Recognizing Movements from the Ground Reaction Force

Robert Headon
Laboratory for Communications Engineering
Cambridge University Engineering Department
rph25@eng.cam.ac.uk

Rupert Curwen
AT&T Laboratories
Cambridge, UK
rcurwen@uk.research.att.com

ABSTRACT

This paper presents a novel approach to movement recognition, using the vertical component of a person's Ground Reaction Force (GRF). Typical primitive movements such as taking a step, jumping, drop-landing, sitting down, rising to stand and crouching are decomposed and recognized in terms of the GRF signal observed by a weight sensitive floor. Previous works focused on vision processing for movement recognition. This work provides a new sensor modality for a larger research effort, that of sentient computing, which is concerned with giving computers awareness of their environment and inhabitants.

1. INTRODUCTION

This paper describes the implementation of a classification system to recognize everyday human movements. The approach is novel in that it uses the forces experienced by a floor to classify movements, as opposed to a number of existing vision systems.

The output of the classification system provides a new sensor modality for a larger research effort, that of *sentient computing* [1], which is concerned with giving computers awareness of their environment and its inhabitants. Through awareness comes the ability of developing new techniques for interacting away from the desk. For example, it is possible to tie the delivery of information to physical actions anywhere in the environment.

Movement is generated by a biomechanical process that exerts forces on the environment. These forces can be observed using a suitable sensing system and used to classify movement patterns. This research uses pressure sensors underneath a floor to observe the vertical component of the ground reaction force (GRF) which is used to classify movements. Pattern recognition can be described as the best possible way of utilizing available sensors, processors and domain knowledge to make decisions automatically [2]. The four best known approaches are: 1) template matching; 2) statistical classification; 3) syntactic or structured matching; 4) neural networks. The *statistical approach* represents a pattern by d-features, viewed as a point in d-dimensional space. Training patterns are used to determine the decision boundaries, which are specified by the probability distributions of the patterns belonging to each class. The Hidden Markov Model (HMM) [3] is one type of statistical model and is chosen for classifying the movements in this research. HMMs are ideal as they implicitly model all of the many sources of variability inherent in real movements, and are shown to work well in practice.

This paper details the implementation of a movement recognition system using the vertical component of the ground reaction force and Hidden Markov Models. Section 2 presents related work, examining approaches to movement recognition and giving context to this research. Section 3 illustrates the GRF patterns generated by typical human movements. Patterns must be sensed from the environment, and this forms the topic of section 4, which looks at the construction of a weight sensitive or "active" floor to sense the GRF. The choice of good features for time-series discrimination is a highly problem dependent task and is one of the fundamental steps in statistical pattern recognition. This is the topic of section 5 which describes the selection of features from the GRF time-series. Section 6 discusses the implementation of the recognition system using HMMs and determining suitable model parameters for each movement considered. Section 7 describes the implementation of the online recognition system in a distributed computing environment. Section 8 considers practical and entertaining uses of movement awareness, and describes two applications in detail. Section 9 concludes.

2. RELATED WORK

Human movement tracking has been an active area of vision research for several years. Movement tracking is the process of observing the position of many parts of the body, and tracking their movement. These motion capture systems have applications in computer animation and biomechanics studies. However, motion tracking should be distinguished from movement recognition. There are several approaches to motion capture using vision; marker-based, edge detection and motion-energy images. With marker based motion capture, passive (e.g. reflective) or active (e.g. IR LEDs) markers are placed on the body, which generate bright points in the image [4]. These points are then used to determine the position of various limbs, or the relative positions of adjoining limbs. Detecting the edges of arms, legs or torsos allows the angles between some limbs to be computed [5]. With these two approaches, it is possible to use a vector of angles to specify the orientation of each limb (or a sub-set) on the body relative to the limb it is joined to. These angles can then used in the pattern recognition process. An interesting untagged system for movement recognition uses motion-energy and motion history images to perform recognition using temporal templates [6]. This technique represents motion as changes between sequences of images and uses the pattern of these changes to model a movement.

These vision systems suffer when occlusion or back ground movement is present, operate in a limited workspace, and tend to be resource intensive. For example, the Microsoft EasyLiving [7] group have recently installed pressure sensors under seat cushions and the floor to determine if someone is sitting or standing. They need to distinguish between a stand and a sit as the occlusion of the chair prevents their vision tracking system from tracking people that are sitting down. Movement awareness of rising to stand or sitting down is a useful input to this system. Furthermore, applications often require more detailed information than just the type of movement, and this information is provided by characterization of the movement. This characterizion is essential for many applications. For example, in gait analysis the parameters of step duration, stride length, rate of force rise and fall, impulse, cadence, contact and stride time, foot progression angle and foot identity are important information. It is not possible to accurately extract all of these parameters using vision techniques alone.

