The ORL Active Floor

M.D. Addlesee, A.H. Jones, F. Livesey, and F.S. Samaria[*],
The Olivetti and Oracle Research Laboratory

Abstract

In this article a novel type of sensor system called the Active Floor is presented that allows the time varying spatial weight distribution of the active office environment to be captured. The properties of the Active Floor are described showing that it differs substantially from other commonly encountered sensor systems. Furthermore classification of the footstep signature of a number of individuals is attempted by application of the hidden Markov model technique.

Introduction

We are concerned in this paper with a weight sensitive floor, to be used as a means of sensing the distribution and time variation of loads within a building. Data obtained from such a floor can be fed into a distributed location system for the Active Office[1]. The Active Floor is a square grid of conventional carpet tiles, each backed by 18mm plywood and 3mm steel plate, supported at the corners by cylindrical load cells which are instrumented to give us the total vertical force. In the data acquisition system described below we have found that the load cells are able to resolve weight changes of about 50 grammes. The grid has a 50cm spacing, and a sampling rate of 500Hz per load cell is employed.

Earlier work at ORL produced the Active Badge system[2] where each person or object wears a device or tag that can communicate via infra-red light with Badge Stations distributed around the working environment. In contrast to this a property of the floor that particularly interests us is that its operation does not require tagging or any special positioning of objects. Strong statements regarding weight and movement can be made about the environment without elaborate interpretation, and there is a wealth of data for more advanced analysis.

To see why the floor is different from other sensing systems that we have encountered, we analyse its properties in comparison with other systems.

Comparison with tagged systems

Sensing means such as radio frequency identification (RFID) tags and infrared active badges carry identification codes so that the particular tag can be directly identified. Bar codes carry sufficient information to distinguish one class of identical object from another. These tagged systems make strong statements about the tags that are sensed. If a tag is not sensed, we cannot deduce that the corresponding object is not present because its tag may be detached, inoperative, or shielded.

The active floor does not identify an object, except in so far as this can be inferred by interpretation of the sensor outputs. Objects which have identical weight signatures cannot be distinguished, but we can track their position and behaviour independently.

It is difficult for an object to be hidden from the floor, it would have to be suspended from the wall or ceiling. Thus we can make strong statements about what is not happening in the environment. For

[*] This article is based on the work carried out when the third and fourth authors were with The Olivetti and Oracle Research Laboratory, Cambridge, UK.

example, we can assert that nothing of significant weight has been removed from a desk if there have been no changes in the sensor outputs under and around that desk.

Comparison with interpretative systems

Analysis systems have been constructed to look for the characteristics of certain object classes in the output from video cameras or microphones[3]. For example, much work has gone into face and voice recognition. The active floor can also be used in this way, for example to attempt to recognize individual people from the way that they walk.

With somewhat less interpretation, we can track objects by their movement in video images, or by the relative arrival times of sound at a microphone array. A corresponding mode for the floor is to track objects by following their centre of pressure as they move. In all of these systems, we may lose track of an object if it passes too close to another of similar properties or changes properties in an undetected way. Finally, at a coarse level, the floor can tell us if there is activity in an area, just as video can detect motion and microphones can detect sound levels.

One unusual property of the floor is that it accurately tells us the total weight in an area. This is because the entire floor and loads on it are supported only by the load sensing elements, and because of the laws of mechanics, we cannot hide a weight inside or behind another. Most other untagged sensor systems only tell us about some of the objects in a room, as one object may be able to mask another to the system. For instance, if three people enter a room, the floor will accurately show an increase in weight corresponding to the sum of their three weights, including anything they are carrying. While a person is walking across the floor, the load will vary with time as the centre of mass of their body moves up and down. However, the centre of mass will on average remain at the same height, so the average load sensed by the floor will converge to the weight of the person.

To sum up, the active floor is a sensing system that does not require tagging, that senses all the objects in a well defined area, and that does not require any complex recognition software for its most fundamental operations.

Applications

Many applications exist for an active floor. It may be used as part of a virtual reality system or integrated into an intelligent building. The view shown on a display could be varied depending on the position of a person in the room. It could act as a people detector for lifts or shops. Switching on or off pieces of equipment as people approach or signaling a warning if someone enters the operating area around a machine tool. In security systems it could be used as an adjunct to tagging systems. For example, if a person identified by a tag enters a room and the net weight of objects in the room is detected to have changed after they leave we can be certain that something has been picked up or delivered and by whom. Furthermore by looking for similar weight changes in other locations we can attempt to track untagged items.

Floor technology

Once we had made the decision to measure the weight distribution across the floor, we reviewed the available technologies with which this could be done. We would like to have used a fine grain pressure sensitive technology such as that from Tekscan of South Boston, Massachusetts, but this seemed to have some disadvantages for our application:

1. Some of the weight is supported by the areas between the pressure sensitive pads, so the total weight on an area of floor cannot be accurately ascertained.
2. High point pressures will damage the small sensors, e.g. high heels. Indeed, the available technology is not recommended for sustained use even with flat shoes.
3. The technology is expensive when applied to a whole office floor (of the order of $3000 per tile).

For these reasons, we opted instead for the load cell approach. In so doing, we lose the opportunity to image process the outputs to see, for example, which way feet are pointing.

We first considered setting each tile on a single load cell. This would be difficult to do mechanically since no vertical load must be transmitted except through the single load cell. We would however be able to measure precisely the weight of an object that was completely supported on that tile. Of course, this would not always be the case. Objects would often span several tiles, and hence several load cells, so we would have to face the problem of interpretation of spread loads. Setting load cells at the corners of the tiles, so that each load cell measures the sum of contributions from the corners of four different tiles, resulted in a much simpler mechanical arrangement, and does not make the problem of interpretation any more difficult.

