1-1-2000

A Real-Time Sitting Posture Tracking System

Lynne A. Slivovsky
Purdue University School of ECE

Hong Z. Tan
Purdue University School of ECE

A REAL-TIME SITTING POSTURE TRACKING SYSTEM

LYNNE A. SLIVOVSKY
HONG Z. TAN

A Real-Time Sitting Posture Tracking System

Lynne A. Slivovsky

Hong Z. Tan

Haptic Interface Research Laboratory

http://www.ece.purdue.edu/HIRL/

School of Electrical and Computer Engineering

1285 Electrical Engineering Building

Purdue University

West Lafayette, IN 47907-1285

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

As computing becomes more ubiquitous, there is a need for distributed intelligent human-computer interfaces that can perceive and interpret a user's actions through sensors that see, hear and feel. A perceptually intelligent interface enables a more natural interaction between a user and a machine in the sense that the user can look at, talk to or touch an object instead of using a machine language. Although research on haptic (i.e., touch-based) interfaces has received less attention in the past as compared to that on visual and auditory interfaces, it is emerging as a new interdisciplinary field that holds much promise for the future.

The goal of the sensing chair project is to enable a computer to track., in real time, the sitting postures of a user through the use of surface-mounted contact sensors. Given the similarity between a pressure distribution map from the contact sensors and a gray-level image, we propose to adapt computer vision and pattern recognition algorithms for the analysis of sitting pressure data. Work in three areas are proposed: (1) data collection for a sitting pressure distribution database, (2) development of a real-time sitting posture tracking system, and (3) performance evaluation of the tracking system. The realization of a robust, real-time tracking system will lead to many exciting applications such as automatic control of airbag deployment forces, ergonomics of furniture design, and biometric authentication for computer security.

THIS PAGE INTENTIONALLY LEFT BLANK

1. INTRODUCTION

The goal of the Sensing Chair Project is to enable a computer to track, in real time, the sitting postures of a user through the use of surface-mounted contact sensors. Given the similarity between a pressure distribution map from the contact sensors and a gray-level image, we propose to adapt computer vision and pattern recognition algorithms for the analysis of sitting pressure data. This work involves: (1) data collection for a sitting pressure clistribution database, (2) development of real-time sitting posture classification and tracking systems, and (3) performance evaluation of the systems. The realization of a robust, real-time tracking system will lead to many exciting applications such as automatic control of airbag deployment forces, ergonomics of furniture design, and biometric authentication for computer security. This work is divided into two parts:

- A real-time classification system for labeling a set of preselected typical sitting postures in an office environment (i.e., Static Posture Classification).

- A real-time tracking system that labels a person's sitting posture at steady-state as well as transitional postures (i.e., Dynamic Posture Tracking).

Preliminary work has focused on the construction of the Static Posture Classification system. First, a Static Posture Database was established. The database contains a total of 1500 sitting pressure distribution maps from 30 individuals. Second, a real-time multi-user Static Posture Classification system was developed. This system classifies sitting pressure distributions with an overall accuracy of 96% correct. Future work involve:; the development of a real-time Dynamic Posture Tracking system. A system like this is needed in order to represent a person's posture during transitions from one static posture to the next (e.g., from seated upright to sitting with one leg crossed over the other).

The system uses pressure sensors (developed by Tekscan, Inc.) that interface with a PC through a special interface board. They are mounted on a Herman Miller Aeron chair and capture pressure distribution information of a user seated in the chair. The setup is shown in Figure 1.1(a). The pressure information from the sensors is converted to an 8 bit raw digital value that can be treated as a greyscale image. A sample sitting pressure distribution map for posture *Seated Upright* is shown in Figure 1.1(b).

(a)   (b)
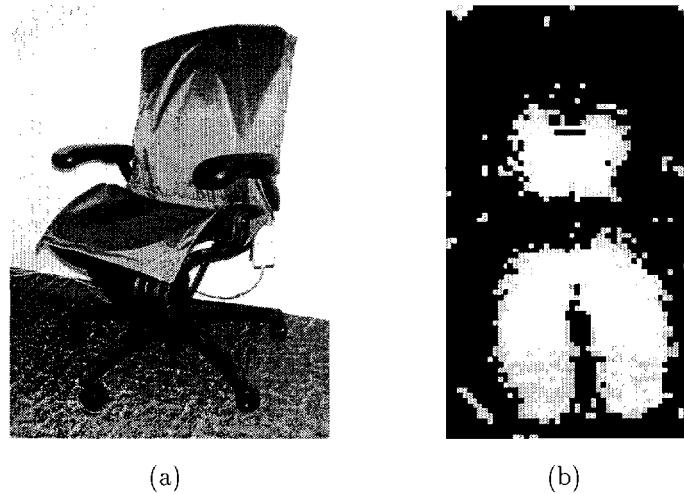
Fig. 1.1. (a) The Smart Chair. (b) An example of sitting pressure distribution maps displayed as an 8-bit greyscale image. The top half of the image shows the pressure distribution on the back of the chair, and the bottom half shows that of the seatpan. The top, bottom, left and right sides of the image correspond to the shoulder area, knee area, right side and left side of the person, respectively.

This work draws on aspects of many different fields. The remainder of this chapter presents related work in computer vision and pattern recognition, intelligent environments, pressure sensing, an overview of human anatomy, and ergonomics as it relates to sitting. Chapter 2 describes the Static posture Classification system and gives performance results. Finally, Chapter 3 concludes with a discussion on the planned implementation of the remaining goals.

1.1   Related Work in Computer Vision

Much emphasis in computer vision has been placed on developing compact representations of objects for recognition and pose identification (e.g., [28], [8]). Model-based methods

such as these perform matching between scene features and model features in their given recognition task. These features typically involve shape and other geometrical properties of objects. These systems work well but can break down when features cannot be extracted from a scene (e.g., due to occlusion) or when multiple object or pose hypotheses cannot be rectified.

This led researchers to study recognition systems using appearance-based approaches. Appearanced-based methods for classification and localization of objects employ an object model that is determined solely by the visual appearance of the object [45]. These models include object and scene parameters that get encoded as greyscale variations in the images. No user-defined feature-based model is needed. Section 1.1.1 will discuss two commonly used appearance-based methods: principal components analysis and linear discriminant analysis. When the scene is dynamic, object modeling and tracking becomes more complex. Section 1.1.2 overviews human modeling and tracking.

1.1.1 Appearanced-Based Methods for Object Recognition

Appearance-based methods use an object model that is dependent solely on an object's visual appearance. This is a combination of its shape, surface reflectance properties, pose in the scene. and illumination conditions. Shape and surface reflectance properties are intrinsic to rigid objects. To obtain an object model, an object is imaged under varying pose and illumination conditions. This process produces a set of training images where image pairs whose imaging conditions were similar are correlated.

The collection of training images can be used as the object model without any further processing. To recognize multiple objects a collection of training images is needed for each object. Recognition of a test object in a scene (imaged under similar conditions to the training images) is a matter of finding the closest training image to the test image. The identity of the closest training image is taken as the identity of the test irnage. To perform well under general conditions where illumination and view angle of the camera to the object can vary, the training set must contain a large number of samples in order to adequately represent these variations. The method for recognition just mentioned is a computationally expensive approach. For images of size RxC, the training set is a collection of points in a

N-dimensional space, where $N = RC$. Computation of the distance to each training image is therefore $O(N)$. If. however, the training samples are images of compact objects, the distribution of the training set can be represented by a low-order subspace of the original N-dimensional space. For training sets of images such as these it is advantageous to represent each training image by a vector in a lower dimensional subspace of the image space.

To obtain this low-dimensional representation, the technique of principal components analysis (Karhunen-Loeve Expansion) can be applied to the set of training images. The Karhunen-Loeve Expansion is given in in detail in [24]. Training of the system involves obtaining an eigen-decomposition of the covariance matrix formed by the training vectors (raster scans of training image samples). The eigenvectors of the covariance matrix form an orthogonal basis that spans the image space of the training vectors. A training sample is then represented by a vector of coefficients. Each element in this vector is the projection of the training image onto one eigenvector. An image can be reconstructed from the coefficient vector by summing the eigenvectors weighted by their corresponding coefficients. To reduce the dimensionality of a set of data, only a subset of the eigenvectors is used for reconstruction. To minimize the mean square error of this representation, the eigenvectors with the greatest variances (those corresponding to the largest eigenvalues) are used. Recognition of an unknown test object involves projecting the test image onto the reduced number of eigenvectors and finding the training instance (projected onto the eigenspace) closest to the test object in the lower dimensional space.

PCA has been used in systems for object recognition and pose estimation [48], face recognition [67], [31], [50], [4], [40], and single user posture classification [62]. Since the training of a system with PCA is based solely on the visual appearance of the training data, recognition (or pose estimation) faulters when trying to recognize objects imaged under different illumination conditions, positions, orientations, scale, with occlusion or with background clutter. Variations in illumination are usually overcome by normalizing the greyscale values in training images and testing images. This only succeeds when there are no other objects whose greyscale values can bias the normalization process in the scene. Training and testing images are scaled to pre-set dimensions to overcome variations in scale. This is inhibited by both occlusion and background clutter.

Early systems such as [48] make the assumption that objects, under no occlusion, can be segmented from a scene. This system trains a single eigenspace, referred to as a parametric eigenspace [47], on all of the training data from multiple objects and poses. Each object is then represented by a manifold in the low-dimensional eigenspace, parameterized by pose. Recognition of an object in a test image proceeds by finding the closest, manifold to the projection of the test image in the eigenspace and its pose is given by the closest point on that manifold. Methods to find the closest point on the manifold include binary search [49] and pattern rejection [2]. This system works well when the training and testing images contain a single object against a solid background. The system cannot handle occlusion and background clutter. The problem of background clutter is handled in [46] by applying an AND window to all images before training. This window indicates the area of each object that is visible over every pose. It is not useful however if too much of the object is AND-ed out after applying the AND window. It also does not address the problem of occlusion.

In [50] the authors introduce the concept of view-based eigenspaces (referred to as face spaces). Here, a face space is formed for all images of faces taken at a particular orientation (e.g., all frontal face images). The reasoning behind using an eigenspace for each view is that it would be better to represent a complex distribution (all face orientations) by multiple clusters (one for each orientation). This is analogous to obtaining an eigenspace for each class of a data set where the best low-order representation for each class is found. Recognition of a person is two-fold. First, the distance of a test image to each face space is computed. This is called the *Distance From Feature Space* (DFFS). It gives the pose of the face in the image (e.g., frontal, angled 45 degrees to the left). space corresponding to the view of the face in the image. Identification of the person then is a matter of finding the closest point in the face space, the one that minimizes the *Distance In Feature Space* (DIFS). This view-based method is used in our work for posture classification. Because we are doing posture classification, and not person identification, we need only to find the minimum DFFS to each of our posture spaces to obtain the posture class (e.g., sitting upright).

While the view-based method in [50] improved recognition accuracy of faces as compared to the parametric approach, it had yet to overcome problems associated with occlusion and background clutter. Systems using PCA to overcome occlusion and background clutter

include the use of eigenfeatures [50], local face areas [40], and a quadtree-based approach [10]. In all of these, training is performed on sub-regions of the training images. Recognition involves combining the results of identified sub-regions of a test image. In [40] a Hidden Markov Model is used on top of the eigendecomposition of the sub-regions of face images. This allows the system to be able to recognize a test face that is not found in the training set but is a compilation of sub-regions in the training set. The quadtree-based approach in [10] gives an hierarchical eigen-representation of each image. Recognition proceeds from a "coarse" level (the entire image) to "finer" levels (subimages). If an object is partially occluded, those subimages that aren't occluded can provide ample information to recognize the object and its pose.