The advent of micro-machined accelerometers allows another approach to determining the angles between the limbs [8]. However these operate in a similar fashion to a tagged system with many accelerometers having to be carefully placed. The accelerometers also require power and networking, are cumbersome and again cannot provide the parameters required for full characterization. These systems tend to be used for static position determination to represent a human model for ergonomic assessment.

The ground reaction force is used in clinical movement analysis (e.g. gait analysis). Force plates (e.g. Kistler [9]) model the force exerted by the user onto the floor, and their use is widespread in the biomechanics literature. Force plates use load-cells to determine the GRF in 1D, 2D or 3D. Force plates tend to be expensive and not suitable for wide-spread deployment in aware environments. Olivetti and Oracle Research (ORL)[1] developed the Active Floor [10] that provides the vertical component of the GRF. Their work focused on person identification using the GRF observed during the gait cycle. This work was later replicated at Georgia Tech using a different feature set to model the GRF signal during gait [11], and was deployed in their Aware Home [12] for person identification and tracking.

Sentient computing is the use of appropriate sensing technologies to maintain a representation of physical space in a world model, which allows shared a perception of the physical world between computers and people. This shared perception allows applications to be more responsive and useful by observing and reacting to the physical world; it allows the environment (physical space) to become the interface with computers. The information in the world model contains information such as the position and orientation of people and objects, environmental conditions, device capability and state, whatever. Movement recognition is another input to the model. This additional awareness can give information on how we live and work, what activities we perform, together with specific applications that use this information. Body motion can be used as a form of nonverbal communication, similar to gesture recognition. An attentive environment can sense these motions and use them as control inputs.

[1] now AT&T Laboratories Cambridge

3. MOVEMENT ANALYSIS

This section shows the GRF for several common movements. These examples illustrate the pattern classes which the system will attempt to recognize. The classes are actually defined by a large set of such sampled motions, captured under a range of conditions. We discuss physical origins of the characteristic features of each motion.

3.1 Ground Reaction Force

The GRF is a three-component vector representing the forces in the vertical, anterior-posterior and medial-lateral planes [13]. Each component measures a different characteristic of movement. The vertical component is primarily generated by the vertical acceleration of the body and is of highest magnitude. In the remainder of this paper the term ground reaction force (GRF) refers to the vertical component only, unless stated otherwise.

As body mass is short term fixed, the force experienced by the floor is dependent on the acceleration of the body acting upon it[2]. If the GRF is less than body weight, then the weight of the body is not being supported by the floor and this signifies acceleration downwards. For example, when you crouch there is a downward acceleration of the body and so the GRF will be reduced. And conversely, when thrusting the body upwards (as in a jump) additional force or acceleration is required to thrust upwards, and this is experienced by the floor. When the vertical force is normalised to body weight, the resultant time-series is the acceleration profile of the movement.

3.2 Some Movements

The GRFs for several movements are discussed in the following sections. The GRF is normalized by body weight (i.e. normalized GRF = GRF/Body weight), and some movements are concatenated in the GRF trace as they exhibit a natural relationship.

3.2.1 Step

The GRF for a single *step* is shown in figure 1. Heel contact is the first event, with maximum load of approximately 110% body weight occurring between 5-10% of the gait cycle (1). As the knee then bends the force is absorbed thus causing an acceleration less than body weight (2). This stage ends with the leg straightening and the foot being flat on the

[2] treating the trunk as center of mass

floor, which results in a second turning point (3). The remainder of the waveform corresponds to the thrust required to continue the step and generate forward propulsion (4).