The load cells used (figure 1) are made out of high strength L168 aluminium alloy and are of the diaphragm type. A 350 Ohm full bridge strain gauge is incorporated which measures the compressive load. The particular load cells used in the prototype active floor have a rated load of 500kg with combined non-linearity and hysteresis error of less than plus or minus 0.5% at full rated output.

Our prototype Active Floor consists of a four by four array of load cells. Each load cell supports the corners of four adjacent floor tiles (or two and one floor tiles at the sides and corners of the Active Floor area respectively. See figure 2) Each floor tile is formed from a 500x500mm sheet of 3mm thick mild steel screwed to a similar sized piece of 18mm thick 5-ply board. Strips of metal at each corner of the mild steel sheet help locate the sheet corner accurately over the central point of the load cell. Carpet tiles are fitted on top of the floor tiles to give the appearance of a normal office floor. The layered composition is shown in figure 3. A wooden ramp on all four sides provides the transition from normal floor level to false floor level (figures 4 & 5). The false floor level being approximately two inches above the normal floor. No special preparation to the normal floor was required prior to laying the Active Floor, it being placed directly over the existing carpet tile floor covering.

The floor proved to be remarkably robust given its simple mechanical construction. In one test a colleague jumped from a desk onto the Active Floor to check for load cell overload. No damage to the load cells occurred, measurements maintained their repeatability, and all floor tiles remained in place with no signs of buckling. The Active Floor feels solid under foot with no tendency to spring. To check for spring and resonances in the floor we fitted a micromachined accelerometer to the lower surface of a tile at its central point. Changes in (vertical) position and velocity were then calculated by integration of the accelerometer sensor signal. Figure 7 shows the position and velocity of a tile as it is stepped on showing that it is free of serious resonances. Further resonance tests were conducted using a hammer blow to simulate an impulse for impulse response measurements.

Instrumentation

The strain gauges are typically provided with a 10VDC excitation voltage on one pair of wires with a second pair of wires providing the output voltage from the full bridge. With our prototype active floor we used VXI based strain gauge data acquisition equipment from Hewlett-Packard[4]. This consists of an E1421A six-slot C-size VXI mainframe populated with HPIB/command module E1406A, 64-channel 16-bit ADC board E1413B and four buffered strain gauge modules which gives the ability to connect a total of sixteen full bridge strain gauges and provides the appropriate excitation voltages. Data acquired by the ADC board is passed via the command module over a GPIB link to an Olivetti Pentium PC where it is processed using Hewlett-Packard's Visual Engineering Environment software. We found this data acquisition environment to be very flexible for experimenting with and the prototyping of various sensor based systems.

Measurements and removal of systematic errors

We initially sampled data from the floor at a variety of different sample rates and established that most of the load cell signal energy lay below 250 Hz for people walking and running over the Active Floor.

In our earliest experiments only four load cells were used, these being the ones supporting the centre tile of the prototype Active Floor. Dummy load cells were used in the other twelve locations. Data was acquired in three stages using three separate HP VEE programs. The first stage involved the collection of real time data from the four load cells and its storage to hard disk. Up to sixty seconds of raw data could be stored at a time using a sampling rate for each sensor of 500 samples per second. For the first few seconds of each sixty second acquisition period no activity was allowed to occur on the Active Floor to allow for later calibration of the system.

During the second stage the first 819 samples of each sensors data set were averaged and the value so produced subtracted from that data set, thus removing any systematic errors due to the static weight of the false floor. Next multiplication by a conversion factor turned the readings from Volts to Kilograms. Adding the conditioned values obtained at each sample time from the four sensors results in a measurement of the total weight of any object placed on the centre tile of the prototype Active Floor. In the biomechanics literature concerning walking and running [5] the reaction produced by a measuring device to the weight and inertial forces of a body in contact with such a device is referred to as the ground reaction force (GRF). Clearly what we are measuring with the prototype Active Floor is the vertical component of the GRF. Previous biomechanics research includes measurement of all three orthogonal components typically obtained using a Kistler force plate[6,7].

The third HP VEE program took the output of stage two and extracted automatically the footstep vertical GRF time traces or signatures. A typical trace is shown in figure 6. During stage one each subject was asked to walk repeatedly across the prototype Active Floor for sixty seconds at a time in as normal a manner as possible. The subjects were also told to target the centre tile aiming to step fully on that tile avoiding any overlap with adjacent tiles. It was further requested that they avoid walking on any of the eight peripheral tiles, which helped in the extraction of individual footstep traces, but this was not always possible for those people with shorter stride length. A database was assembled consisting of twenty footstep vertical GRF traces for each person from a group of fifteen people. This database was later extended to forty traces for ten members of the group.

Our earliest attempts at classifying different people's footstep traces focused on the application of the Hidden Markov Model technique as we had previously been using this method with some success in the context of face recognition.