Another popular appearance-based method is that of Linear Discriminant Analysis (LDA) This technique is also known as Fisher Discriminant Analysis (FDA) and Discriminant Analysis. This technique, like PCA, reduces the dimensionality of a problem to make it more manageable.

Training consists of finding the best projection of an N-dimensional set of data of K classes onto M dimensions, where $M < N$. LDA finds the M features that best discriminate between the different classes in the data. This is done by minimizing the within-class distance while maximizing the between-class distance. LDA has been used for face recognition [61], [4], [21] and mobile robotics [68]. In order to train a system using LDA, $N + K$ (the dimension of the data N plus the number of classes K) training samples are required. For vision problems where images can contain thousands of pixels, this constraint is rarely met. To overcome this, data is first projected onto an eigenspace to reduce the dimension. LDA can then be applied to the reduced dimension data.

When applying LDA after PCA, or using PCA in classification, care must be taken in the selection of eigenvectors used to represent the data in the eigenspace. Those features (eigenvectors) that best represent the data by minimizing the mean-square error are not necessarily those that can bst discriminate between classes in the data [24]. For example, one can describe a car as having 4 wheels and a windshield but cannot use these features to differentiate between a sports car and a luxury car.

1.1.2 Human Modeling and Tracking

This section describes work related to the tracking and understanding of human action in image sequences. Those methods that employ a 3-D description generally believe that the 3-D description is necessary and sufficient for interpreting human motion ([25], [26], [32], [54], [56], [70]). In contrast to this approach, others use appearance-based methods [1], [16], [71]. Still others [52], [60], [37], [17] use motion of regions of the body (as a whole or parts) to interpret action without reference to an underlying sequence of static images or poses. These are termed directional motion recognition [17].

Those systems using a 3-D description utilize a 3-D object model to recover the pose of a scene object in each frame of a sequence. As with any model-based system. it is necessary to be able to accurately extract features from the scene for tracking. Systems that maintain a history of the 3-D model alignment with the scene object over time can predict future poses of the scene object using a technique like Kalman filtering [26], [54], [56]. In [70], observed 2-D blob features are probabilistically integrated into a dynamic 3-D skeletal model. The 3-D model can be used to track the 2-D blob features through an extended Kalman filter. This approach directly couples 2-D and 3-D information. It also incorporates learned behaviors of human:; while performing a given task.

In [26] an arm is tracked against a solid background using a two cone model. A full body cylindrical model is used in [56] and a 22 degree of freedom super-quadric model is used in [25] to track human motion against complex backgrounds. It is not clear that these methods can be directly extended to our problem of tracking a human:: by their sitting pressure distributions since, in addition to modeling human shape, a moclel would need to take into account the weight distributiori of a person. In order to extend this type of approach to our domain, in the context of dynamic posture tracking, the 3-D model would need to incorporate deformations since our object to be tracked, the human body, is not projected onto an image plane but deforms against the surfaces of a chair. Such deformation could be captured with a finite element model coupled with a physics-based model of the human body.

Appearance-based systems interpret action from a sequence of 2-D images. Systems that use greyscale images to represent action include [16] and [69]. Here again, differences in body type can affect system performance. In [71] body silhouettes are used in a Hidden Markov Model framework. In [1] body contours and a simple 2-D body model are used to extract body parts. Drawbacks to these approaches include dealing with complex backgrounds and clothing and the necessity to examine the body as a whole, as opposed to regions of interest. In our case, we do not suffer from cluttered backgrounds as there is a single person sitting in the chair. Applying an appearance-based approach to our problem would be possible. A set of features in the pressure distribution map could be fed into an HMM to learn desired motion sequences.

Directional motion approaches examine blob-like and predefined body regions [52], [60], [37], [17]. In [52] cyclic walking motions are recognized through repetitive motion. A feature vector containing optical flow and periodicity measurements is used to track humans. In [60] an ellipsoid model of the body is used with optical flow measurements. These are combined into a phase portrait from which force, rotation, and strain dynamics are computed. In [37], two ellipsoids are used to model the body. One models motion region silhouettes and the other motion magnitudes. Gait is characterized by measures such as centroid movement and torque of the two ellipses. [17] use Motion Energy Images and Motion History Images to form a two component view-specific temporal template. This method temporally segments action sequences automatically and in real-time. The blob-like approach does not seem well suited for our problem. Problems can arise when, for example, the back is no longer leaning against the seat back.

1.2 Intelligent Environments

Computers can perform astronomical computations, yet they aren't intelligent. They do what we tell them to do. We enter commands and the computer executes them. A computer is aware of its surroundings only if the necessary information has been entered by a user through the low bandwidth devices of mouse and keyboard. Machine intelligence would allow a computer to interpret and anticipate the needs of a user while interacting with his/her surroundings. Intelligence can be achieved by giving a computer human-like senses

such as sight, hearing. and touch. This can be done by interfacing digital cameras for eyes, microphones for ears, and pressure sensors for skin along with algorithms that can interpret sensory input from these devices and from it, predict the user's wishes.

We can extend the notion of intelligent machines to include typically non computational aspects of our surroundings. such as a room, a chair, a desk, and clothing. The range of these items is endless. Most of us are familiar with the TV series Star Trek and the chronicles of its crew on the Starship Enterprise. The crew is able to naturally interact with the ship, which can locate and identify members of the crew and interpret their commands. This fictional intelligent setting is within reach of modern reality. Systems are being developed that can recognize, locate and track people, interpret gestural commands, understand natural language, and use all of this information to predict and assist the needs of a user (e.g., [12], [51], [39], [23], [22]).

Before creating any intelligent environment it is necessary to define what an intelligent environment entails (i.e., what it should be, what types of sensory capabilities it should possess and what roles it can play in the environment). There is debate as to what an intelligent, environment is. To some [12], it must be invisible to the user and require minimal hardware. To others [51], [39], it and the user are intertwined as in the case of smart clothes and cyborgs. Still others fall in between these two extremes [23], [62], [22]. To all, an intelligent environment must be able to assist the user in various tasks. These tasks depend on the type of human-computer interfacing defined, which include gesture, speech, affect, context. and intent, and on the desired roles to be performed by the intelligent environment, such as climate control.

Work in creating *rooms with intelligence* is currently being investigated by many researchers. This includes the smart room [51], the intelligent room [12], [66], the intelligent classroom [23], and intelligent houses [42].Cameras are placed in a room for person tracking, identification, and gesture recognition. Speech recognition systems are used to interpret user commands and for sound localization. Pressure sensors can be used to track posture in chair and for person localization. Communicating from environment to user is accomplished through visual displays and synthesized speech.

The current state of individual subsystems (e.g., person tracker, speech recognition) can be used to aid other systems by restricting their possible inputs (e.g., where to look or a subset of the vocabulary). Knowledge about the location of a person can be used to restrict and predict a person's next command. If a person were standing by a visual display, that would be a good indication that the user's speech and to where on the display the user is pointing could be coupled together [22], [51], [12]. In the case of a classroom, multi-media records of lectures can be obtained that link simultaneous events such as gesture, audio, and note taking together.

Coupling the input of two sensors is not trivial, as it is for a human who can combine verbal and spatial information effortlessly. In addition to the difficulties of implementing individual sensory systems are the difficulties in integrating these systems. An architecture to integrate sensory systems must also provide an easy manner to add new systems. In [11] the author argues for a distributed architecture, called Scatterbrain, for combining sensory systems based on the subsumption architecture of [6], [7] and on [41]. This is advantageous over a monolithic approach because it permits individual systems to be relatively independent of each other and because integrating new systems is easier. With interacting sensory systems, an environment can act as a personal assistant would. For example, if the location of a city was requested by a user, a smart room could check if any other information about that city was stored in its knowledge base and ask the user if it should be supplied as well.

In adtlition to making an entire room or house intelligent, much research is being done on giving intelligence to typically non-computational objects such as clothing [51], [39] and chairs [62], [30], [15]. Many of the same sensory systems found in smart rooms have been embedded in clothing and cyborg-like augmentation. These include vision and speech recognition systems. A system to indicate the presence of a person in a car seat for situation appropriate airbag deployment was developed in [36]. In [62], a real-time system to classify sitting postures based on pressure distributions was developed. Work in virtual reality-like environments includes chairs with dome-like helmets onto which images can be projected, interaction through either a joystick, keyboard, or touch screen [30]. Airline seats containing air chambers that inflate and deflate have been developed in [15]. A system that could automatically detect pressure peaks would be ideal to regulate the air chambers.

1.3 Pressure Sensing

The design of products, such as seats, is a lengthy and expensive process. It iterates through prototyping and evaluation, the latter of which is often in the form of subjective assessments by test subjects. There is now way to do long term systematic evaluation of seat comfort. Subject testing is inadequate due to the lack of nerves in and around the thighs and buttocks areas. Comfort is perceived more as a lack of discomfort. This has led to the development of pressure sensing devices to measure pressure distributions in products that support the body (e.g., seats, shoes, and beds).

Some early pressure sensing methods involved the use of pressure sensitive inks, thermographs, mechanical springs, and capacitance bridges [29], [13]. Recently there has been commercial developments in pressure sensing devices [34], [63]. Application of these devices include the measurement of: (1) bite pressure (with U-shaped sensors), (2) shoe pressure [5], [14], [43], (3) pressure from clinical support surfaces (e.g., hospital beds) [3], and (4) sitting pressure distributions. It is the understanding and interpretation of sitting pressure distributions to which this work is devoted.

Current work in the use of pressure sensing devices to measure sitting pressure distributions is focused on relating these distributions to seat comfort [34], [58] [64], [65], [3], [13], [33], [55]. The exception to this is the work in [62] which determines sitting posture (e.g., sitting upright, leaning forward) from sitting pressure distribution data and is the basis of this work..

1.4   Anatomy

This section is meant to give the reader a bare bones understanding of anatomical structure and its relation to sitting pressure distributions. Covered are anatomical terminology, the pelvis, the spine, and the legs. A glossary of terms can be found in Appendix B.

Starting with terminology, the three fundamental planes of reference are used to depict the structural arrangement of organs, bones, etc, in the body. Figure 1.2 depicts the commonly used coronal, midsagittal. and *transverse* planes. Directional terms such as anterior (toward the front) vs. posterior (towards the back), medial (toward the midline of the body) vs.

lateral (toward the side of the body), and *superior* (toward the head) vs. *inferior* (away from the head) are used to locate and relate different parts of the body.
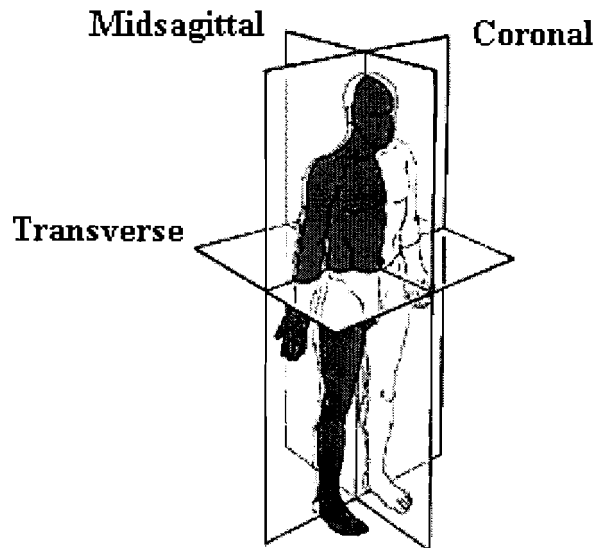
Fig. 1.2. Three fundamental planes.