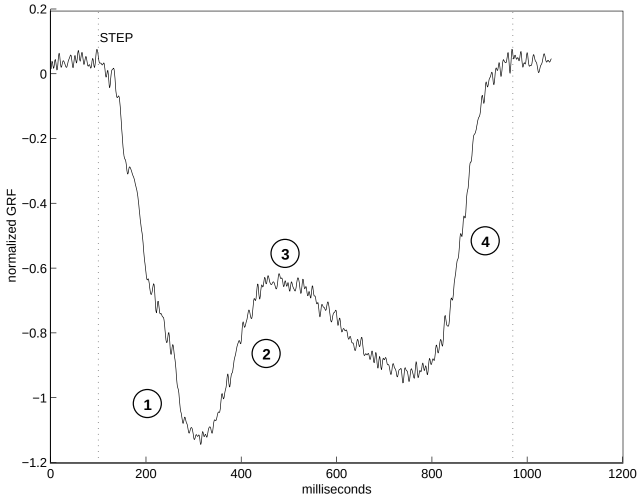

Figure 1: Step

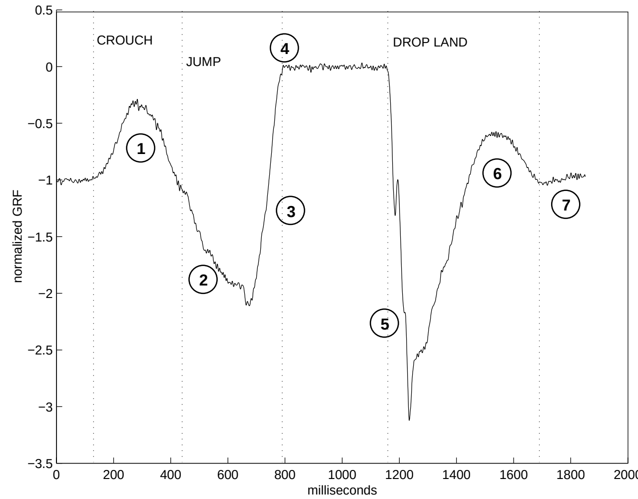

Figure 2: Crouch, jump and dropland

3.2.2 Crouch, jump and dropland

The movements of crouch, jump and dropland are presented in figure 2. Prior to jumping a human is required to lower the centre of mass (COM) by hip and knee flexion. The lower the COM drops the more distance is available for the push of the actual jump. This *crouching* movement produces a downward acceleration and reduces the force experienced by the floor, settling back to body weight upon completion (1). During the first phase of a jump, an increase in GRF is caused by exerting a force, through the leg muscles, to cause an upwards acceleration (2). This force continues until it ultimately exceeds the force of gravity and body lifts off the floor (i.e. *jump*). As the body begins to leave the floor there is a rapid unloading of the forces acting and the GRF quickly subsides (3). Following a jump no load is experienced by the floor during the in-flight phase (4). A *drop-landing* is characterized by a rapid and intense force being exerted on the floor (5). In the case of human movement, the toes make first contact, followed by the heels. The knees then bend to absorb this energy (6), followed by a settling period (7). The maximum force experienced during a dropland can exceed three times body weight, and this depends on the height from which the object falls (e.g. height of the jump).

3.2.3 Rise to stand and sit

The movements of rising to stand and sitting down are depicted in figure 3. *Rising* from a chair can be decomposed into phases of leaning forward, ascent, and establishing stability in standing. There is an initial requirement to create a force greater than body weight to cause an upward acceleration of the body (1). This maximum force decreases after the thighs leave the seat (2). Once the body is up to velocity there is an increase in GRF back to body weight. A similar pattern is apparent for *sitting* from standing. With a sit, there is an increase in GRF prior to seat contact caused by the breaking activity to control descent (4). The GRF rapidly declines as the thighs contact with the seat.

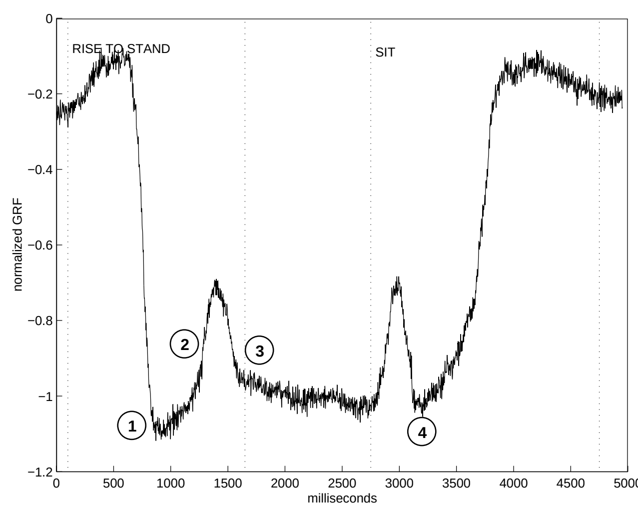

Figure 3: Rise to stand and sit

4. SENSOR SYSTEM FOR GRF ACQUISITION

The development of a recognition system requires the ability to sense the environment. The concern in this research is to sense the GRF from the floor. In an Active Floor the basic sensor is the strain gauge. Arranging matrices of load-cells and placing floor tiles on top allows the measurement of the GRF.