Introducing Hidden Markov Models

An HMM is a doubly-embedded stochastic model based on an underlying Markov chain of states and an output probability function associated with each state in the chain. HMMs are generally used for modeling non-stationary vector time-series. The temporal nature of the signals generated from the floor makes HMMs an obvious choice for the purpose of classification. HMMs have been extensively used in speech recognition[8] and recently in image analysis[9]. A comprehensive tutorial on HMMs can be found in Rabiner[10]. In brief, the time signal is sampled to generate a vector time-series of finite length T denoted as $\boldsymbol{o}_1 \ldots \boldsymbol{o}_T$. A continuous HMM consists of a sequence of N states. Observations are generated every time a transition from a state to another occurs. The probability of a transition from state i to state j is a_{ij}. The probability of observing the symbol $\boldsymbol{o}_t$ at time t in state j varies according to a multivariate Gaussian distribution denoted by $b_j(\boldsymbol{o}_t)$. Allowing for a start and end state, an HMM can be fully defined by the following set of parameters:

- N is the number of states in the model.
- $A=\{ a_{ij} : 1 \le i, j \le N \}$ is the state transition matrix.
- $B=\{ b_j(.) : 1 \le j \le N \}$ is the output probability function.

In short-hand notation, a given model can be summarized as $\lambda=\{A,B\}$. Various state topologies are possible and a common one adopted is the left-to-right configuration illustrated in figure 8, where transitions are only allowed in a forward direction.

For a given model λ, the joint likelihood of a state sequence $Q = q_1...q_T$ and the corresponding observation sequence $\boldsymbol{O} = \boldsymbol{o}_1...\boldsymbol{o}_T$ is obtained by computing the probability of traversing the HMM trellis with Q and $\boldsymbol{O}$ as:

$$P(\boldsymbol{O},Q \mid \lambda) = b_{q1}(\boldsymbol{o}_1)\left[\prod_{t=2}^{T} a_{q_{t-1},q_t} b(\boldsymbol{o}_t)\right]$$

Equation 1

In practice, the state sequence Q is unknown and Equation 1 cannot be evaluated. Instead, the likelihood $P(\boldsymbol{O} \mid \lambda)$ can be computed by summing over all possible state sequences:

$$P(\boldsymbol{O} \mid \lambda) = \sum_{Q} P(\boldsymbol{O},Q \mid \lambda)$$

Equation 2

There is a simple procedure for finding the parameters of model λ which maximise Equation 2. The procedure is known as the Baum-Welch algorithm[11]. This is the fundamental procedure used for training the HMM. Identification is carried out simply by computing the likelihood of each HMM generating the unknown observation sequence and then selecting the model which gives the highest value. The exact details of the training and testing procedures are beyond the scope of this paper and a full explanation can be found in the references provided above.

In the work presented here, the foot signature of individuals will be represented by left-to-right models. The parameterisation of the models plays a fundamental part in designing a successful classifier. For this reason, a careful study of possible parameterisations will be presented through a set of experiments where various ways of generating individual models are explored.

HMM parameterisation

Assuming that a signature is a discrete signal with Z points, an observation sequence $\boldsymbol{o}_1...\boldsymbol{o}_T$ can be generated by sampling the signature using a window of width L and allowing for M-point overlap between successive observations. The sampling technique is illustrated in figure 9.

The L points contained in the window are arranged in an L-element column vector $\boldsymbol{o}_t$. The length T of the observation sequence is then obtained as:

$$T = \left\lceil \frac{Z - L}{L - M} \right\rceil + 1$$

Equation 3

where $\lceil x \rceil$ is the largest integer r such that $r \leq x$. A parameterised model can be specified in short-hand notation as:

H=(N,L,M)

Experimental Set-up

All the HMM-based experiments reported in this paper were carried out using the *HTK: Hidden Markov Model Toolkit V1.3* developed by Young[12]. The footstep database consisted of 20 signatures for each of 15 different individuals. Of these, 10 signatures were used for training and 10 for testing. Hence, a total of 150 signatures were used for training and another 150 for testing.

Model Parameterisation

The parameterisation of a model determines how successful a model is. Success is measured by the number of times the model correctly identifies one of the 150 test signatures. A corpus of approximately five hundred recognition experiments were carried out, with model parameters from the following range:

$$\begin{aligned} \boldsymbol{N} &= \{2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\} \\ \boldsymbol{L} &= \{5, 10, 15, 20, 25\} \\ \boldsymbol{M} &= \{0, 2, 5, 8, 10, 12, 15, 18, 20, 22, 25, 28\} \end{aligned}$$

Clearly, only values of *M* less than *L* were tried. Also, for large values of *N* and *L*, the overlap that was used started from larger values in order to generate a reasonable number of observations in the sequence. The results are shown in the figure 10. Each graph shows the error rate (y-axis) plotted against the discrete set of overlap values (x-axis) for various values of window size (the multiple line plots for each value of *L*). The following sections analyse the influence of each of the parameters on the error rate.

Window Overlap (M)

When sampling the time signal, overlap is normally used to avoid an arbitrary, rigid partitioning of the signal, which could result in sample boundaries cutting across important features. The value *M* determines primarily how likely feature alignment is and a large value implies a smaller probability of cutting across features. From Equation 3, it is evident that the overall length of the observation sequence depends on *M*: a larger value of *M* implies a larger value for *T*, as the signal is over-sampled, hence increasing the number of observations. A larger value of *T* is generally preferable, because the ability of an HMM to develop an accurate data model depends on having a sufficient number of events from which to derive its parameters. However, a larger observation sequence length causes longer computing times at training and testing. It is expected that for large *N* and *L*, increasing the overlap should reduce the error rate. The results shown do not offer any conclusive evidence on the influence of *M*, with the exception of the cases of maximum *L* and *N*=11,12. Here it appears that increasing the overlap improves the overall recognition performance.

Window Size (L)

The sampling window size determines the size of the features that the HMM models. From Equation 3 it is also evident that *L* determines the length of the observation sequence. The plotted results indicate that as *L* increases, the recognition performance improves marginally for all number of states. There is a clear downward trend for the case with *N*=12, while for the other cases the trend is not as obvious.