Two key elements involved with sitting are the pelvis and spine. One important function of the pelvis is that it transmits weight from the upper body to the ischial tuberosities while seated. The ischial tuberosities are commonly referred to as the sitting bones. Figure 1.3 depicts the human pelvis. The ischial tuberosities are the two bottom-most protrusions of the pelvis.

Figure 1.4 shows the natural curve of the spine. The lumbar region of the spine is of particular importance in sitting posture. Proper support of this region while seated to maintain the lumbar curve, sometime referred to as lumbar lordosis, is necessary for both comfort and health [27].

The vertebrae of the spine are connected together by intervertebral discs (Figure 1.5). These are the cartilage discs that separate the vertebrae and provide flexibility and cushioning for the spine. Improper posture can place stress on these discs which, over time, can cause serious damage to them and the surrounding joints and nerves.

Another key component in sitting posture are the legs. Figure 1.6 shows the major bones in the leg. While seated, pressure is distributed along the length of the thigh.

Fig. 1.3. An x-ray of the pelvic area.
(http://www.scar.rad.washington.edu/RadAnatoy/Pelvis/Pelvis.html)
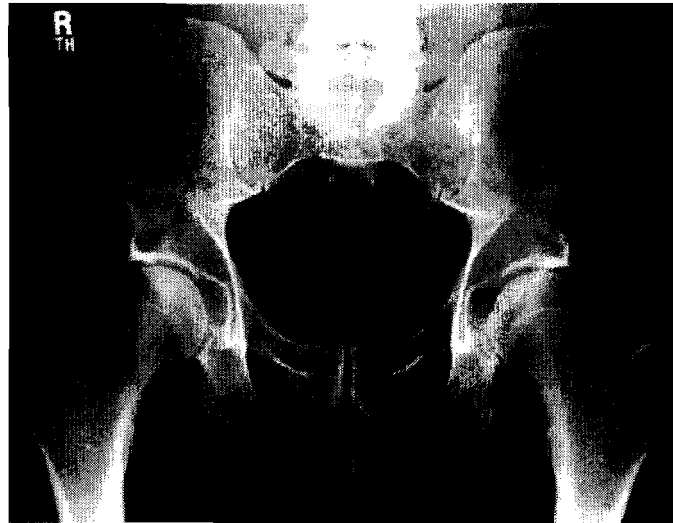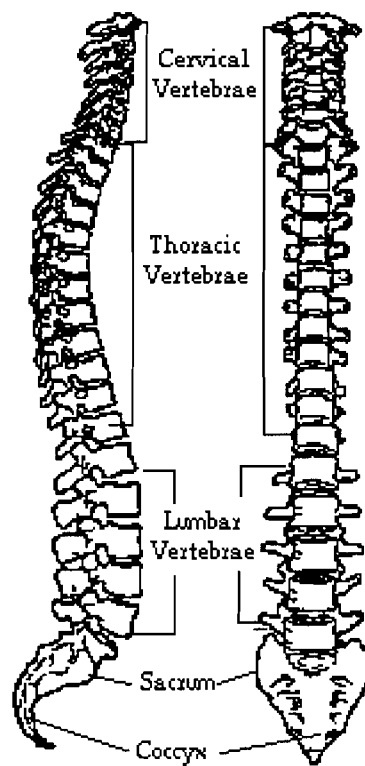
Fig. 1.4. The normal, healthy curve of the spine.
(http://www.mc.maricopa.edu/academic/cult_sci/anthro/origins/webanatomy/spine.html)
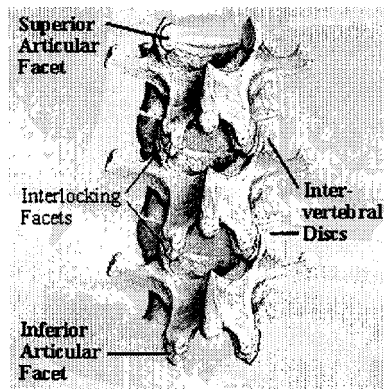
Fig. 1.5. A section of the spine showing the intervertebral discs. (www.scoi.com)

Fig. 1.6. The leg. (http://www.mc.maricopa.edu/anthro/origins/webanatomy/leg.html)
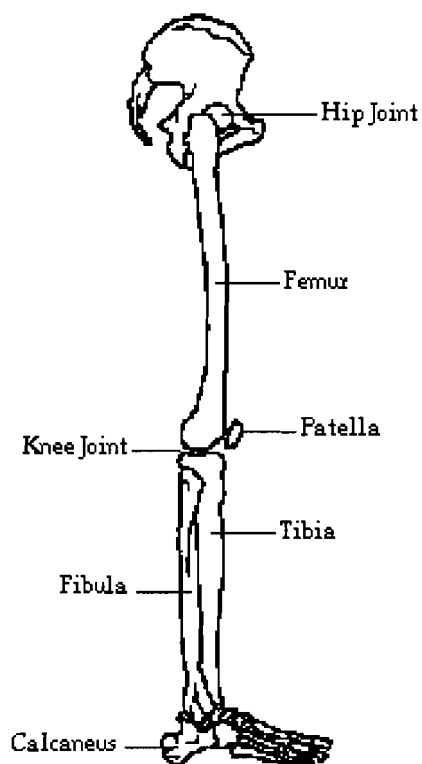
This discussion, the figures in particular, has been centered around the skeletal system. Also of importance to sitting posture is the muscular system. When the body is in an unstable position, a position in which it is not balanced over the pelvis such as leaning forward, muscles contract to keep the body from falling over. Though a person appears still, their muscles may be hard at work. The effects of this will be discussed further in the section in this chapter on Sitting Postures at Work.

1.5   Ergonomics and Sitting

In recent years there has been an increase in *cumulatiue trauma disorders* associated with working environments. These are disorders of the muscles, tendons, ligaments, and/or nerves caused or aggravated by repetitive motion activity that applies stress to the body. This has precipitated the need to design work environments that are healthy for people.

Here vie discuss the role of ergonomics as it applies to seating. Ergonomics deals with the relationships between workers and their environments. Section 1.5.1 contains a discussion on anthropometrics. Section 1.5.2 describes proper chair fit and the results of an incorrectly fitted chair. The final section on ergonomics is Section 1.5.3 which discusses sitting postures in a work environment.

1.5.1   Anthropometrics

The Oxford English Dictionary defines *anthropometry* as the measurement of the human body with a view to determine its average dimensions, and the proportions of its parts, at different ages and in different races or classes. Examples of anthropometric measurements include *popliteal height* (The vertical distance between the floor and the crease just behind the knee of a seated person.), *seat depth* (the distance from the buttocks to the back of the knee), and *elbow rest height* (the height of the elbow above the surface on which a person sits when the torso and thighs form a right angle) [35].

Ergonornicists use anthropometric measurements to designs product:; to fit a desired range of people. This range is usually inclusive of a 5th percentile female and 95th percentile male. (This assumes that all females below the 5th percentile have dimensions that are smaller than those of the 1st percentile of males and all males above the 93th percentile

have dimensions that are larger than those of the 100th percentile female). A person that falls into the 50th percentile is an *average-sized* person. By designing to fit all people in the 5th-to-95th percentile range it is hoped that the product will fit 95 percent of the population. But this may not be attainable.

All males above the 95th percentile for one measurement are not necessarily the same as those who fall above the 95th percentile for a different measurement. Much of the anthropometric data available to ergonomicists comes from military populations. The entire population has greater amounts of extremely small and extremely large people. This results in the effective percentile range being smaller than the desired range.

Ways to compensate for this reduced range are to make products adjustable to differently sized users and to offer the same product in different sizes. These steps rnake the product better fitted to an individual.

1.5.2 Chair Fit

Office workers spend a large percentage of their time at work seated in their chairs. A chair that adjusts to a person's size and does not restrict a person as they move from one task to another is important for a healthy and productive work environment. The chair we use in our work, the Herman Miller Aeron chair, comes in three sizes to better fit people on the small and large ends of the size spectrum.

A chair's height should be adjustable to fit a wide range of people. Ideally, the chair height (as measured from floor to the top of the seatpan) should be close to the popliteal height of the person sitting in it. This is to distribute pressure from the seatpan over the entire back of the thigh. A chair that is adjusted too low would increase the pressure under the ischial tuberosities. A chair that is too high would increase the pressure under the distal thigh. Both of these, over time, would constrict blood flow in the area of increased pressure and make the person uncomfortable, and most likely decrease their productivity [38].

Other common chair adjustments include seatpan angle, seat back tilt.,and lumber support. The later is important to maintain the lumbar curve.

1.5.3 Sitting Postures at Work

A variety of tasks induce a variety of sitting postures at work. People lean forward when performing work on their desks such as writing. They sit up when typing on a keyboard. When their phone rings, they lean to the side to answer it. They lean back and slouch when they want to rest. And they cross their legs when talking with a colleague.

Changing posture, whether due to a change in the task at hand or simply a need to break from the task, has benefits. Alternatingly contracting and relaxing muscles increase circulation. Movement also allows the spine to be nourished. This is because there are no blood vessels in the spine. The only way to provide nutrients is by moving the fluid around the spine through body movement. Joints and ligaments also benefit therapeutically from joint movement. Continual change in sitting posture prevents maintaining awkward or non-neutral positions of the spine. prolonged compression forces on the discs and localized contact stresses.

THIS PAGE INTENTIONALLY LEFT BLANK

2. PRELIMINARY RESULTS

This chapter describes the preliminary results that have been obtained so far on the *Sensing Chair* project. The overall goal of this project is to develop a robust real-time system for tracking a person's sitting posture using surface-mounted pressure distribution sensors on the seatpan and the back of the chair. This work is divided into two parts:

- A real-time classification system for labeling a set of preselected typical sitting postures in an office environment (i.e., Static Posture Classification).

- A real-time tracking system that labels a person's sitting posture at any given time, whether it is a transitional or steady-state posture (i.e., Dynamic Posture Tracking).

The preliminary work focused on the construction of the Static Posture Classification system. First, a Static Posture Database was established. The database contains a total of 1500 sitting pressure distribution maps from 30 individuals (half male and half female) who contributed 5 samples for each of 10 preselected postures. Section 2.2 describes the postures that were selected, the anthropometrics of the subjects, the procedures used for data collection, and the data manipulation. Second, a real-time multi-user Static Posture Classification system was developed. This system classifies sitting pressure distributions from individuals who either contributed samples to the database, or those whose anthropometry is represented in the database, with an overall accuracy of 96% correct. Section 2.3 presents the implementation of the Static Posture Classification system, and its evaluation. The hardware system for capturing pressure distribution maps is briefly described in Section 2.1.

2.1 Sitting Pressure Distribution Measurement System

The sensing system used in the Sensing Chair Project is the Body Pressure Measurement System (BPMS) (Tekscan, Inc, South Boston,MA). It consists of two identical surface-mounted pressure-sensitive transducer sheets, their interface electronics, and a PC interface board. Tekscan also supplies Windows software that captures and displays pressure maps from the sensor sheets.

Each sensor sheet contains a flexible printed circuit array of 42x48 pressure sensing elements (sensels). The sensels are uniformly spaced 10 mm apart. The overall effective sensing area is 41 x 47 cm. For the Sensing Chair Project, the two sensor sheets have been mounted on the seat back and seatpan of a Herman Miller Aeron chair (Figure 2.1).

