09] L. Gorgu, G.M.P. O'Hare, and M.J. O'Grady. "Towards mobile collaborative exergaming." In *Advances in Human-oriented and Personalized Mechanisms, Technologies, and Services, 2009. CENTRIC'09. Second International Conference on*, pp. 61–64. IEEE, 2009.

[GOS08] Regina Guthold, Tomoko Ono, Kathleen L Strong, Somnath Chatterji, and Alfredo Morabia. "Worldwide variability in physical inactivity: a 51-country survey." *American journal of preventive medicine*, **34**(6):486–494, 2008.

[GPH09] D.L. Graf, L.V. Pratt, C.N. Hester, and K.R. Short. "Playing active video games increases energy expenditure in children." *Pediatrics*, **124**(2):534–540, 2009.

[Hag10] Hagop Hagopian. *"Health and Gaming: Making the Body More Active."*. Master's thesis, University of California, Los Angeles, 2010.

[HAS10] M-C Hsiao, Navid Amini, and Majid Sarrafzadeh. "A Wireless Home Automation System For Childhood Obesity Prevention." In *Proceedings of the International Conference on Biomedical Electronics and Devices*, 2010.

[HB11] Jinhui Hu and Nikolaos V Boulgouris. "Fast human activity recognition based on structure and motion." *Pattern Recognition Letters*, **32**(14):1814–1821, 2011.

[HFH09] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. "The WEKA data mining software: an update." *ACM SIGKDD explorations newsletter*, **11**(1):10–18, 2009.

[HLJ08] Z. He, Z. Liu, L. Jin, L.X. Zhen, and J.C. Huang. "Weightlessness feature a novel feature for single tri-axial accelerometer based activity recognition." In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pp. 1–4. IEEE, 2008.

[HMB00] D. Hendelman, K. Miller, C. Baggett, E. Debold, and P. Freedson. "Validity of accelerometry for the assessment of moderate intensity physical activity in the field." *Med Sci Sports Exerc*, **32**(9 Suppl):S442–449, Sep 2000.

[HPT12] S. Hilton, C. Patterson, and A. Teyhan. "Escalating coverage of obesity in UK newspapers: The evolution and framing of the obesity epidemicfrom 1996 to 2010." *Obesity*, 2012.

[HTK11] Paul A Heidenreich, Justin G Trogdon, Olga A Khavjou, Javed Butler, Kathleen Dracup, Michael D Ezekowitz, Eric Andrew Finkelstein, Yuling Hong, S Claiborne Johnston, Amit Khera, et al. "Forecasting the future of cardiovascular disease in the United States a policy statement from the American heart association." *Circulation*, **123**(8):933–944, 2011.

[HTM11] A. Henpraserttae, S. Thiemjarus, and S. Marukatat. "Accurate Activity Recognition using a Mobile Phone regardless of Device Orientation and Location." In *Body Sensor Networks (BSN), 2011 International Conference on*, pp. 41–46. IEEE, 2011.

[HTW04] W. Hu, T. Tan, L. Wang, and S. Maybank. "A survey on visual surveillance of object motion and behaviors." *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, **34**(3):334–352, 2004.

[KCK11] Oliver Kramer, David Echeverría Ciaurri, and Slawomir Koziel. "Derivative-free optimization." In *Computational Optimization, Methods and Algorithms*, pp. 61–83. Springer, 2011.

[KLH10] S.L. Kozey, K. Lyden, C.A. Howe, J.W. Staudenmayer, and P.S. Freedson. "Accelerometer output and MET values of common physical activities." *Medicine and science in sports and exercise*, **42**(9):1776, 2010.

[KLM13] Haik Kalantarian, Sunghoon Ivan Lee, Anurag Mishra, Hassan Ghasemzadeh, and Majid Sarrafzadeh. "Multimodal Energy Expenditure Calculation for Pervasive Health: A Data Fusion Model using

Wearable Sensors." In *Proceedings of IEEE PerCom Workshop on Smart Environments and Ambient Intelligence*. ACM, 2013.