4.1 The Active Floor

The Active Floor was developed by Olivetti and Oracle Research (ORL) in 1995. It provides a mechanism for determining the vertical ground reaction force (GRF) experienced by a floor. Early work at ORL focused on identifying people by examining their gait using Hidden Markov Models [10]. This research focuses on using the Active Floor to provide ground reaction force information thus enabling the classification of basic human movements.

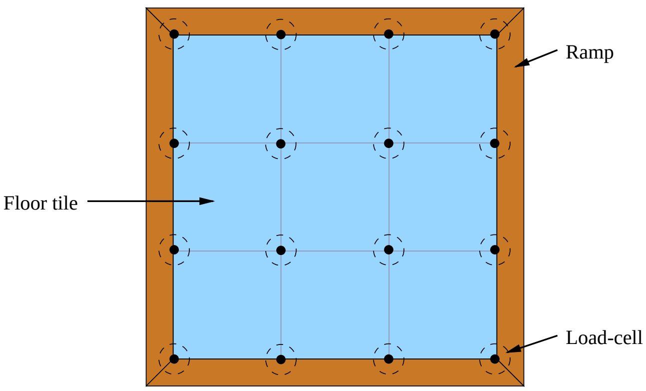

Figure 4: Plan view of the Active Floor

The current implementation of the Active Floor consists of a four by four array of load-cells with a three by three array of false tiles resting on top (figure 4). Each tile is composed of a steel plate, with a three quarter inch plywood board bonded to it, and dimensionally matches a standard floor tile (i.e. 500mm $\times$ 500mm). Standard floor carpet tiles then rest on these tiles. The corner of each tile sits on a load-cell. At the corners the load cell supports only one tile, and each vertex on the outer edge supports two. The inner load-cells support four tiles each.

The sensor bus is extended using a CORBA interface to the data acquisition device which makes the information accessible to interested applications. Services then operate on this data, for example to extract the GRF per tile. The GRF is the sum of the reaction forces acting upwards in an opposite direction to the applied load, irrespective of the position of the load.

5. FEATURE SELECTION

The aim of feature selection is to formulate a feature vector that captures the important information of the GRF signal, specifically that information which allows classification.

Initially, three basic features are extracted from the signal. Firstly the signal is normalized by body weight (BW) prior to feature extraction. This normalized signal is then segmented into windows of sample size n=20 observations, which corresponds to a duration of 20ms. For each window the features are extracted and represent this segment of the waveform. The primary three features extracted per window are: the mean ($\bar{x}$), the standard deviation (s), and the slope (m). The duration of the movement is inherently incorporated by the choice HMM type (left-right model) used in the classifier. A left-right model has the property that as time increases, the state index increases or remains the same, therefore the duration is modelled by the number of states and the state transition probabilities.

The *mean* (equation 1; where x_i represent the GRF data points of the window) provides information on the intensity of the signal and is a simple description of the underlying trend. The intensity is useful in describing the evolution of the force. For example, a drop-landing will exhibit a force well exceeding 250% BW, whereas a step will exhibit a maximum component at about 110% BW. The *standard deviation* (equation 2) is a measure of the variation of the signal. The acceleration has direct correlation with the acceleration of the body in the biomechanical process. The *slope* is calculated from the acceleration curve (i.e. normalized GRF) and captures the rate of change of acceleration. Some movements, e.g. a dropland, have a quite high initial acceleration, while others, such as the acceleration phase of a jump, have a more gentle slope. The slope is that of a least-squares straight line fit of the windowed data (equation 3; where y_1 and y_n are the end-point values of the fitted line). A slope of zero or near zero represents a static vertical force acting upon the floor. If this force is less than body weight, the body is accelerating downwards, and vice versa.