Number of States (N)

The number of states represents the number of features used by the HMM to model the signal. Choosing the correct number of states is crucial for obtaining an accurate model of the data. However, the choice is not always simple and most models rely on subjective intuition, where the analyst determines by inspection how many different states are appropriate for the data. A small number of states forces the HMM to group together observations which may introduce a substantial variance and that would be better modeled by separate states. On the other hand, if too many states are used, the model may not have enough distinct events to estimate the parameters of each state correctly. The result graphs are plotted for various values of *N* and in particular figure 11 shows the average error rate (with error bars at $\pm\sigma$) from which it appears that the best performance is obtained for *N*=7. Indeed the best overall recognition was obtained with H=(7,25,2), where 13 footsteps out of 150 were misclassified.

Criticism of the HMM results

In later experiments with the HMMs we normalized the traces for weight. The weight estimator used was simply calculated as the average of the GRF over the step interval. This caused the error rates obtained with the HMMs to increase to 50% and higher. From this we concluded that direct application of HMMs in the way described above is inadequate. A better approach would probably be to segment the footstep traces into different phases such as initial impact, thrust maximum etc. and then apply HMMs to these combining the results obtained with other features such as weight, gait distance and gait period in a multidimensional cluster analysis.

Final Comments

We have seen that an Active Floor has certain useful properties that distinguish it from other sensor systems.

- Objects directly detected, no special gadgets need be worn
- Objects cannot be shielded
- No fields to give rise to EMC problems
- Can be used in combination with other sensor systems

Our HMM results, although not wholly successful, are nonetheless an interesting first step in the classification of footsteps. A much harder problem is that of recognizing individuals performing more complex actions such as turning, jumping, running, etc. However, the footstep signatures used above represent only a small part of the information that could be obtained from the Active Floor.

An array of sensors can provide information on the distribution of vertical GRF over the area of the floor. In circumstances where only one object is in contact with the floor at a time it is possible to calculate precisely the centre of pressure. This is the average point of application of the applied force due to the object. However, when two or more objects are in contact with the floor it is impossible to uniquely calculate a set of centres of pressure for these objects from the sensor array readings unless one has a sufficiently dense array relative to the size of the objects to be located. Obtaining the spatial position of multiple objects (using such techniques as Maximum Entropy), identifying and tracking them is a subject of current research at ORL. We are also looking at Universal Serial Bus as a possible means of networking the sensors in a larger scale deployment of the Active Floor.

Biographies

Dr. Michael D. Addlesee (C.Eng. MIEE MIEEE) received the B.Sc. (Hons) degree in Physics from the University of Nottingham in 1981, the M.Sc. degree in Communications Engineering from Imperial College London in 1986 and the Ph.D. degree in Electrical Engineering from the University of Cambridge in 1991. In 1985 whilst working for Marconi Communication Systems, he was awarded a GEC Fellowship for a years study at the University of London. His Ph.D. research on Image Compression using the Subband Technique was, in part, funded by Thorn-EMI Central Research Laboratories, the IEE and the Fellowship of Engineering. In May 1991 Dr. Addlesee joined ORL, as a research engineer, where he continues to pursue research interests in digital signal processing, image compression and portable multimedia computer networking.

Dr. Alan Jones is a Principal Research Engineer at ORL. He was awarded a BA in Physics in 1981, Computer Science Diploma in 1982, and a PhD in Computer Science in 1986, all from Cambridge University, England. He has worked on control systems, sensor devices, optimisation, simulation, real-time multimedia and a variety of mathematically oriented projects. Current interests include communication and navigation systems, as both a keen user and a developer.

Finbarr Livesey received his B.Sc. in Physics from University College Cork, Ireland, in 1994 and the Diploma in Computer Studies from the University of Cambridge, UK, in 1995. On finishing the Diploma, he worked with Olivetti Research Labs in Cambridge for a short time, after which he joined Cambridge Consultants Ltd., Europe's leading technology consultancy. His main work concentrates on design for future Air Traffic Management systems, and his other interests are signal processing, human machine interaction and innovation management.

Dr Ferdinando Samaria - Ferdinando Samaria holds a PhD in image processing from the University

of Cambridge, England. He was with Olivetti Research Ltd (Cambridge) between 1994-96, where he worked on software agents and data classification. He has published various papers in face recognition and multimedia. He is presently working in the area of financial derivatives for Credit Suisse Financial Products in their London office.

Pictures

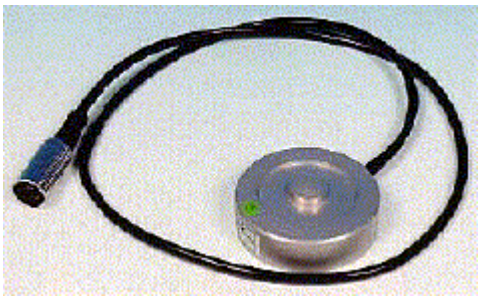

Figure 1 Photograph of one of the load cells used.

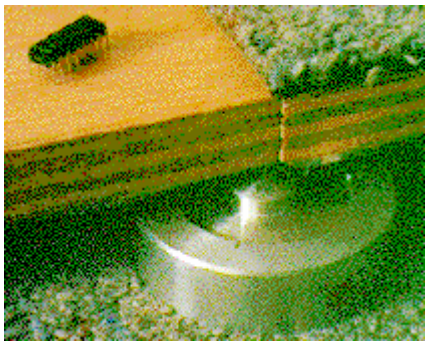

Figure 2 Load cell supporting two tiles.