[KMH13] Matthias Kranz, Andreas Möller, Nils Hammerla, Stefan Diewald, Thomas Plötz, Patrick Olivier, and Luis Roalter. "The mobile fitness coach: Towards individualized skill assessment using personalized mobile devices." *Pervasive and Mobile Computing*, **9**(2):203–215, 2013.

[KMK11] A. Koivisto, S. Merilampi, and K. Kiili. "Mobile exergames for preventing diseases related to childhood obesity." In *Proceedings of the 4th International Symposium on Applied Sciences in Biomedical and Communication Technologies*, p. 29. ACM, 2011.

[KNM06] D.M. Karantonis, M.R. Narayanan, M. Mathie, N.H. Lovell, and B.G. Celler. "Implementation of a real-time human movement classifier using a triaxial accelerometer for ambulatory monitoring." *Information Technology in Biomedicine, IEEE Transactions on*, **10**(1):156–167, 2006.

[KPT10] K. Kiili, A. Perttula, P. Tuomi, M. Suominen, A. Lindstedt, and I.A. Sánchez. "Designing mobile multiplayer exergames for physical education." In *Proceedings of IADIS International Conference on Mobile Learning*, pp. 163–169, 2010.

[LFH13] Chuanjiang Li, Minrui Fei, Huosheng Hu, and Ziming Qi. "Free Weight Exercises Recognition Based on Dynamic Time Warping of Acceleration Data." In *Intelligent Computing for Sustainable Energy and Environment*, pp. 178–185. Springer, 2013.

[LGF12] Shaopeng Liu, Robert Gao, and Patty Freedson. "Computational Methods for Estimating Energy Expenditure in Human Physical Activities." *Medicine and science in sports and exercise*, 2012.

[LGJ12] Shaopeng Liu, Robert X Gao, Dinesh John, John W Staudenmayer, and Patty S Freedson. "Multisensor Data Fusion for Physical Activity Assessment." *Biomedical Engineering, IEEE Transactions on*, **59**(3):687–696, 2012.

[LKS11] Kate Lyden, Sarah L Kozey, John W Staudenmeyer, and Patty S Freedson. "A comprehensive evaluation of commonly used accelerometer energy expenditure and MET prediction equations." *European journal of applied physiology*, **111**(2):187–201, 2011.

[LL13] Oscar D Lara and Miguel A Labrador. "A survey on human activity recognition using wearable sensors." *Communications Surveys & Tutorials, IEEE*, **15**(3):1192–1209, 2013.

[LSL12] I-Min Lee, Eric J Shiroma, Felipe Lobelo, Pekka Puska, Steven N Blair, and Peter T Katzmarzyk. "Effect of physical inactivity on major non-communicable diseases worldwide: an analysis of burden of disease and life expectancy." *The Lancet*, **380**(9838):219–229, 2012.

[LTB12] A.S. Lu, D. Thompson, J. Baranowski, R. Buday, and T. Baranowski. "Story Immersion in a Health Video Game for Childhood Obesity Prevention." *GAMES FOR HEALTH: Research, Development, and Clinical Applications*, 2012.

[LTK12] Elizabeth J Lyons, Deborah F Tate, Stephanie E Komoski, Philip M Carr, and Dianne S Ward. "Novel approaches to obesity prevention: effects of game enjoyment and game type on energy expenditure in active video games." *Journal of diabetes science and technology*, **6**(4):839–848, 2012.

[Mac12] Andrew P Macvean. "Developing adaptive exergames for adolescent children." In *Proceedings of the 11th International Conference on Interaction Design and Children*, pp. 339–342. ACM, 2012.

[MAL13] Bobak Mortazavi, Nabil Alsharufa, Sunghoon Ivan Lee, Mars Lan, Majid Sarrafzadeh, Michael Chronley, and Christian K Roberts. "MET calculations from on-body accelerometers for exergaming movements." In *Body Sensor Networks (BSN), 2013 IEEE International Conference on*, pp. 1–6. IEEE, 2013.

[MB11] J. Marshall and S. Benford. "Using fast interaction to create intense experiences." In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pp. 1255–1264. ACM, 2011.