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{1}$$

$$s = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})^2}{n}} \tag{2}$$

$$m = \frac{y_n - y_1}{n} \tag{3}$$

These features ($\bar{x}$, s, m) represent a single window. To complement these features, *delta* and *acceleration* (i.e. first and second order regression) coefficients are appended to each feature vector. These coefficients add time-derivatives between the windows. The delta coefficients, d_t, are computed (eqn. 4) for each of the basic features $\bar{x}$, s, m, with c_t representing the basic feature to which the regression is applied.

$$d_t = \frac{\sum_{\theta=1}^{2}\theta(c_{t+\theta} - c_{t-\theta})}{2\sum_{\theta=1}^{2}\theta^2} \tag{4}$$

Equation 4 is then applied to the computed delta coefficients (i.e. $c_t = d_t; \forall\ \bar{x}, s, m$) to obtain acceleration coefficients.

6. CLASSIFIER DEVELOPMENT

The Hidden Markov Model (HMM) is used to classify the GRF generated by movements. In particular, the HMM toolkit (HTK) [14] is used to implement the classification system. This section describes the development of a recognition system using HMMs and a procedure used to determine suitable model parameters.

6.1 Model Development

An HMM can be defined by $\lambda = (A, B, \Pi)$, where $A = \{a_{ij}\}$ represents the transition probability distribution; $B = \{b_j(k)\}$ the observation probability distribution at time t in state j; and $\Pi = \{\pi_i\}$ the initial state probabilities. In the classification of movements, each movement is represented by an HMM. The first step is to specify a prototype model that defines the model topology (i.e. L-R or ergodic, number of states, mixtures, etc.). In a L-R model, the state transition probabilities are zero for previous states, i.e. there is no back-tracking. In an ergodic model, it is possible to transit from any state to any other state in the network.

Training is performed to adjust the model parameters to maximise the probability of the observation sequence (O) given the model, i.e. max $P(O|\lambda)$. This training is performed using the Baum-Welch algorithm, which is an iterative process to improve the model by locally maximising $P(O|\lambda)$. Embedded re-estimation is a technique to train several models from a unique source of data by updating all the models simultaneously.

Recognition is achieved by computing the likelihood of each model generating the unlabelled observation sequence and selecting the one with maximum probability. That is, given an observation sequence for an unspecified movement, compute the probability of the sequence given each of the models, and choose the one with maximum probability.

6.2 Determining Model Parameters

The choices relating to the type of model (e.g. L-R or ergodic), number of states and choice of observation symbols (i.e. discrete or continuous), depend on the signal being modelled. This research uses continuous densities in a L-R model configuration in which it is only possible to remain in the current state or progress to the next (i.e. Bakis-1 model). With a L-R model, the number of states is proportional to the average duration, and thus the state duration is inherently modelled, and it is neither necessary nor useful to explicitly include state duration probabilities [3].

Having specified the features (section 5) and the model type, the window size in feature extraction and the number of states must be specified. An iterative search was performed with window sizes from 10-30 milliseconds, and with the number of states ranging from 1 to 20, both with unit increments. The window duration was specified between 10ms and 30ms as the mean is not correctly estimated with a window duration of less than 10ms, and features are not well defined with the large window duration of 30ms. The maximum number of states (20) was not reached by any movement and deemed to be a sufficient upper limit. The observation sequences were manually labelled. For each window, feature vectors were generated. Then for each state a new model topology was specified and the models were trained using embedded re-estimation. Recognition was performed on labelled test sequences and performance evaluated.

Movement	# States	Model
Crouch	11	Bakis-1
Sit	6	Bakis-1
Jump	8	Bakis-1
Step	8	Bakis-1
Rise to Stand	8	Bakis-1
Dropland	7	Bakis-1
Static	1-2	Ergodic

Table 1: Number of states per movement

Movement	# Samples	Hit	Ins.
Crouch	102	102	1
Sit	34	34	0
Jump	102	102	0
Step	27	27	0
Rise to Stand	34	34	0
Dropland	102	102	0
Static	315	312	0

Table 2: Results

A window duration of 20ms was found to represent the signal sufficiently. The number of states in the HMMs can differ. Table 1 lists the movements and the corresponding number of states.

6.3 Evaluation

It is necessary to show the performance of the classifier with data not previously used for training the models. To this end, several minutes of new GRF data were recorded and the movements also video-taped. These GRF sequences were then put through the classifier which generates a list of recognized movements. This list is compared to the actual sequence of movements performed. Table 2 shows the number of test sequences for each movement with the number of correct hits and insertion errors. An insertion error occurs when a movement was incorrectly recognized as being present when it was not physically performed. These results show excellent classification performance.