Fig. 2.1. The Aeron Chair. (http://www.hermanmiller.com)

Each sensel acts as a variable resistor. Its resistance is determined by the normal force being applied to its location. When unloaded, its resistance is high. As an applied force is increased. its resistance decreases. The resistance is converted to an 8-bit digital value. The interface electronics and PC interface board can capture the two pressure distribution maps at rates up to 127 Hz. The pressure maps can be visualized as a greyscale image. Figure 2.2 is an example of the pressure maps captured with a person seated upright in the chair.

The irnage shown in Figure 2.2 is subject to noise due to two sources: inherent *Sensor Noise*, and Sensor Sheet *Deformation.* Sensor noise can be seen as the local abrupt changes in greyscale values. Sensor sheet deformation introduces pressure artifacts into the sitting pressure distribution map that are the result of the sensors bending around and conforming to the chair. The pressure sensors in the Body Pressure Measurement System were designed

Fig. 2.2. An example of sitting pressure distribution maps displayed as an 8-bit greyscale image. The top half of the image shows the pressure distribution on the back of the chair, and the bottom half shows that of the seatpan. The top, bottom, left and right sides of the image correspond to the shoulder area, knee area, right side and left side of the person, respectively.

to be placed on firm flat surfaces. The Aeron chair is contoured to fit the human body. To affix the pressure sensors to the chair, their corners and edges have been wrapped around the edges of the chair. This causes pressure artifacts to appear in the sitting pressure distribution maps (e.g., see the small pressure areas in the upper-left and upper-right corners of Figure 2.2). Removal of sensor noise and pressure artifacts is performed by a process called *cleaning* and will be discussed in detail in Section 2.2.5.

2.2   Static Posture Database

There is no known publicly accessible database of sitting pressure distribution data. We have therefore collected a small database containing Static Sitting Pressure Distribution Maps. It provides the necessary training data for the development of a Static Posture Classification system, as well as data needed for the evaluation of the clasification system. The collection of a database containing *dynamic* sitting pressure distribution data will be discussed in Chapter 3.

Section 2.2.1 describes the set of postures we have chosen for the Static Posture Database. That is followed by a description of the software used to collect the pressure distribution data. Section 2.2.3 provides information on the subjects from whom data was collected. The procedure for data collection is given in Section 2.2.4. Next is a section on the preprocessing of the raw sitting pressure maps to remove sensor noise and sensor sheet deformation pressure artifacts. Section 2.2.6 describes the sitting pressure distribution maps for all postures. Section 2.2.7 focuses on feature extraction of sitting pressure distribution maps. Finally, Section 2 2.8 summarizes data clustering results using the K-Means algorithm.

2.2.1   Postures

The postures contained in the Static Posture Database are *Upright, Leaning Forward, Leaning Left, Leaning Right, Right Leg Crossed, Left Leg Crossed, Leaning Left with Right Leg Crossed, Leanzng Right with Left Leg Crossed, Leaning Back, and Slouching.* These postures are representative of the typical sitting postures that can be found in an office environment [38].

What follows is a general description of each posture. The posture names are self-explanatory so skipping this section should not affect the understanding of the rest of this document. The reader should look at Table 2.1 to become familiar with the abbreviations used for the posture names as postures will commonly be referred to by their abbreviations.

For posture *Upright* a person is sitting comfortably upright in the chair with both feet flat on the floor. Their hands and forearms may either be on the lap or on the armrests. The back may rest against the back of the chair but does not push against it. In the *Leaning Forward* posture a person's trunk is angled forward from the waist. There is usually no pressure applied on the backrest of the chair except for a small pressure area near the lumbar region. In posture *Leaning Left* a person has their weight centered over the left sitting bone. Placing the left arm on the armrest is optional for the person. In *Leaning Right,* a person's weight is centered over the right sitting bone. Placing the arm on the armrest is again optional.

Posture *Right Leg Crossed* is when the right leg is crossed on top of the left leg. Usually, for women' the right knee is over the left knee and for men, the right ankle is on the left knee. Posture *Left Leg Crossed* has the left leg crossed on top of the right; leg. Usually, for women, the left knee is over the right knee and for men, the left ankle is on the right knee. For posture *Leaning Left with Right Leg Crossed* the person's weight is over the left sitting bone while having their right leg crossed. The arm may rest on the left armrest. In posture *Leaning Right Left Leg Crossed* the person's weight is over the right sitting bone while their left leg is crossed over their right leg. The arm may rest on the right armrest.

In the: *Leaning Back* posture, a person's upper torso presses against the back of the chair. Posture *Slouching* is when a person's pelvis is positioned toward the front edge of the seatpan. For brevity, throughout the rest of this document the postures will commonly be referred to by their abbreviations listed in Table 2.1.

2.2.2 Static Posture Acquisition Software

This section describes the software developed to collect sitting pressure clistribution maps. This software was needed to obtain data for the Static Posture Database. A program was written in Microsoft Visual C++ 6.0 to run under Windows 98. It uses an API library

Posture Name	Abbreviation
Upright	N
Leaning Forward	LNF
Leaning Left	LNL
Leaning Right	LNR
Right Leg Crossed	RLC
Left Leg Crossed	LLC
Leaning Left with Right Leg Crossed	LLRLC
Leaning Right with Left Leg Crossed	LRLLC
Leaning Back	LNB
Slouching	SL

Table 2.1 Posture names and their abbreviations.

supplied by Tekscan that permits direct access to the BMPS interface board. This program, called Sitting Posture Acquisition, automatically records and stores sitting pressure distribution maps for the Static Posture Database.

Figure 2.3 shows the main window of the *Sitting Posture Acquisition* program. The experimenter is able either to enter information of the subject or to collect data from the subject by clicking on the desired button.
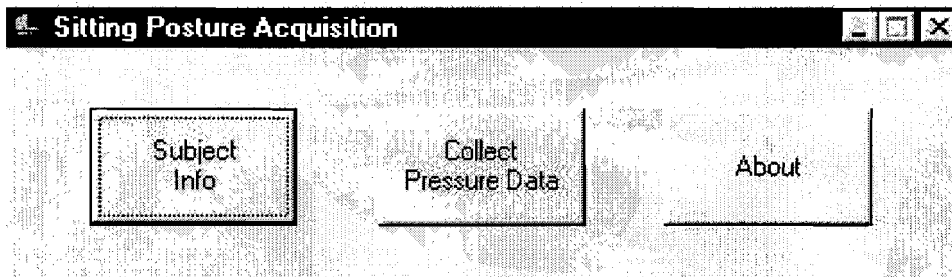
Fig. 2.3. Static Posture Acquisition program.

The *Subject Info* button opens a dialog box in which the experimenter can enter the following information about the subject: (1) a unique identification name (for anonymity of the subject), (2) subject's height, (3) subject's weight, (4) subject's age, (5) subject's gender, (6) height of the chair[1], and (7) comments by the experimenter (e.g., how the subject crosses his or her legs). For female subjects the unique identification name (Subject ID) is $F\#$, where $\#$ is a number. For male subjects the Subject ID is $M\#$. The subject information is saved in a text file with the name of the unique Subject ID of the subject.

The *Collect Pressure Data* button opens a dialog box that assists the experimenter by prompting her with the posture name for the nest sample of pressure distribution maps. Steps involved in the data collection process will be explained in more detail in Section 2.2.4.

2.2.3 Subjects

A total of 30 subjects (15 females and 15 males) participated in data collection. Five samples for each of the ten postures were collected from each subject. Therefore, a total of 150 samples were collected for each posture. Subjects were selected on the basis of their overall size. The goal was to obtain subjects with a wide distribution of weight and height.

---

[1]see Section 2.2.4 for the importance of chair height

The range, mean and standard deviation for subject's height, weight, and age, chair height, and subjective assessment of paddedness have been computed for all of the subjects as well as for the female and male subjects separately (see Table 2.2). Paddedness is a subjective assessment of the subject's build by the experimenter to describe the fitness in the torso and upper leg region of the subject. The range of paddedness is 1 to 3, where 1 is not *paddeded* and 3 is *well padded.* Subjects that are well padded tend to be overweight while those that aren't edge towards being skinny and/or muscular.

Subject by Group	Statistic	Subject Height (in.)	Subject Weight (lbs.)	Subject Age (years)	Chair Height (in.)	Paddedness
Female	Range	60-70	100-185	18-60	0-3.125	1-2
	Mean	65.867	139.067	30.300	0.833	1.333
	StdDev	2.475	21.171	14.446	1.082	0.488
Male	Range	66-75	146-260	19-37	0-4.75	1-3
	Mean	70.600	177.400	28.133	2.476	1.60
	StdDev	2.667	33.032	6.334	1.220	0.828
All	Range	60-75	100-260	18-60	0-4.75	1-3
	Mean	68.233	158.233	29.267	1.655	1.467
	StdDev	3.491	33.513	11.020	1.408	0.681

Table 2.2 Database statistics computed for female subjects, male subjects and all subjects.

A plot of the subjects' height vs. weight can be seen in Figure 2.4. Male subjects' data are displayed with an '.' and female subjects' data are displayed with a 'x'. The plot shows that we succeeded in collecting data for subjects over a wide range of height and weight.

2.2.4 Procedure

Since the building of the Static Posture Database took place over several weeks, a collection procedure was developed to standardize the collection of data. The steps in this procedure are outlined below.
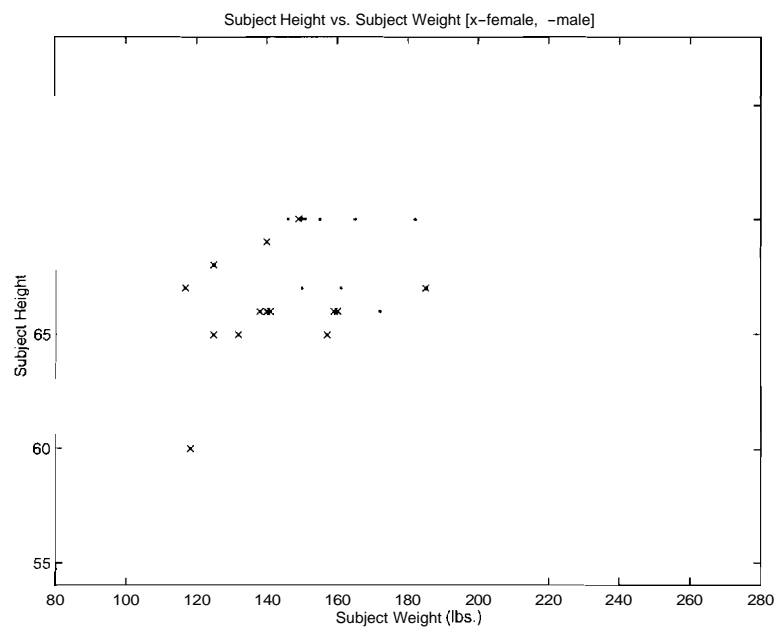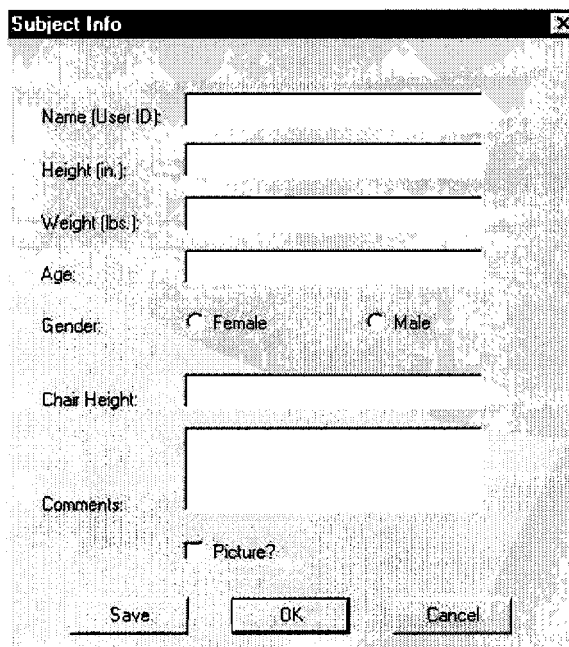
Fig. 2.4. The distribution of subject height and weight.