[MCK08] Ann E. Maloney, T. Carter Bethea, Kristine S. Kelsey, Julie T. Marks, Sadye Paez, Angela M. Rosenberg, Diane J. Catellier, Robert M. Hamer, and Linmarie Sikich. "A Pilot of a Video Game (DDR) to Promote Physical Activity and Decrease Sedentary Screen Time." *Obesity*, **16**(9):2074–2080, Jul 2008.

[MCL12] Bobak Mortazavi, Kin Chung Chu, Xialong Li, Jessica Tai, Shwetha Kotekar, and Majid Sarrafzadeh. "Near-Realistic Motion Video Games with Enforced Activity." In *Wearable and Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*, pp. 28–33. IEEE, 2012.

[Mem] Memsense. "Memsense Wireless IMU Product Specification User Guide." `http://memsense.com/docs/wIMU/%scPSD-0821_Rev-J_Wireless_IMU_Product_Specification_User_Guide.pdf`. [Online; accessed 13-September-2012].

[MEV11] F.F. Mueller, D. Edge, F. Vetere, M.R. Gibbs, S. Agamanolis, B. Bongers, and J.G. Sheridan. "Designing sports: a framework for exertion games." In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pp. 2651–2660. ACM, 2011.

[MFM11] Ralph Maddison, Louise Foley, Cliona Ni Mhurchu, Yannan Jiang, Andrew Jull, Harry Prapavessis, Maea Hohepa, and Anthony Rodgers. "Effects of active video games on body composition: a randomized controlled trial." *The American journal of clinical nutrition*, **94**(1):156–163, 2011.

[MHK06] T.B. Moeslund, A. Hilton, and V. Krüger. "A survey of advances in vision-based human motion capture and analysis." *Computer vision and image understanding*, **104**(2):90–126, 2006.

[MHS01] J. Mantyjarvi, J. Himberg, and T. Seppanen. "Recognizing human motion with multiple acceleration sensors." In *Systems, Man, and Cybernetics, 2001 IEEE International Conference on*, volume 2, pp. 747–752. Ieee, 2001.

[MHW11] Bobak Mortazavi, Hagop Hagopian, Jonathan Woodbridge, Behrooz Yadegar, and Majid Sarrafzadeh. "A Wireless Body-wearable Sensor System for Designing Physically Interactive Video Games." In *BIODEVICES*, pp. 62–69, 2011.

[MNL14] B. Mortazavi, S. Nyamathi, S.I. Lee, T. Wilkerson, H. Ghasemzadeh, and M. Sarrafzadeh. "Near-Realistic Mobile Exergames With Wireless Wearable Sensors." *Biomedical and Health Informatics, IEEE Journal of*, **18**(2):449–456, March 2014.

[MPA14] Bobak Mortazavi, Mohammad Pourhomayoun, Gabriel Alsheikh, Nabil Alshurafa, Sunghoon Ivan Lee, and Majid Sarrafzadeh. "Determining the Single Best Axis for Exercise Repetition Recognition and Counting with SmartWatches." In *2014 Eleventh IEEE Body Sensor Networks Conference(BSN)*. IEEE, 2014.

[MR12] Andrew Macvean and Judy Robertson. "iFitQuest: a school based study of a mobile location-aware exergame for adolescents." In *Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services*, pp. 359–368. ACM, 2012.

[MRD12] A Moller, Luis Roalter, Stefan Diewald, Johannes Scherr, Matthias Kranz, Nils Hammerla, Patrick Olivier, and T Plotz. "Gymskill: A personal trainer for physical exercises." In *Pervasive Computing and Communications (PerCom), 2012 IEEE International Conference on*, pp. 213–220. IEEE, 2012.

[MTM12] Nicole M Mullins, Kathryn A Tessmer, Michele L McCarroll, and Brian P Peppel. "Physiological and perceptual responses to Nintendo® Wii Fit in young and older adults." *International Journal of Exercise Science*, **5**(1):9, 2012.

[MVG12] Florian Mueller, Frank Vetere, Martin Gibbs, Darren Edge, Stefan Agamanolis, Jennifer Sheridan, and Jeffrey Heer. "Balancing exertion experiences." In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, pp. 1853–1862. ACM, 2012.

[MW09] Jorge J Moré and Stefan M Wild. "Benchmarking derivative-free optimization algorithms." *SIAM Journal on Optimization*, **20**(1):172–191, 2009.