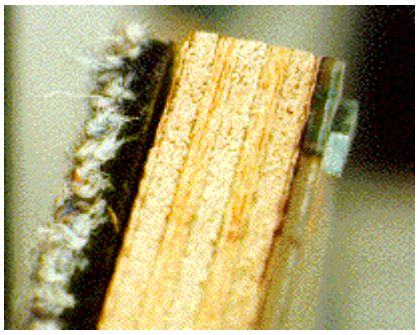

Figure 3 Corner of an Active Floor tile showing construction.

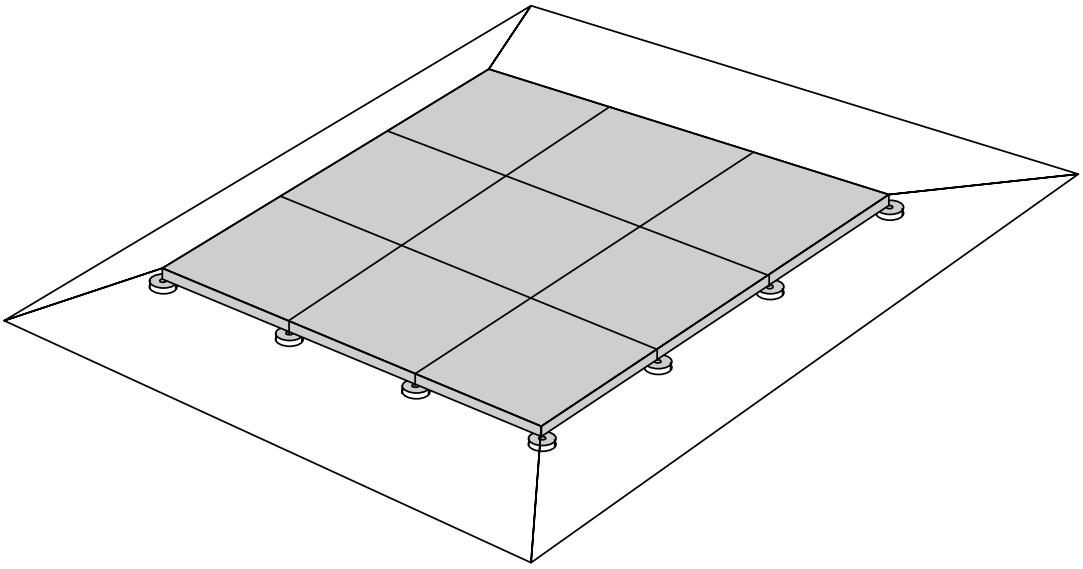

Figure 4 Arrangement of the prototype Active Floor.

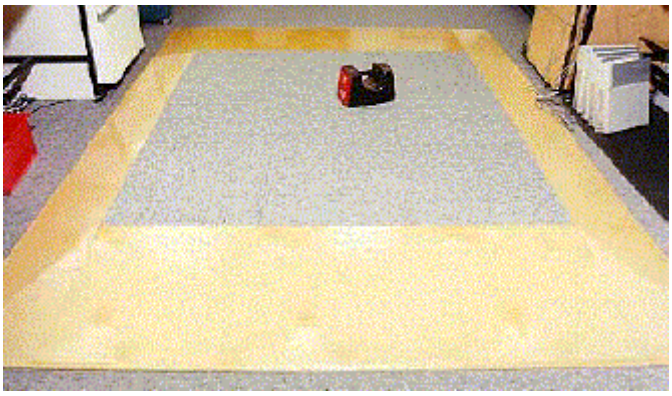

Figure 5 Prototype Active Floor with test weight.