1. The subject is given a brief description of the Static Posture Classification Project and informed of the types of postures and number of samples of each posture that will be collected.

2. The subject is then familiarized with sitting pressure distribution maps by viewing their sitting pressure distribution data on a computer monitor.

3. The chair height is adjusted to properly fit the subject. The height of a chair can dramatically change the pressure distribution pattern on (especially) the seatpan. When the chair is too low or too high, a person's weight is mostly supported by the back or the front edge of the seatpan, respectively. The experimenter adjusts the height of the chair such that pressure seems to be evenly distributed across the subject's thighs.

4. Data collection begins by opening the Sitting Posture Acquisition program (see Figure 2.3). The experimenter clicks on the *Subject Info* button and enters the relevant information in the dialog box shown in Figure 2.5.

5. To collect data, the experimenter clicks on the *Collect Pressure Data* button. The *Data Collection* dialog box (Figure 2.6a) opens and the experimenter click:; the *Start* button to begin the collection process (see Figure 2.3).

6. The program prompts the subject to sit in a specified posture. When the subjects complies, the experimenter clicks the *Capture* button to capture the current sitting pressure distribution map and then the *Save Map* button to save the sitting pressure distribution map to a file (Figure 2.6 b and c).

   When a sitting pressure distribution map is saved, the Sitting Posture Acquisition program saves the pressure map in an ASCII file with a name that is indicative of the subject, posture and sample number. All data files for a given subject are stored in one directory. The directory name is taken as the value of the Subject ID field of the Subject Info dialog box. The naming convention of the data files is as follows:

   $< SubjectID >< PostureAbbreviation > . < SampleNumber >$

Fig. 2.5. The Subject Info dialog box.
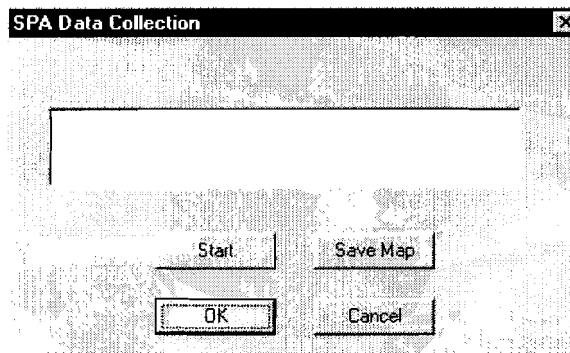
(a)
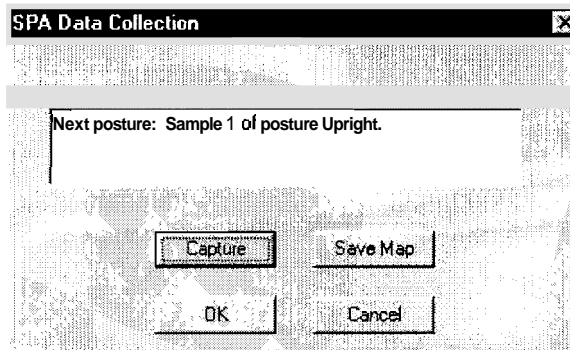
(b)
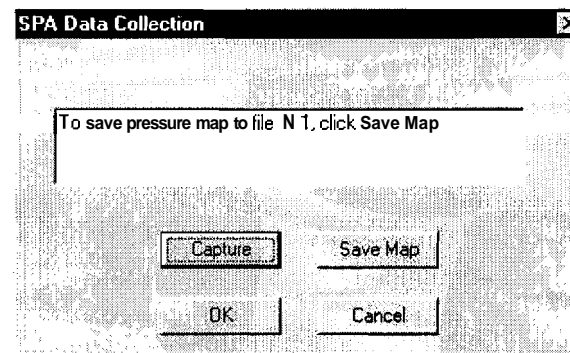
(c)

Fig. 2.6. (a) Click Start to begin data collection. (b) The data acquisition program prompts for a posture. (c) The current sitting pressure distribution map is stored in a file after the Save Map button is clicked.

For example, the five data files for posture Right Leg Crossed for subject F04 would be given the names F04RLC.1, F04RLC.2, F04RLC.3, F04RLC.4, F04RLC.5.

7. Step 6 is repeated until all samples for all postures have been collected.

2.2.5 **Data** Preprocessing

As described at the beginning of Section 2.1, there are two types of noise found in the sitting pressure distribution maps due to either *sensor noise* or *sensor sheet deformation.* This section describes the methods used to clean the sitting pressure distribution maps. Cleaning the sitting pressure distribution maps involves smoothing the raw sitting pressure distribution maps to remove sensor noise, and removing pressure artifacts from the smoothed maps.

Figure 2.7 shows a raw sitting pressure distribution map for posture N. This is the first sample of posture *Upright* collected for female subject number four (F04N.1) The image is shown as a 3-D height map, where the height above the z=0 plane indicates the pressure value. Large height values indicate high pressure values in the sitting pressure distribution map.
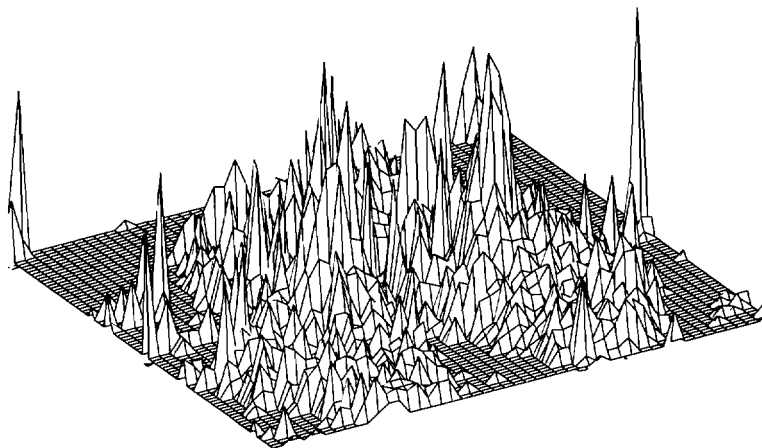
Fig. 2.7. Raw sitting pressure distribution map for posture Upright (sample F04N.1).

As one can see, the raw sitting pressure distribution map is extremely noisy. The first step in data preprocessing is to smooth the sitting pressure distribution map. The 3x3 smoothing kernel shown below is applied to the sitting pressure distribution map.

$$\frac{1}{7} \begin{bmatrix} 0.5 & 1.0 & 0.5 \\ 1.0 & 1.0 & 1.0 \\ 0.5 & 1.0 & 0.5 \end{bmatrix} \tag{2.1}$$

The dramatic affects of applying the smoothing operator to the pressure map in Figure 2.7 can be seen in Figure 2.8. This does not remove all of the noise in the sitting pressure distribution map. In addition to noisy pressure values, there are pressure artifacts throughout the sitting pressure distribution map.
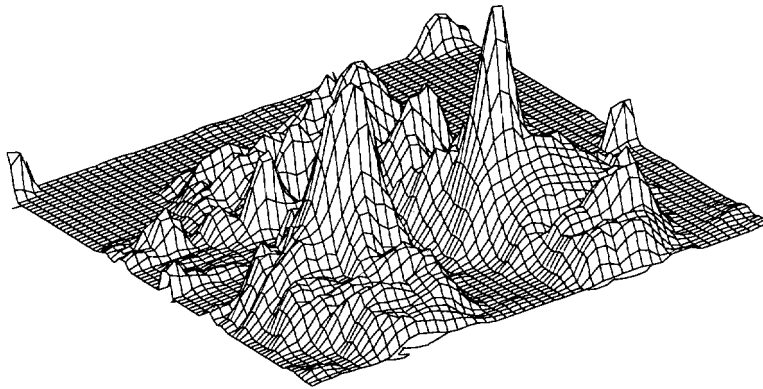
Fig. 2.8. Sample F04N.1 after smoothing.

The pressure artifacts due to Pressure Sheet Deformation occupy significantly smaller areas than any of the subject induced pressure components in the sitting pressure distribution map. They usually lie near the corners and along the edges of the sitting pressure distribution map where the sensors fold around the chair. These artifacts are removed in a two-step process. First, the connected components in the sitting pressure distribution map are found

by a standard component labeling algorithm [57]. Second, those components whose areas are smaller than ten percent of the largest component area of the sitting pressure distribution map are deemed to be pressure artifacts and are removed. The cleaned version of the sample upright posture is shown in Figure 2.9.
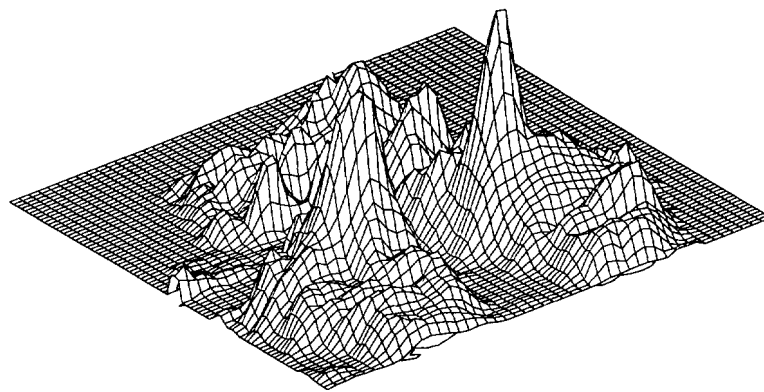
Fig. 2.9. Sample F04N.1 after artifact removal.

2.2.6 Description of Sitting Pressure Distribution Maps

What follows is a general description of the sitting pressure distribution maps after they have been cleaned, that is after smoothing and pressure artifact removal. For all postures, there are usually peaks in the sitting pressure distribution map that correspond to the the sitting bones of the pelvis. Most sitting pressure distribution maps show two distinct peaks, one for each sitting bone. Occasionally, the peaks are not distinct as compared to the pressure values around the posterior distal thigh. This case generally occurs with more padded subjects. Figure 2.10 shows a sitting pressure distribution map for a padded subject.