[MYO10] M. Miyachi, K. Yamamoto, K. Ohkawara, and S. Tanaka. "METs in adults while playing active video games: a metabolic chamber study." *Medicine & Science in Sports & Exercise*, **42**(6):1149, 2010.

[Nac09] Lennart Nacke. "From playability to a hierarchical game usability model." In *Proceedings of the 2009 Conference on Future Play on@ GDC Canada*, pp. 11–12. ACM, 2009.

[NAP03] B. Najafi, K. Aminian, A. Paraschiv-Ionescu, F. Loew, C.J. Bula, and P. Robert. "Ambulatory system for human motion analysis using a kinematic sensor: monitoring of daily physical activity in the elderly." *Biomedical Engineering, IEEE Transactions on*, **50**(6):711–723, 2003.

[OCK12] C.L. Ogden, M.D. Carroll, B.K. Kit, and K.M. Flegal. "Prevalence of Obesity and Trends in Body Mass Index Among US Children and Adolescents, 1999-2010." *JAMA: The Journal of the American Medical Association*, **307**(5):483–490, 2012.

[OF06] N. Oliver and F. Flores-Mangas. "MPTrain: a mobile, music and physiology-based personal trainer." In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, pp. 21–28. ACM, 2006.

[Pap09] M. Papastergiou. "Exploring the potential of computer and video games for health and physical education: A literature review." *Computers & Education*, **53**(3):603–622, 2009.

[PBW13] G Plasqui, AG Bonomi, and KR Westerterp. "Daily physical activity assessment with accelerometers: new insights and validation studies." *Obesity Reviews*, 2013.

[PEA07] Juha Parkka, Miikka Ermes, Kari Antila, Mark van Gils, Ari Manttari, and Heikki Nieminen. "Estimating intensity of physical activity: a

comparison of wearable accelerometer and gyro sensors and 3 sensor locations." In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pp. 1511–1514. IEEE, 2007.

[PEK06] J. Parkka, M. Ermes, P. Korpipaa, J. Mantyjarvi, J. Peltola, and I. Korhonen. "Activity classification using realistic data from wearable sensors." *Information Technology in Biomedicine, IEEE Transactions on*, **10**(1):119–128, 2006.

[PFB00] Michael L Pollock, Barry A Franklin, Gary J Balady, Bernard L Chaitman, Jerome L Fleg, Barbara Fletcher, Marian Limacher, Ileana L Piña, Richard A Stein, Mark Williams, et al. "Resistance exercise in individuals with and without cardiovascular disease benefits, rationale, safety, and prescription an advisory from the committee on exercise, rehabilitation, and prevention, council on clinical cardiology, American Heart Association." *Circulation*, **101**(7):828–833, 2000.

[PHK13] Igor Pernek, Karin Anna Hummel, and Peter Kokol. "Exercise repetition detection for resistance training based on smartphones." *Personal and ubiquitous computing*, **17**(4):771–782, 2013.

[PHL12] Taiwoo Park, Inseok Hwang, Uichin Lee, Sunghoon Ivan Lee, Chungkuk Yoo, Youngki Lee, Hyukjae Jang, Sungwon Peter Choe, Souneil Park, and Junehwa Song. "ExerLink: enabling pervasive social exergames with heterogeneous exercise devices." In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pp. 15–28. ACM, 2012.

[PJF12] Leslie S Prichep, Arnaud Jacquin, Julie Filipenko, Samanwoy Ghosh Dastidar, Stephen Zabele, Asmir Vodencarevic, and Neil S Rothman. "Classification of traumatic brain injury severity using informed data reduction in a series of binary classifier algorithms." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, **20**(6):806–822, 2012.

[PLC11] Wei Peng, Jih-Hsuan Lin, and Julia Crouse. "Is playing exergames really exercising? A meta-analysis of energy expenditure in active video games." *Cyberpsychology, Behavior, and Social Networking*, **14**(11):681–688, 2011.

[PLL09] L. Prévost, O. Liechti, and M.J. Lyons. "Design and Implementation of a Mobile Exergaming Platform." *Intelligent Technologies for Interactive Entertainment*, pp. 213–220, 2009.