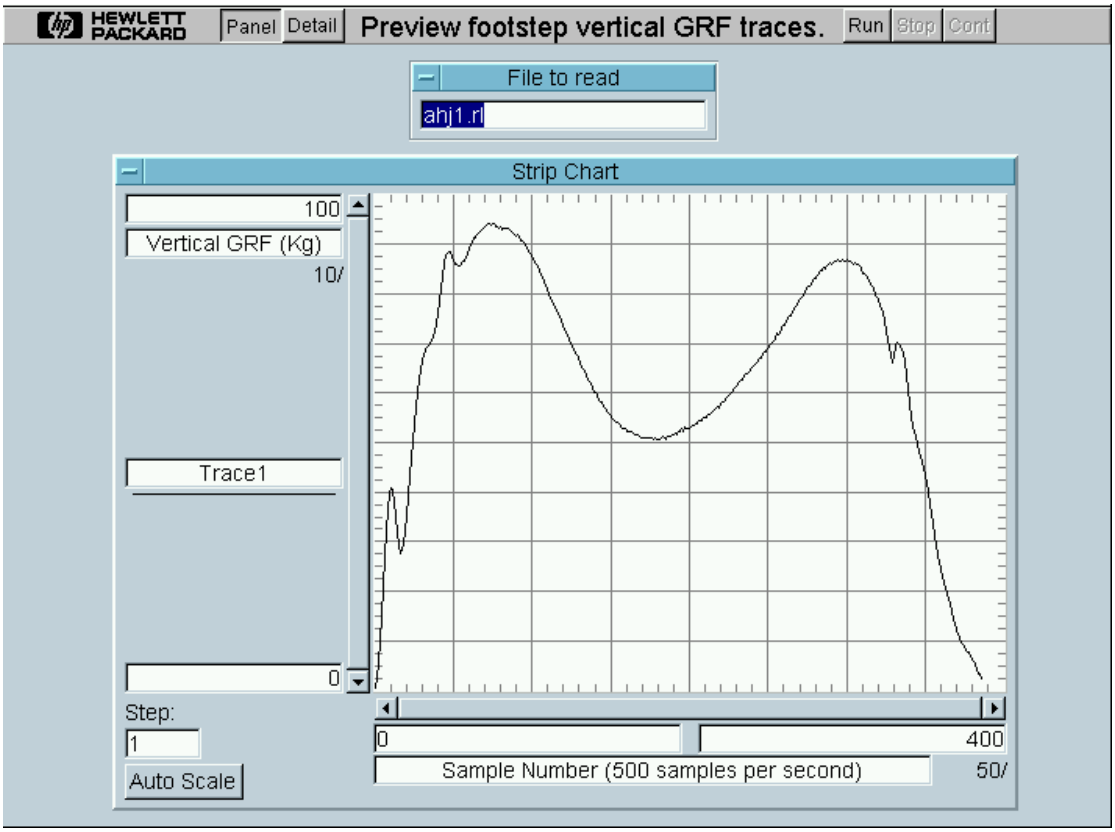

Figure 6 Typical footstep vertical GRF trace.

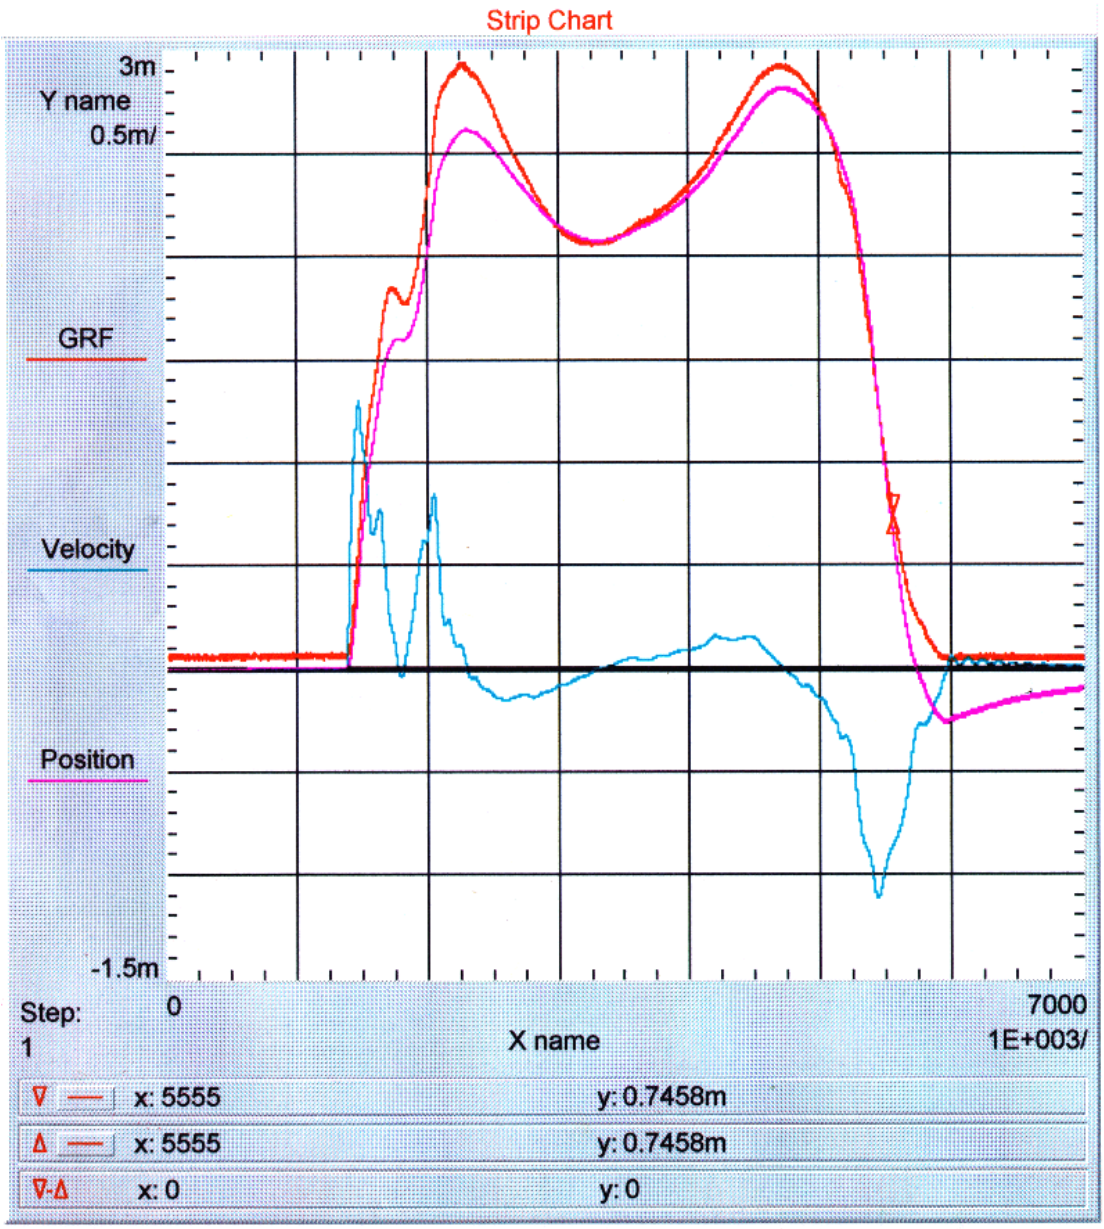

Figure 7 Results obtained with the Accelerometer.