Sometimes there are secondary peaks located down the thigh and to the outside of the leg to peaks from the sitting bone. These are believed to result from a pressure increase due to the bony bulge at the end of the femur near the hip joint called the greater trochanter. which is distal and lateral to the sitting bones. Figure 2.11 shows an x-ray of the pelvic area
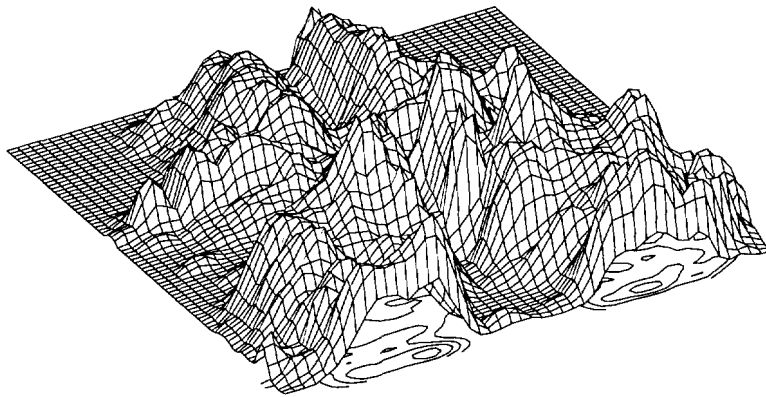
Fig. 2.10. A sitting pressure distribution map of a *padded* person.

with the greater trochanters labeled. Figure 2.12 shows a typical sitting pressure distribution map for posture N compared to one that has secondary peaks.
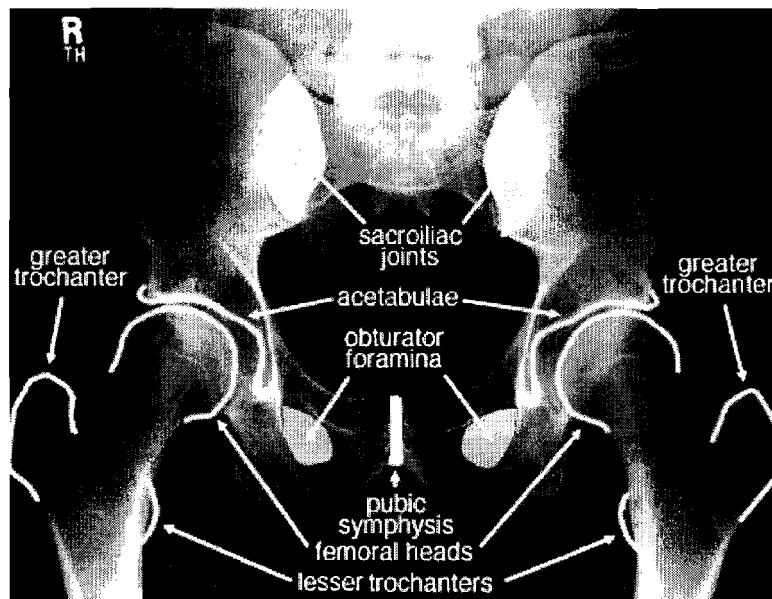
Fig. 2.11. X-ray of a pelvis. The greater trochanters are labeled on each femur. (http://www.scar.rad.washington.edu/RadAnatomy/Pelvis/PelvisLabelled.html)
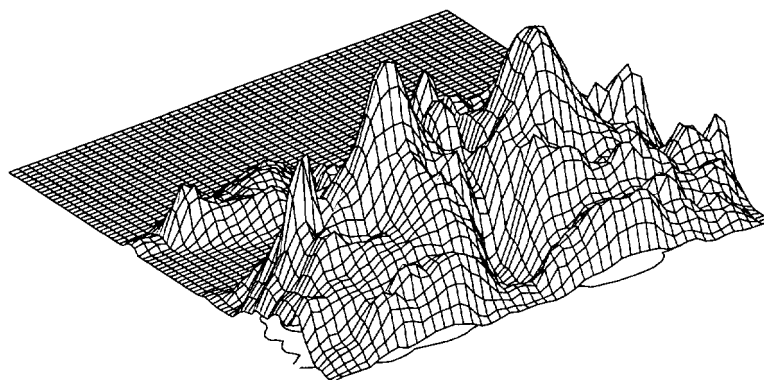
Fig. 2.12. A sitting pressure distribution map showing secondary peaks.

The generalities mentioned so far are applicable to all postures. Below are descriptions of the sitting pressure distribution maps for each posture. Descriptions for both the seatpan and the seat back are given.

N Posture *Upright*: Figure 2.13a

Seatpan: Usually one connected component where areas of the pressure distribution corresponding to the left and right thighs diverge. It is possible to have two separate components.

Seatback: One trapezoidal or butterfly-like region is common, usua'lly with two vertically symmetric lobes. There can be two components separated by the area on the chair over which the spine is positioned.

LNF Posture *Leaning* Forward: Figure 2.13b

Seatpan: Similar to N but pressure is more evenly distributed across the length of the leg.

Seatback: There is very little pressure on the seat back.

LNL Posture *Leaning* Forward: Figure 2.13c

Seatpan: There is an increase in pressure and area on the left leg and a decrease on the right leg as compared to posture N.

Seatback: Pressure shift to the left side. Back pressure region can resemble that of N or can show only one of the vertical lobes.

LNR *Posture Leaning Forward*: Figure 2.13d

Similar to LNL but flipped.

RLC *Posture* Right Leg *Crossed*: Figure 2.13e

Seatpan: Left leg is of same shape as that in posture N. The right leg reduces to a sphere with an increase of pressure around the sitting bones. There can also be an increase of pressure around the left knee.

Seatback: Similar to posture N.

LLC Posture Left *Leg Crossed*: Figure 2.13f

Similar to RLC but flipped.

LLRLC *Posture Leaning Left with Right Leg Crossed:* Figure 2.13g

Seatpan: The pressure on left leg is greater than that of posture RLC. There is also a decrease in pressure of right sitting bone.

Seatback: Similar to posture LNL.

LRLLC *Posture Leaning Right with Left Leg Crossed:* Figure 2.13h

Similar to LLRLC but flipped.

LNB *Posture Leaning Back:* Figure 2.13i

Seatpan: Similar in shape to posture N. The pressure of the sitting bone peaks is normally reduced. There is an increase in pressure along the entire leg.

Seatback: There is an increase in pressure compared to posture N. The area of the back component may increase. In the case of a lightweight person, there may not be a great amount of pressure placed on the seat back.

SL *Posture Slouching:* Figure 2.13j

Seatpan: The sitting bones slide forward towards the middle of the chair. The main component looks like an elliptical blob. Any sharp pressure peaks from the sitting bone:; are reduced from those of posture N.

Seatback: The area of the main component decreases. What corresponds to the shoulder area in posture N is located slightly lower down the back of the chair. There is the possibility of the neck and/or head pressing against the back of the chair as either an extension of the main component or as a separate component.

2.2.7   Feature Extraction

Many features can be extracted from the sitting pressure distribution maps. Some of these are the following:

- Ischial tuberosity (sitting hone) localization

- Total force

- Average pressure

Fig. 2.13. Samples of cleaned sitting pressure distribution maps for all postures: (a) N (b) LNF (c) LNL (d) LNR (e) RLC (f) LLC (g) LLRLC (h) LRLLC (i) LNB (j) SL.

- Maximum pressure value

- Number of components

- Areas of the largest components in the pressure map

- Angle of divergence of the legs, major axis, minor axis

Once the pressure peaks that correspond to the sitting bones are known, the distance between them and the orientation of the line connecting them can be calculated. These features can be used to rotate and scale a sitting pressure distribution map. The sitting bone distance for females usually is larger than that for males (F: 15.517 ± 0.550, M: 13.352 ± 1.623). This measurement could be used to predict the gender of a person sitting in a sensing chair.

Sitting bone localization. total force, average pressure, maximum pressure value, number of components, component areas, leg divergence angle, and leg orientations can be automatically calculated for 90 percent of the pressure distribution maps. They will be useful for a feature-based static posture classification or tracking system.

2.2.8   K-Means

K-means is a clustering algorithm that takes a set of N data samples of dimension d and splits it into K clusters. This algorithm was applied to the Static Posture Database to investigate whether the natural clustering of the database samples corresponds to that according to their posture labels[2].

K-means works by randomly initializing K clusters. The means of each of the K clusters are calculated. Then for each sample, its distance to each cluster is computed as the Euclidean distance to each cluster mean and it is reassigned to the cluster corresponding to the smallest distance. The means of the new clusters are computed and the reassignment step is repeated until the clusters do not change [20].

The K-means algorithm was applied to the data in the Static Posture Database under various scenarios. It was first applied to cluster samples into different postures and then to find clusters within a given posture

---

[2]The author thanks Tim Stough and Jennifer Dy for the K-means MATLAB code
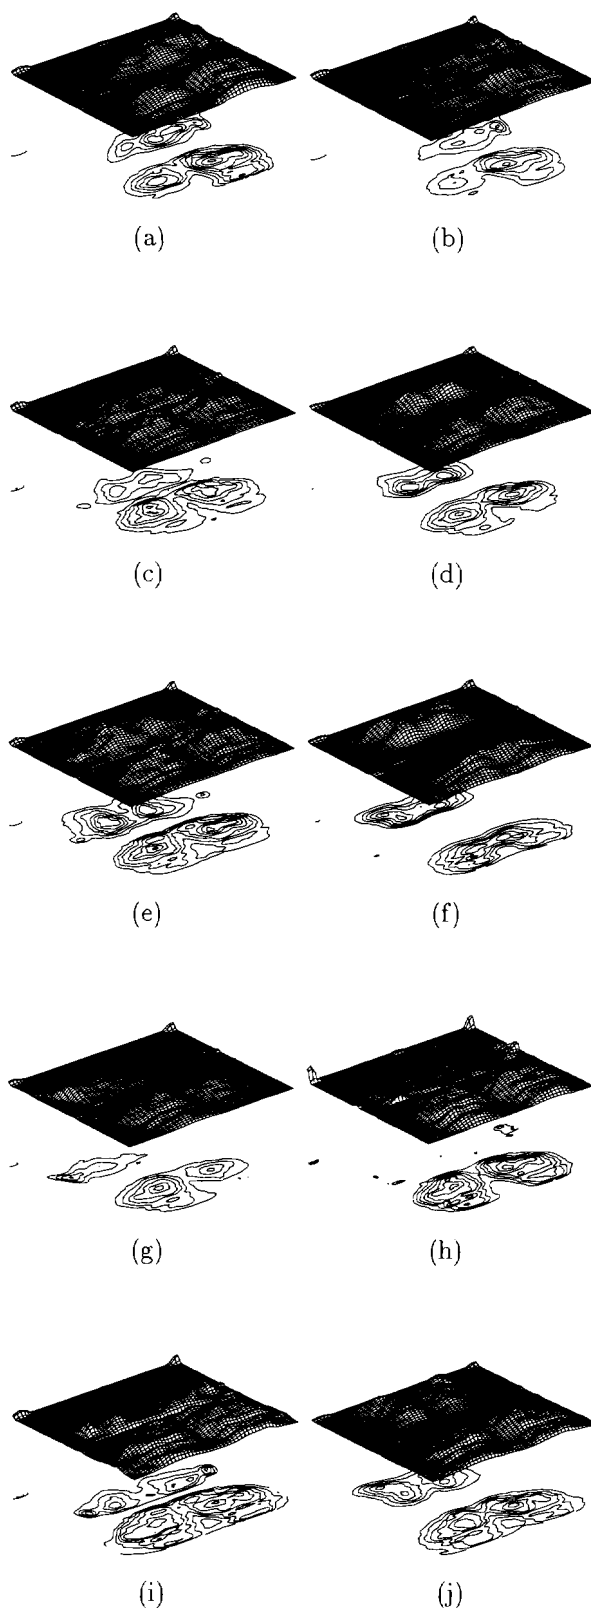
Fig. 2.14. Output from running K-means on the data in the Static Posture Database with K = 10.

To find clusters of samples representing different postures K-means was run on smoothed sitting pressure distribution maps. The number of clusters was varied from 10 to 20, the dimension of the data was 4032, and the number of samples was 1500. Figure 2.14 shows the mean:; it computed for $K = 10$. Whereas some of the cluster means (e.g., Figure 2.14f) are very similar to the posture mean of the database samples (in this case Slouching, see Figure 2.16j), others appear to be averaged from samples with different posture labels. To investigate this further, the number of samples from each posture that contributed to each cluster mean are tallied in Table 2.3. As can be seen, the clusters do not necessarily correspond to individual postures.