[PMJ08] C.A. Paw, JM Marijke, W.M. Jacobs, E.P.G. Vaessen, S. Titze, and W. van Mechelen. "The motivation of children to play an active video game." *Journal of Science and Medicine in Sport*, **11**(2):163–166, 2008.

[PSM07] J. Pansiot, D. Stoyanov, D. McIlwraith, B. Lo, and G. Yang. "Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems." In *4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, pp. 208–212. Springer, 2007.

[PTS05] Fabio Pitta, Thierry Troosters, Martijn A Spruit, Vanessa S Probst, Marc Decramer, and Rik Gosselink. "Characteristics of physical activities in daily life in chronic obstructive pulmonary disease." *American journal of respiratory and critical care medicine*, **171**(9):972–977, 2005.

[QMX10] H. Qian, Y. Mao, W. Xiang, and Z. Wang. "Recognition of human activities using SVM multi-class classifier." *Pattern Recognition Letters*, **31**(2):100–111, 2010.

[RAG11] Aiying Rong, Renzo Akkerman, and Martin Grunow. "An optimization approach for managing fresh food quality throughout the supply chain." *International Journal of Production Economics*, **131**(1):421–429, 2011.

[RAO08] K. Ridley, B. E. Ainsworth, and T. S. Olds. "Development of a compendium of energy expenditures for youth." *Int J Behav Nutr Phys Act*, **5**:45, 2008.

[RDM05] N. Ravi, N. Dandekar, P. Mysore, and M.L. Littman. "Activity recognition from accelerometer data." In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, p. 1541. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.

[RKH11] M. Raptis, D. Kirovski, and H. Hoppe. "Real-time classification of dance gestures from skeleton animation." In *Proceedings of the 10th Annual ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA 2011*, pp. 147–156, 2011.

[Rob99] Thomas N. Robinson. "Reducing Children's Television Viewing to Prevent Obesity: A Randomized Controlled Trial." *JAMA*, **282**(16):1561–1567, 1999.

[RVB08] J.P. Rey-López, G. Vicente-Rodríguez, M. Biosca, and L.A. Moreno. "Sedentary behaviour and obesity development in children and

adolescents." *Nutrition, Metabolism and Cardiovascular Diseases*, **18**(3):242–251, 2008.

[RWB05] Brent W Roberts, Kate E Walton, and Tim Bogg. "Conscientiousness and health across the life course." *Review of General Psychology*, **9**(2):156, 2005.

[Sam] Samsung. "Samsung Galaxy Gear.".

[SBV11] Christian Seeger, Alejandro Buchmann, and Kristof Van Laerhoven. "myHealthAssistant: a phone-based body sensor network that captures the wearer's exercises throughout the day." In *Proceedings of the 6th International Conference on Body Area Networks*, pp. 1–7. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011.

[SC11] Amanda E Staiano and Sandra L Calvert. "The promise of exergames as tools to measure physical health." *Entertainment Computing*, **2**(1):17–21, 2011.

[SFC11] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake. "Real-time human pose recognition in parts from single depth images." In *CVPR*, volume 2, p. 7, 2011.

[SGB09a] T. Stach, TC Graham, M. Brehmer, and A. Hollatz. "Classifying input for active games." In *Proceedings of the International Conference on Advances in Computer Entertainment Technology*, pp. 379–382. ACM, 2009.

[SGB09b] Tadeusz Stach, T. C. Nicholas Graham, Matthew Brehmer, and Andreas Hollatz. "Classifying input for active games." In *ACE '09: Proceedings of the International Conference on Advances in Computer Entertainment Technology*, pp. 379–382, New York, NY, USA, 2009. ACM.

[SHD09] S.R. Siegel, B.L. Haddock, A.M. Dubois, and L.D. Wilkin. "Active video/arcade games (exergaming) and energy expenditure in college students." *International journal of exercise science*, **2**(3):165, 2009.

[shi] "Shimmer3 Wireless Sensor Platform." `http://www.shimmersensing.com/images/uploads/docs/Shimmer3_Spec_Sheet_-_Jan_2014.pdf`. [Online; last-accessed 31-March-2014].