The results show that the classifier can correctly recognize a sequence of movements and provides no information regarding the alignment of the movement boundaries with the physical movement. It is possible to manually label each movement in the test data and compare the alignment, though this is a laborious process. Instead, to show the alignment accuracy of movement recognition, a new GRF time-series composed of several movements was classified. Figure 5 shows this GRF on which the movement boundaries are indicated.

7. ONLINE MOVEMENT RECOGNITION

For the majority of sentient applications that use movement information, it is desirable to get the events in real-time. The system is composed of data acquisition, data processing (including classification), and event notification. Each of

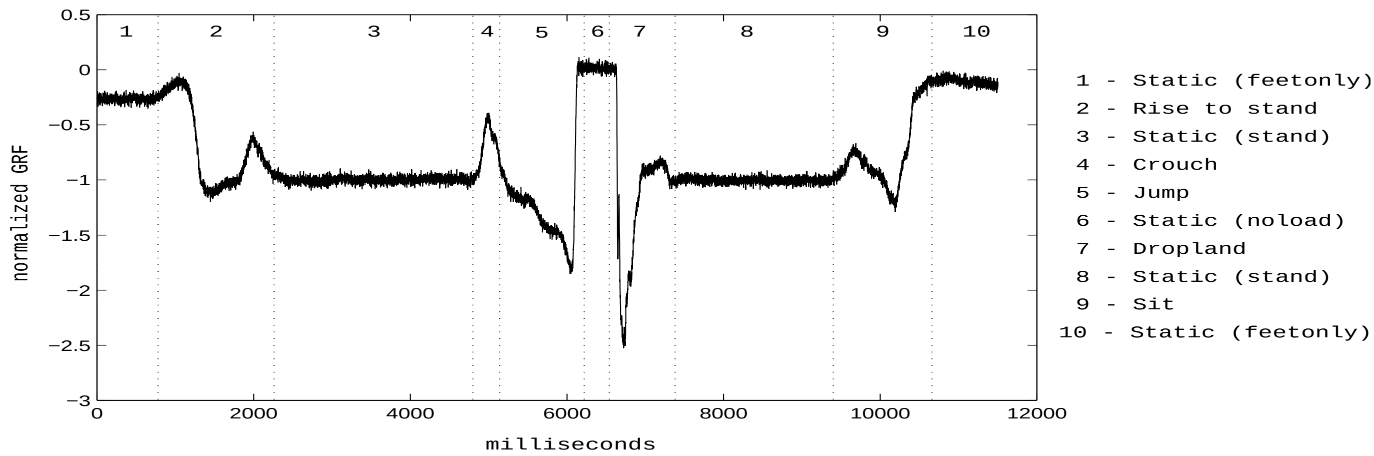

Figure 5: Labelled recognition

these stages are composed of several processes, as shown in figure 6, and operate in a distributed environment.

Data Acquisition

The load-cells of the Active Floor are connected to a data acquisition device (DAQ) that provides the signal conditioning and A/D conversion required for strain gauge measurement, and is connected to a server. As the DAQ is tightly coupled to the sensors of the Active Floor, which is in turn distributed throughout the environment, the DAQ is represented by a corresponding CORBA software object. This interface allows control of the DAQ and data acquisition over the network, effectively extending the sensor bus.

Data Processing

There are several stages of data processing, cumulating into a classification process to recognize movements. The process starts with *ground reaction force (GRF) extraction*, in which the GRF is the convolution of several sensors, depending on where the force is exerted on the floor.

A *movement detector* is employed to determine when movement is occurring on the floor, and to send this data for further processing. Initiating a new motion will involve a change in force, resulting in a non-zero slope. When there is a sufficient load, but no movement on the floor, an *adaptive weight estimator* determines the body weight of the user and uses this as the the normalization factor. As the user moves on the floor they are tracked and this normalization factor remains.

Feature extraction is the next stage of processing. The features are generated and then passed into the *classification* engine. Following the successful recognition of a movement, it is then *characterized* to extract information such as the duration of a step, which foot made the step, the height of a jump, the weight of the object lifted, etc. The next stage is to present this information to clients.