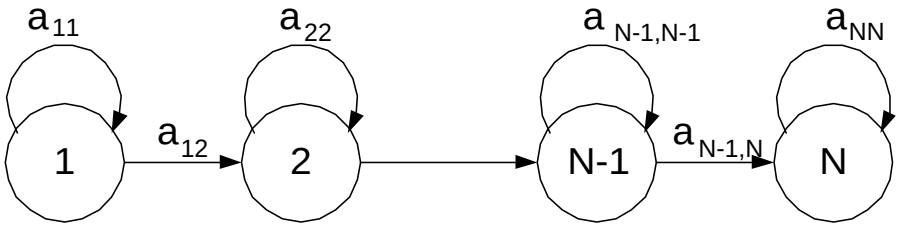

Figure 8 A left-to-right HMM.

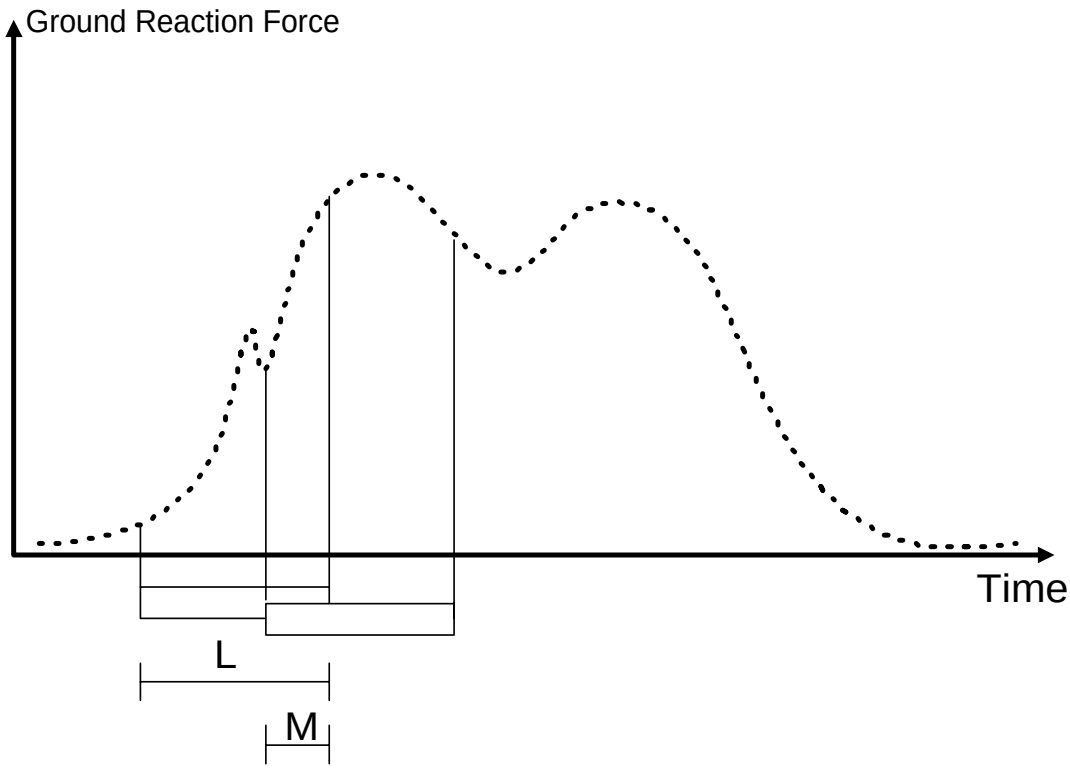

Figure 9 The Sampling technique.