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Total
a	0	0	0	34	9	2	68	0	0	0	113
b	0	5	0	77	6	2	60	0	2	0	152
c	37	24	16	4	36	29	0	2	7	0	155
d	31	0	5	5	53	69	8	36	25	4	236
e	41	0	21	2	25	20	6	17	13	1	146
f	0	0	0	0	0	1	1	0	9	111	122
g	5	11	76	0	7	10	0	81	12	0	202
h	7	82	5	2	0	0	0	0	1	0	97
i	24	28	27	26	14	15	7	13	1	0	155
j	5	0	0	0	0	2	0	1	80	34	122

Table 2.3 Number of samples from each posture that belong to each cluster from K-means with K=10. The cluster labels in the table correspond to the image of the cluster mean of Figure 2.16.

Appendix A contains figures of the cluster means and tables of the number of samples from each posture that contributed to each cluster mean for $K = 12, 14, 16, 18$, and 20.

We then investigated clustering the data in the eigenspace. A single eigenspace was trained on the 1500 samples in the Static posture Database. The training samples were then projected onto the first D eigenvectors of the eigenspace to obtain a vector of weights for each training sample. The value of D varied from 5 to 250. The weight vectors were then

clustered using K-means for various values of K. The results of clustering in the eigenspace also produced cluster means that did not visually resemble one of the 10 postures. Looking at the means and standard deviations of the weight vector elements across different postures shows that the means are similar for different postures and that the variances are relatively large. These both imply that the projection weights are not good features by which to cluster.

The second method of clustering was done to examine how by which the data samples within a single posture cluster. The samples for each posture usually form 3 clusters. These clusters represent differences in size and weight of a person and correspond to small, medium, and large sized people.

2.3   Posture Classification

The problem of classification can be defined as follows. Assume, we are given a set of labeled training samples containing several examples of each posture by each of our subjects. The associated label indicates to which class each sample belongs (i.e. the posture that the sample represents). The first step is to select a learning algorithm that can be used to construct some type of model of the data, whether it be appearanced-based or feature-based. Then, given a new unknown or unlabeled sample X_{new}, the problem of classification is to identify to which model (or group or class) that the new sample belongs.

Classification systems fall under one of two categories: appearanced-based and model-based. Appearanced-based systems include those that use techniques such as principal components analysis and discriminant analysis for classification. Model-based systems match scene features to model features in order to classify an object in a scene.

A real-time static posture classification system that uses the appearanced-based technique of principal components analysis (PCA) has been developed. Section 2.3.1 describes PCA and its formulation for the classification of sitting postures. Section 2.3.2 describes the posture data in the new space defined by PCA. The real-time sitting posture classification program is described in Section 2.3.3. Lastly, Section 2.3.4 gives the classification results.

2.3.1 Posture-Based Eigenspaces

Principal components analysis is a well known technique in computer vision. It has been applied to many vision problems with much success [67], [50]. It has also been used for a single user posture classification system [62].

A common approach to PCA consists of projecting the set of training vectors X_i to a single eigenspace. The classification step can be described as follows. First, we project the new vector X_{new} onto the eigenspace. We then use the L2-norm to search for the nearest-neighbor of X_{new}. This nearest-neighbor is the one that infers the class (in our case, posture) to our unknown vector X_{new}. We can reduce the dimensionality of our problem by using those eigenvectors that correspond to the N largest eigenvalues and discarding the rest [67].

Another approach (and the approach we take) would correspond to generating a separate eigenspace for each posture, what we call a posture space (i.e., the view-based eigenspace in [50]). If 10 different postures are to be discriminated, 10 different posture spaces are generated. In this case, the classification step is more complicated than before. One possible way of doing this consists of searching for the best reconstruction of the unknown sample. That is to say, we first project our vector X_{new} onto each of the posture spaces. Second, we reconstruct the new vector from each of the 10 projections. Last, we search for the best reconstruction of the unknown sample. This step is a simple matter of computing the Euclidean distance between the unknown vector X_{new} and each of the reconstructions (X_{new1}, X_{new2}, .. X_{new10}). The label of the reconstruction that is closest to the unknown sample is used to classify X_{new}. Once again, we can limit ourselves to a low dimensional subspace by keeping those eigenvectors that correspond to the largest eigenvalues. We call these eigenvectors *Eigen Pressure Maps* (EPMs).

Eigen-decomposition for the Posture Classificatiori System starts with preprocessing the sitting pressure distribution maps in the Static Posture Database. The smoothing step described in Section 2.2.5 is applied to the raw sitting pressure distribution maps. Smoothing reduces the mean-square error in the eigen representations and increases the overall accuracy in classification by 7% (See Section 2.3.4 for more details). Pressure artifacts are not removed because (1) they are common to all pressure maps, and therefore, do not affect

the performance of the Posture Classification System; (2) the removal of pressure artifacts reduces the effective sampling rate of the whole system.

After the sitting pressure distribution maps are smoothed, they are then normalized. The pressure values on the seatpan are usually much higher than those of the seat back. Therefore, the values in the seat back and seatpan are normalized independently. Figure 2.*15* shows the result of normalizing the cleaned sample posture $F04N.1$ (see Figure 2.9).
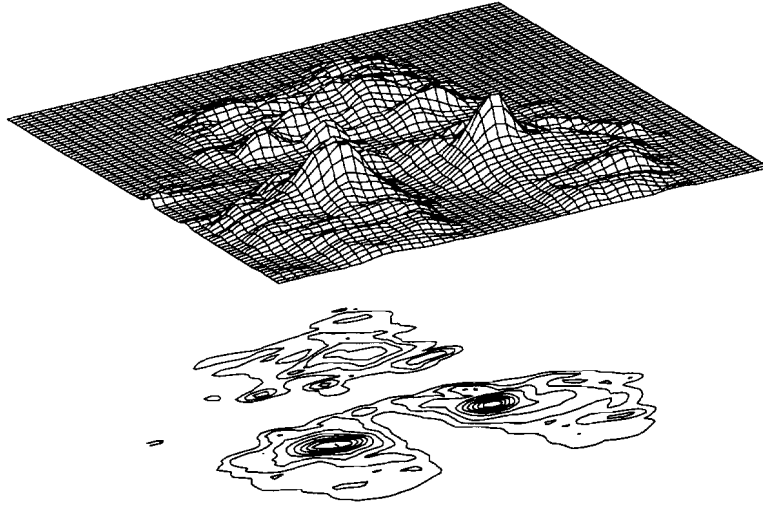
Fig. 2.*15*. Normalized sitting pressure distribution map. Also shown are the equal pressure contours.

The sitting pressure distribution map is raster scanned to form a *4032x1* vector. Eigen-decomposition for each posture k, k = 1..N with $N = 10$, is computed as follows.

Let $P_k = \{X_{1,1}, ...X_{1,S}, X_{2,1}, ...X_{2,S}, ..., X_{H,1}, ...X_{H,S}\}$ be the set of training samples for posture k, where $H = 30$ is the number of subjects. $S = 5$ is the number of samples per posture per subject. There are $HS = 150$ training samples of each posture in the Static Posture Database. For simplicity, the training set P_l, is rewritten as $P_l = \{X_1, X_2, ..., X_M\}$ where $M = 150$. The mean sitting pressure distribution map vector for posture k is calculated as:

$$\bar{X}_k = \frac{1}{M} \sum_{i=1}^{M} X_i \qquad (2.2)$$

The mean sitting pressure distribution map vectors for all postures art: shown in Figure 2.16. The vectors have been converted back to arrays for display purposes.

Zero-mean data is obtained by subtracting the mean from each of the training samples:

$$\phi_i = X_i - \bar{X}_k \tag{2.3}$$

The zero-mean training samples are stacked to form a large matrix Φ, of size 4032x150, whose covariance matrix C is:

$$C = \frac{1}{M} \sum_{i=1}^{M} \phi_i \phi_i' = \frac{1}{M} \Phi \Phi' \tag{2.4}$$

where $\Phi = [\phi_1, \phi_2, ... \phi_M]$ is of size 4032xM.

The eigenvectors and eigenvalues of C determine the eigen-decomposition of the training samples. Seeing that C is such a large matrix, computation of its eigenvectors is computationally expensive. Since $M = 150 < 4032$ no more than M eigenvalues are non-zero. The eigenvectors of our large 4032x4032 matrix $C = \frac{1}{M} \Phi \Phi'$ and those of the MxM matrix $C\prime = \frac{1}{M} \Phi' \Phi$ matrix are related in the following manner [44]:

Let vectors v, of size Mx1 be the eigenvectors of $C\prime = \frac{1}{M} \Phi' \Phi$. Then,

$$\frac{1}{M} \Phi \Phi' v_i = \mu_i v_i \tag{2.5}$$

where μ_i are eigenvalues. Premultiplying by Φ yields:

$$\frac{1}{M} \Phi \Phi' (\Phi v_i) = \mu_i (\Phi v_i) \tag{2.6}$$

So, the eigenvalues and eigenvectors of C are μ_i and Φv_i, respectively.

The eigen-decomposition of the training samples consists of finding the M eigenvectors v, of $C' = \frac{1}{M} \Phi' \Phi$. The eigenvectors of C u_i, of size 4032x1, are given by:

$$u_i = \Phi v_i = \sum_{j=1}^{M} v_{i,j} \phi_i \quad \text{for} \quad i = 1, 2, ... M \tag{2.7}$$

where $v_{i,j}$ is the j-th component of v_i.
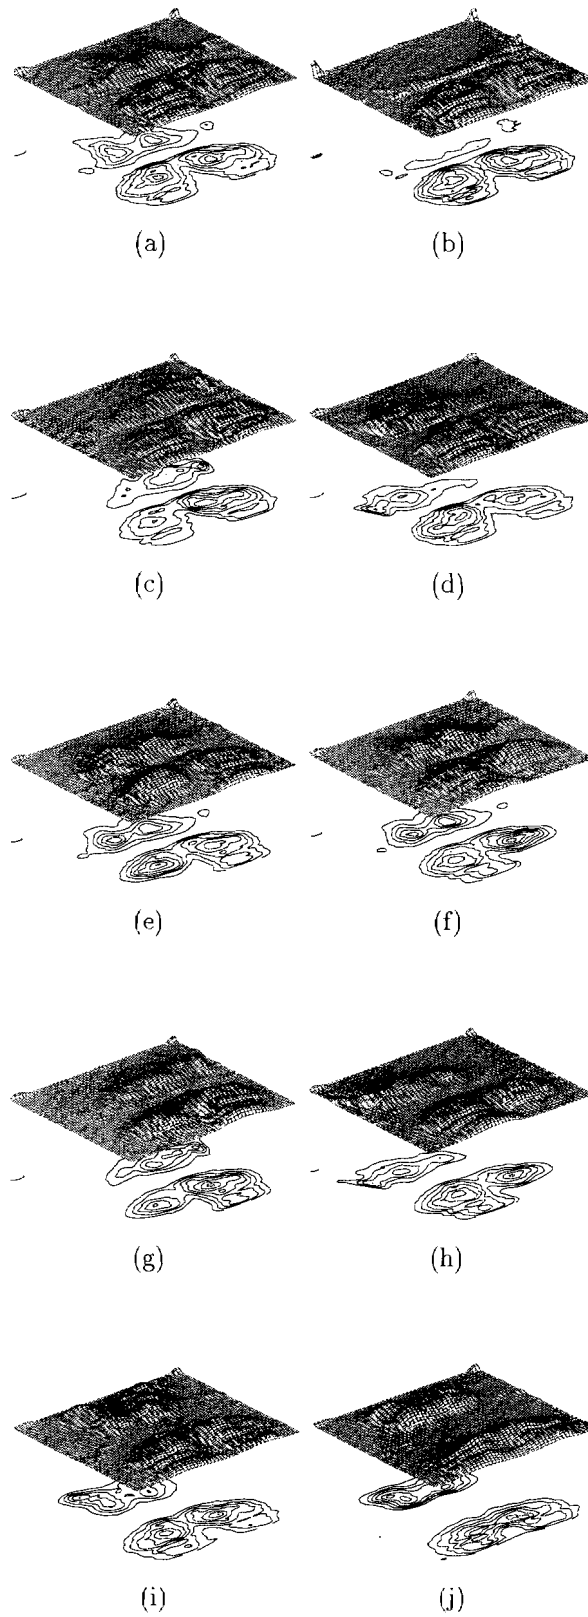
Fig. 2.16. Mean sitting pressure distribution maps: (a) N (b) LNF *(c)* LNL (d) LNR *(e)* RLC (f) LLC (g) LLRLC (h) LRLLC (i) LNB (j) SL.