Event Notification

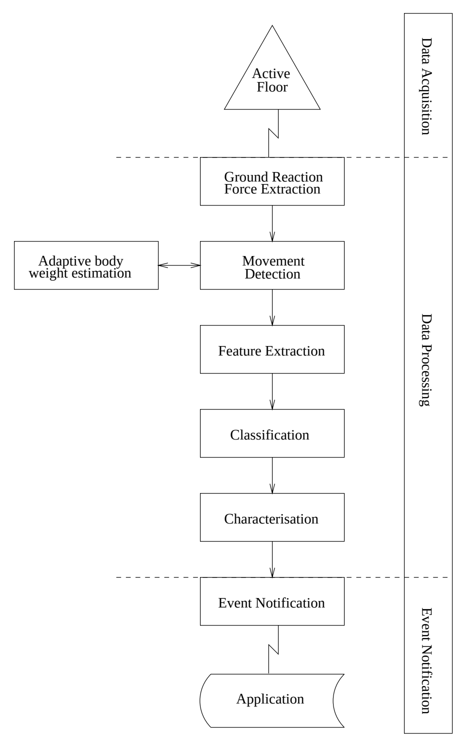

Figure 6: Block diagram of online-recognition components

A mechanism is desired that can notify interested clients of recognized movements. These clients and applications are likely to be distributed. Applications are notified using an event mechanism; we choose the CORBA notification service [15]. The notification service provides filtering of structured events, so clients can specify the movements and attributes they are interested in. Structured events, which consists of name-value pairs, are generated by the classification and characterization processes and are then either pushed or pulled onto the notification channel. Event filtering can be applied to these name-value pairs, ensuring that the client only gets the events they are interested in.

8. APPLICATIONS

The use of context in applications is slowly emerging. The output of the classifier provides additional context to an attentive system. This new interface, enabled by movement events, allows application developers to use more awareness when developing applications, making them more perceptive of the environment. As mentioned, this research is part of the larger effort of sentient computing, and this movement recognition is not in isolation, but contributes to the awareness of the system.

However many fun and practical applications are possible using movement recognition alone. Applications exist in safety, entertainment and the development of context-aware applications. Applications in safety extend from monitoring tasks for manual materials handling injuries, to fall detection of the elderly by examining their gait and balance.

The majority of entertaining applications stem from the ability to control a virtual character from physical movements or having an application respond to your location on the floor. For example, controlling the movement of a character in a game through your own movements gives an exciting new degree of interactivity, making the game more involving, and a new motivation for exercise.

8.1 Controlling Quake with real movements

The ability to control a virtual environment through physical movements allows for the development of entertaining applications. In particular, we use this exciting new awareness to directly control movement in a virtual environment, namely the game Quake. The control set is relatively small. Through movement, it is necessary to control direction, rate of progression, jumping, and shooting of weapons.

The direction can be forward, backward, left, right, and combinations. The direction is determined by the location of the force on the floor. Regions are defined that represent the directions, and the direction is set when the center of pressure is contained within a region. These direction control regions can be specified anywhere on the floor, and are set during a calibration stage. Both single and multiple tile configurations have been used. Using the three tile configuration (figure 7), the region of one tile represents going straight, with the other two representing left and right. On each tile a region is specified for forward and backward, thus facilitating the choice of any direction determined by the location of the force on the tile.

Progression is achieved by on-the-spot stepping on any of the

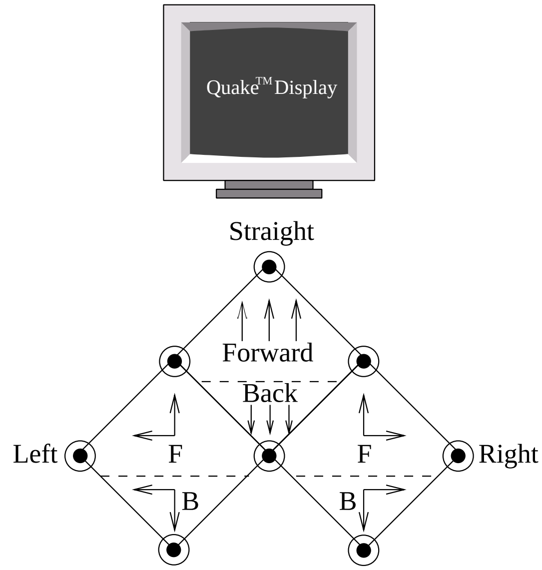

Figure 7: Plan view of selected floor tiles with markings indicating the regions used in determining direction

defined progression regions (tiles). The frequency of steps is used to control the speed at which the character moves, if at all. A physical jump in a progression region represents a jump in the virtual environment. The application uses the online classification of movements to determine that a jump has occurred. Finally an EMG sensor, placed on the arm, is used to detect a virtual trigger pull and to simulate a shoot event.