[SKM06] Cristian Sminchisescu, Atul Kanaujia, and Dimitris Metaxas. "Conditional models for contextual human motion recognition." *Computer Vision and Image Understanding*, **104**(2):210–220, 2006.

[SKV07] H.M. Schepers, H. Koopman, and P.H. Veltink. "Ambulatory assessment of ankle and foot dynamics." *Biomedical Engineering, IEEE Transactions on*, **54**(5):895–902, 2007.

[SMS05] A.M. Sabatini, C. Martelloni, S. Scapellato, and F. Cavallo. "Assessment of walking features from foot inertial sensing." *Biomedical Engineering, IEEE Transactions on*, **52**(3):486–494, 2005.

[Son] Sony. "Sony Smartwatch 2.".

[SSS04] N. Stettler, T.M. Signer, and P.M. Suter. "Electronic Games and Environmental Factors Associated with Childhood Obesity in Switzerland* &ast." *Obesity*, **12**(6):896–903, 2004.

[SVS08] Maja Stikic, Kristof Van Laerhoven, and Bernt Schiele. "Exploring semi-supervised and active learning for activity recognition." In *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, pp. 81–88. IEEE, 2008.

[SVS12] José Luis González Sánchez, Francisco Luis Gutiérrez Vela, Francisco Montero Simarro, and Natalia Padilla-Zea. "Playability: analysing user experience in video games." *Behaviour & Information Technology*, **31**(10):1033–1054, 2012.

[SWC05] SW Su, L Wang, BG Celler, E Ambikairajah, and AV Savkin. "Estimation of walking energy expenditure by using support vector regression." In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pp. 3526–3529. IEEE, 2005.

[SWY09] Ju Sun, Xiao Wu, Shuicheng Yan, Loong-Fah Cheong, T-S Chua, and Jintao Li. "Hierarchical spatio-temporal context modeling for action recognition." In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 2004–2011. IEEE, 2009.

[SZC12] JOHN Staudenmayer, WEIMO Zhu, and DIANE J Catellier. "Statistical considerations in the analysis of accelerometry-based activity monitor data." *Medicine and science in sports and exercise*, **44**(1 Suppl 1):S61, 2012.

[TBB] Debbe Thompson, Tom Baranowski, Richard Buday, Janice Baranowski, Victoria Thompson, Russell Jago, and Melissa Juliano Griffith. "Serious Video Games for Health: How Behavioral Science Guided the Design of a Game on Diabetes and Obesity.".

[TWG11] A.G. Tsai, D.F. Williamson, and H.A. Glick. "Direct medical cost of overweight and obesity in the USA: a quantitative systematic review." *Obesity Reviews*, **12**(1):50–61, 2011.

[UMB06] Bülent Üstün, Willem J Melssen, and Lutgarde MC Buydens. "Facilitating the application of support vector regression by using a universal Pearson VII function based kernel." *Chemometrics and Intelligent Laboratory Systems*, **81**(1):29–40, 2006.

[VAS11] A. Vahdatpour, N. Amini, and M. Sarrafzadeh. "On-body device localization for health and medical monitoring applications." In *Pervasive Computing and Communications (PerCom), 2011 IEEE International Conference on*, pp. 37–44. IEEE, 2011.

[VBG13] Eduardo Velloso, Andreas Bulling, Hans Gellersen, Wallace Ugulino, and Hugo Fuks. "Qualitative activity recognition of weight lifting exercises." In *Proceedings of the 4th Augmented Human International Conference*, pp. 116–123. ACM, 2013.

[VBV96] P.H. Veltink, H.B. Bussmann, W. de Vries, W.L. Martens, R.C. Van Lummel, et al. "Detection of static and dynamic activities using uniaxial accelerometers." *Rehabilitation Engineering, IEEE Transactions on*, **4**(4):375–385, 1996.

[VLR10] Arnold POS Vermeeren, Effie Lai-Chong Law, Virpi Roto, Marianna Obrist, Jettie Hoonhout, and Kaisa Väänänen-Vainio-Mattila. "User experience evaluation methods: current state and development needs." In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*, pp. 521–530. ACM, 2010.