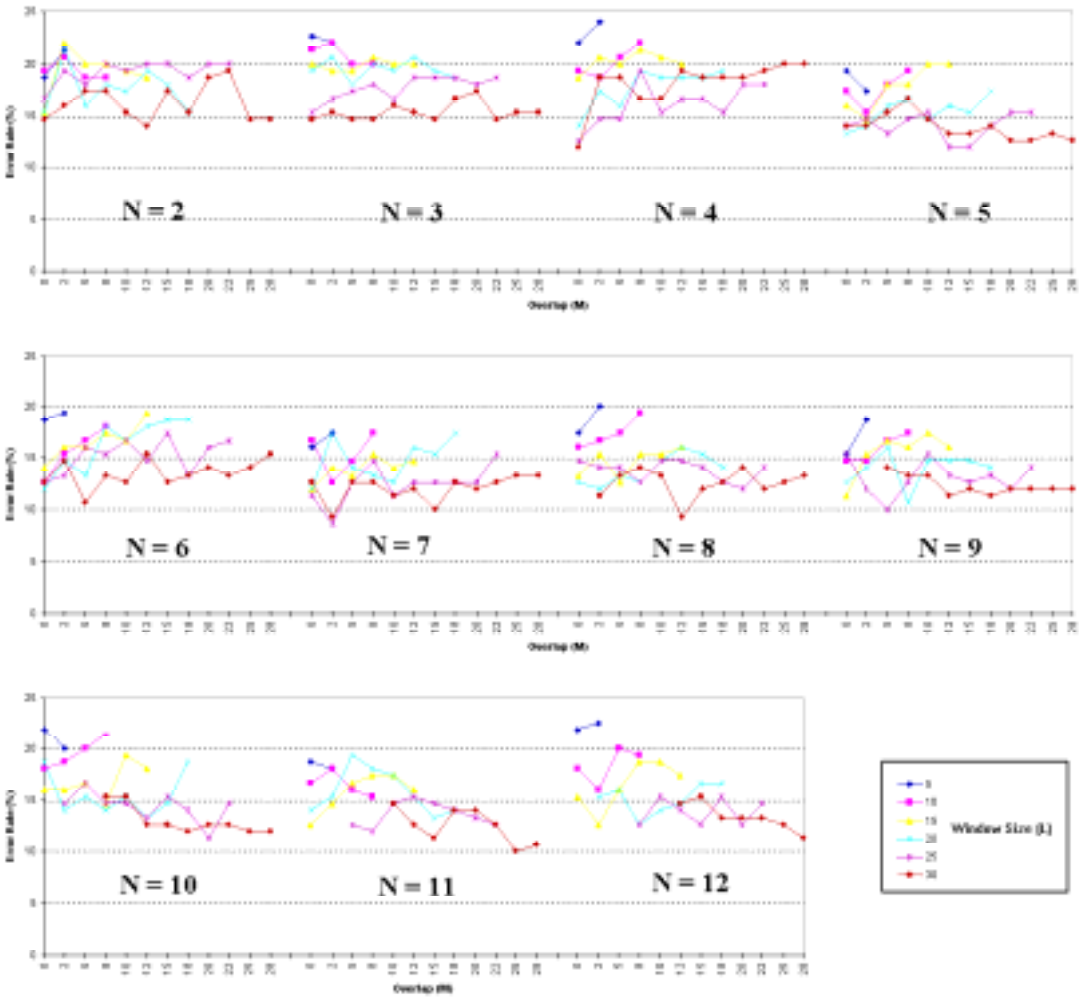

Error rates obtained with different sets of values for L, M and N.

Figure 10 Results for different parameterisations H.

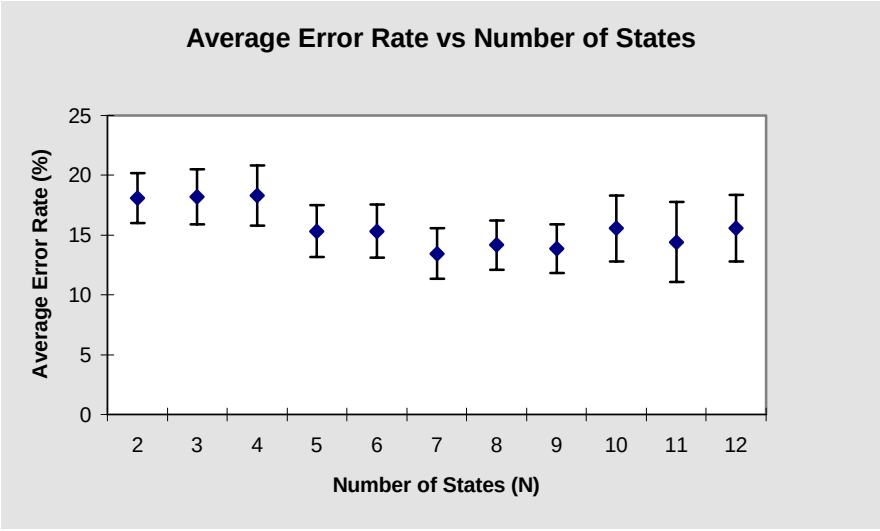

Figure 11 Average error rate versus number of states in HMM .

Figure 12 A more complex problem to solve!

References

[1] A. Harter and A. Hopper, *A Distributed Location System for the Active Office.* IEEE Network, Vol. 8, No. 1, January 1994.

[2] A. Hopper, A Harter and T Blackie, *The Active Badge System.* INTERCHT93, Amsterdam, April 1993.

[3] F. Samaria, H. Syfrig, A. Jones and A. Hopper, *Enhancing Network Services through Multimedia Data Analysers.* Proceedings of ACM Multimedia'96, 4th International MultiMedia Conference and Exhibition, Boston, USA, November 1996

[4] Hewlett Packard *HP75000 Family of VXI Products and Services.* 1995 Catalog, HP pub. no. 5963-3718E.

[5] P.R. Cavanagh, *The shoe-ground interface in running.* American Academy of Orthopaedic Surgeons Symposium on "The foot and leg in running sports", Coronado, California, September 1980. Published by The C.V. Mosby Company 1982.

[6] A. Pedotti, *Simple Equipment Used in Clinical Practice for Evaluation of Locomotion.* IEEE Trans. Biomedical Engineering, Vol. BME-24, No. 5, pp 456-461, September 1977.

[7] G.C. Santambrogio, *Procedure for Quantitative Comparison of Ground Reaction Data.* IEEE Trans. Biomedical Engineering, Vol. 36, No. 2, pp 247-255, February 1989.

[8] L.R. Rabiner and B-H. Juang. *Fundamentals of Speech Recognition.* Prentice-Hall, 1993.

[9] F. S. Samaria. *Face Recognition Using Hidden Markov Models*. PhD Thesis. Cambridge University Engineering Department. Available from: ftp.cam-orl.co.uk:/pub/users/fs/thesis.ps.Z

[10] L.R. Rabiner. A tutorial on Hidden Markov Models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257-286, 1989.

[11] L.E. Baum. An inequality and associated maximisation technique in statistical estimation for probabilistic functions of Markov processes. *Inequalities*, III:1-8, 1972.

[12] S.J. Young. *The HTK Hidden Markov Model Toolkit: Design and Philosophy.* Technical Report TR.153, Department of Engineering, Cambridge University (UK), 1993.