These control events are exported to the application by generating the corresponding keyboard events to the operating system. For the Quake control application a simple mapping between the movements and required keys is all that is required. This mapping can easily be altered to control the browsing or movement though any virtual environment (e.g. a VRML world).

8.2 Task Monitoring

The awareness of movement can be used for task monitoring during manual materials handling (MMH). MMH relates to people's interactions with objects that are lifted, carried, pushed or pulled. Lifting tasks make up a large proportion of MMH tasks and are involved in far more back injuries than other types of MMH tasks [16]. There are two approaches to the analysis of MMH for lifting tasks; biomechanical and physiological, and equations exist that give recommended limits for each. The biomechanical limits are determined in the NIOSH guidelines [16]. These guidelines specify the maximum permissible limit by considering the weight and position of the object being lifted, the height and range of lift, and the method and frequency of lifting. Sometimes we are not aware of our actions during MMH and a sentient environment has the potential to reduce these work related injuries. Using the Active Floor and movement recognition it is possible to identify that a lift has occurred, and the lift is characterized to extract the parameters for the NIOSH equation. Through the use of awareness it is possible to alert the potential of injury and to build a history of lifting tasks.

9. CONCLUSION

Sentient systems depend on awareness of their environment and inhabitants. This paper describes a new technique for recognizing whole body human movement using the GRF and HMMs. This new awareness is provided to sentient environments. This information enables the development of novel applications that respond in real-time to human movements.

10. ACKNOWLEDGEMENTS

The authors would like to thank Andy Hopper, Mike Addlesee, Diego López de Ipiña and AT&T Laboratories Cambridge for their support and advice in this work.

11. REFERENCES

[1] Andy Hopper. The Clifford Paterson Lecture, 1999 Sentient Computing. *Phil. Trans. R. Soc. Lond.*, A (2000)(358):2349–2358, 2000.

[2] Anil K. Jain, Robert P.W. Duin, and Jianchang Mao. Statistical Pattern Recognition: A Review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1):4–37, January 2000.

[3] Lawrence R. Rabiner. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 77(2):257–286, February 1989.

[4] Axel Mulder. Human movement tracking technology. Technical Report 94-1, School of Kinesiology, Simon Fraser University, July 1994. ftp://fas.sfu.ca:/pub/cs/graphics/vmi/HMTT.pub.ps.gz

[5] Christopher R. Wren, Ali Azarbayejani, Trevor Darrell, and Alex Pentland. Pfinder: Real-Time Tracking of the Human Body. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 19, pages 780–785, July 1997.

[6] G Bradski and J Davis. Motion Segmentation and Pose Recognition with Motion History Gradients. In *WACV*, pages 238–244, 2000.

[7] B. Brummit et al. EasyLiving: Technologies for Intelligent Environments. In *Handheld and Ubiquitous Computing*, September 2000.

[8] C.M. Gross and J.D. Lloyd. HumanTRAC: A New Method for Performing Rapid Ergonomic Product Assessments. Technical note, Centre for Product Ergonomics, Tampla, Florida, USA, 1996.

[9] Kistler force plate. http://www.kistler.com/

[10] M.D. Addlesee et al. The ORL Active Floor. *IEEE Personal Communications*, 4(5):35–41, October 1997.

[11] Robert J. Orr and Gregory D. Abowd. The Smart Floor: A Mechanism for Natural User Identification and Tracking. In *Proceedings of the 2000 Conference on Human Factors in Computing Systems*, The Hague, Netherlands, 1-6 April 2000.

[12] Cory D. Kidd et al. The Aware Home: A Living Laboratory for Ubiquitous Computing Research. In *Proceedings of the Second International Workshop on Cooperative Buildings (CoBuild'99)*, October 1999.

[13] Brian R. Durward, Gillian D. Baer, and Philip J. Rowe, editors. *Functional Human Movement.* Butterworth-Heinemann, 1999.

[14] Steve Young et al. *The HTK Book (for HTK Version 3.0)*. Microsoft Corporation, 2000. http://htk.eng.cam.ac.uk/docs/docs.shtml

[15] Object Management Group. Notification Service Specification, June 2000. http://cgi.omg.org/cgi-bin/doc?formal/00-06-20.ps

[16] Mark S. Sanders and Ernest J. McCormick. *Human Factors in Engineering and Design.* Psychology Series. McGraw-Hill, 7 edition, 1992.