[WA00] SY Wu and Green A.P. "Projection of chronic illness prevalence and cost inflation.", 2000.

[WA10] D Withrow and DA Alter. "The economic burden of obesity worldwide: a systematic review of the direct costs of obesity." *Obesity Reviews*, **12**(2):131–141, 2010.

[WAH10] Wanmin Wu, Ahsan Arefin, Zixia Huang, Pooja Agarwal, Shu Shi, Raoul Rivas, and Klara Nahrstedt. "Im the Jedi!-A case study of user experience in 3D tele-immersive gaming." *Proc. IEEE ISM*, 2010.

[WBL08a] Y. Wang, M.A. Beydoun, L. Liang, B. Caballero, and S.K. Kumanyika. "Will all Americans become overweight or obese? Estimating the progression and cost of the US obesity epidemic." *Obesity*, **16**(10):2323–2330, 2008.

[WBL08b] Youfa Wang, May A. Beydoun, Lan Liang, Benjamin Caballero, and Shiriki K. Kumanyika. "Will All Americans Become Overweight or Obese? Estimating the Progression and Cost of the US Obesity Epidemic." *Obesity*, **16**(10):2323–2330, 2008.

[WC08] C.G. Wylie and P. Coulton. "Mobile exergaming." In *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*, pp. 338–341. ACM, 2008.

[WCF07] Anthony Whitehead, Nick Crampton, Kaitlyn Fox, and Hannah Johnston. "Sensor networks as video game input devices." In *Future Play '07: Proceedings of the 2007 conference on Future Play*, pp. 38–45, New York, NY, USA, 2007. ACM.

[WJN10] Anthony Whitehead, Hannah Johnston, Nicole Nixon, and Jo Welch. "Exergame effectiveness: what the numbers can tell us." In *Proceedings of the 5th ACM SIGGRAPH Symposium on Video Games*, pp. 55–62. ACM, 2010.

[WWC13] Xiaoxu Wu, Yan Wang, Chieh Chien, and Greg Pottie. "Self-calibration of sensor misplacement based on motion signatures." In *Body Sensor Networks (BSN), 2013 IEEE International Conference on*, pp. 1–5. IEEE, 2013.

[XCC13] J Xu, H Chang, C Chen, W Kaiser, and G Pottie. "Context-driven, Prescription based Personal Activity Classification: Methodology, Architecture and End-to-End Implementation." 2013.

[XPK13] James Y Xu, Gregory J Pottie, and William J Kaiser. "Enabling Large-Scale Ground-Truth Acquisition and System Evaluation in Wireless Health." *Biomedical Engineering, IEEE Transactions on*, **60**(1):174–178, 2013.

[XSW11] J.Y. Xu, Y. Sun, Z. Wang, W.J. Kaiser, and G.J. Pottie. "Context guided and personalized activity classification system." In *Proceedings of the 2nd Conference on Wireless Health*, p. 12. ACM, 2011.

[XZS12] Wenyao Xu, Mi Zhang, Alexander A Sawchuk, and Majid Sarrafzadeh. "Co-recognition of human activity and sensor location via compressed sensing in wearable body sensor networks." In *Wearable and Implantable Body Sensor Networks (BSN), 2012 Ninth International Conference on*, pp. 124–129. IEEE, 2012.

[YH10] Che-Chang Yang and Yeh-Liang Hsu. "A Review of Accelerometry-Based Wearable Motion Detectors for Physical Activity Monitoring." *Sensors*, **10**(8):7772–7788, 2010.

[YYC12] X. Yi, C. Yin, M. Chang, and Y. Xiao. "Prevalence and risk factors of obesity among school-aged children in Xian, China." *European journal of pediatrics*, pp. 1–6, 2012.

[ZG13] Guangya Zhang and Huihua Ge. "Support vector machine with a Pearson VII function kernel for discriminating halophilic and non-halophilic proteins." *Computational biology and chemistry*, **46**:16–22, 2013.

[ZSF07] P. Zappi, T. Stiefmeier, E. Farella, D. Roggen, L. Benini, and G. Troster. "Activity recognition from on-body sensors by classifier fusion: sensor scalability and robustness." In *Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on*, pp. 281–286. Ieee, 2007.