The u_i (i $= 1, 2, ... M$) are therefore the EPMs. We let e_{ki} denote the ith eigenvector (the eigenvector corresponding to the ith largest eigenvalue) of the posture space for posture k. A training sample can be reconstructed from the EPMs. This is done by projecting the sample onto each EPM to obtain a vector of weights. The EPMs are then added together, weighted by the respected projection weight, to reconstruct the sample (i.e., $\phi_i = \sum_{j=1}^{M}(\phi_i \cdot e_{kj})e_{kj}$). We can choose D $\leq M$ eigenvectors to reduce the dimension of our problem. The error in the reconstruction caused by eliminating eigenvectors is minimized, in the mean-square sense, if we use the D eigenvectors corresponding to the D largest eigenvalues.

The procedure for obtaining the EPMs described above is repeated for each of the 10 postures. Figure 2.17 shows the eigenvalues calculated for each of the posture spaces.

Classification of a test sitting pressure distribution map is performed as follows. First, the test map undergoes the same preprocessing steps the training data has undergone (smoothing and normalization). The following steps are repeated for each posture space k. The posture mean $\bar{X}_k$ of the training samples for posture k is subtracted from the test map X_t.

$$\phi_t = X_t - \bar{X}_k \tag{2.8}$$

The mean-subtract test sample ϕ_t is projected onto the first D eigenvectors of posture space k to obtain a D-dimensional weight vector $W_k = [w_{k1}...w_{kD}]^T$, where the ith element of W_k is the projection of ϕ_t onto the ith eigenvector of posture space k. The test map is then reconstructed as follows:

$$\hat{\phi}_t = \sum_{i=1}^{D} w_{ki}e_{ki} \tag{2.9}$$

where w_{ki} is the ith element of W_k and e_{ki} is the ith eigenvector of posture space k

The distance $\phi_t - \hat{\phi}_t$ (i.e., the Distance From Feature Space - DFFS [67]) is used as a distance measure between the test map ϕ_t and posture space k. The posture space yielding the smallest DFFS in the reconstruction is taken as the posture label for classification. If $\min_k$ DFF'S $\geq$ threshold then the sitting pressure distribution map is labeled as *unknown posture*. The threshold is determined empirically

All 150 eigenvalues for each of 10 postures

(a)

First 10 eigenvalues for each of 10 postures
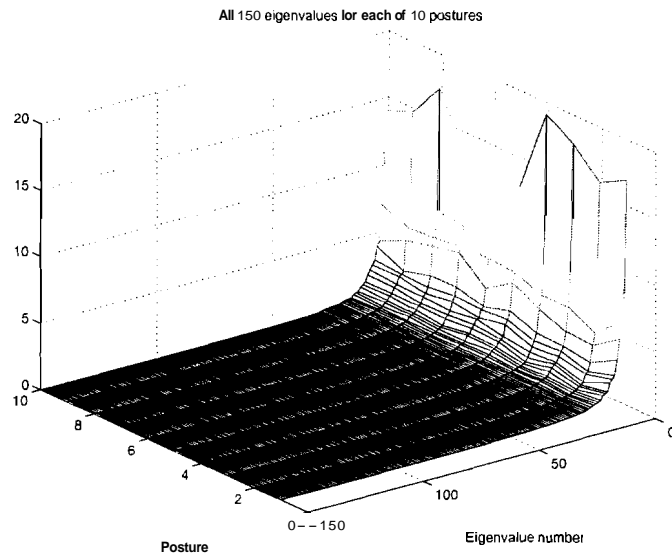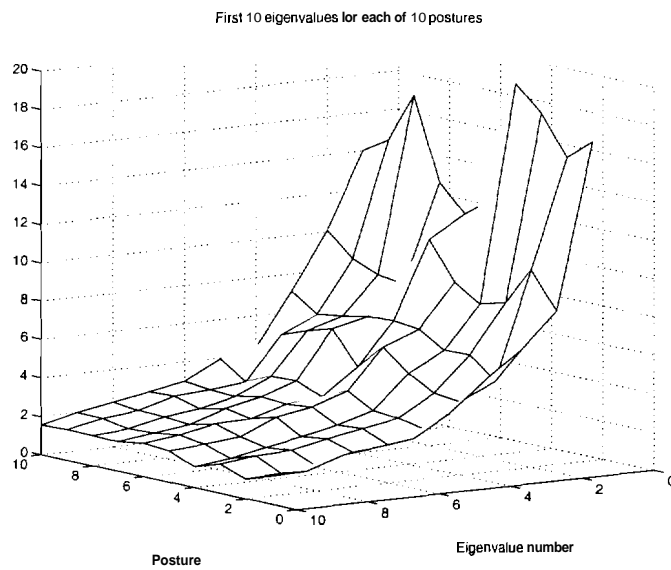
(b)

Fig. 2.17. Eigenvalues for all ten postures. (a) All 150 eigenvalues. (b) The ten largest eigenvalues for each posture.

2.3.2 Data Visualization

Since the sitting pressure distribution samples are of high dimension (4032) it is hard to visualize the distribution of training samples in that posture space. Figure 2.18 shows the projections of all of the training samples for each posture onto the three eigenvectors (eigen pressure maps) corresponding to the three largest eigenvalues. This figure shows that even in just three dimensions, the posture spaces are complex.

2.3.3 A Real-Time Static Posture Classification System

This section describes the real-time static posture classification software. Figure 2.19 shows the dialog box that is created when the program starts. The *Initialize* button must first be pressed in order to initial the hardware and software variables containing the means, eigenvectors, and eigenvalues of each posture space.

Once initialized, the program continually loops through capturing a sitting pressure distribution map and classifying the captured sitting pressure distribution map. The text window displays a message stating the classification of the current sitting pressure distribution map. If the maximum pressure value in the sitting pressure distribution map is below a threshold, the seat is said to be empty (see Figure 2.20).

When the seat is not empty, the sitting pressure distribution map is classified as posture k if the reconstruction of the projection of the current pressure distribution map onto the kth posture space is closest to the current map. This label and the distance to the reconstruction is displayed in the text window along with the labels and distances of the second and third closest reconstructions. Figure 2.21 shows the classification for a subject sitting in posture LLRLC. The classification of the test posture is correctly given as Leaning Left Right Leg Crossed. Note that the next two closest posture spaces are those postures most similar to LLRLC: LNL and RLC.

2.3.4 Classification Results

The real-time static posture classification system was evaluated in three ways. First, execution 5me as a function of number of eigenvectors used was measured. The results are shown in Table 2.4. It was observed when using 20 eigenvectors in classification, there is a
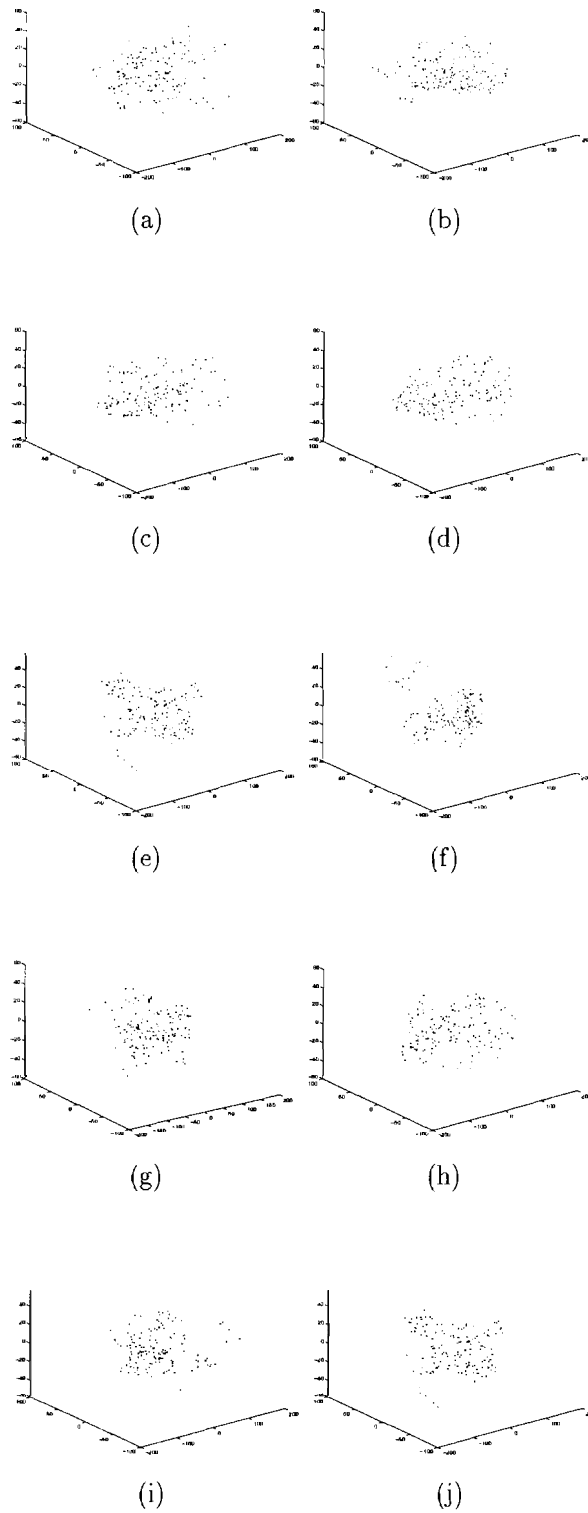
Fig. 2.18. Projections of the training samples for each posture onto the first three eigenvectors of its posture space. (a) N (b) LNF (c) LNL (d) LNR (e) RLC (f) LLC (g) LLRLC (h) LRLLC (i) LNB (j) SL.
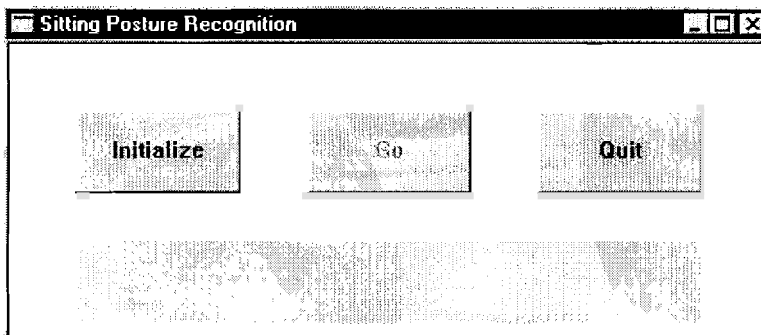
Fig. 2.19. Main window of the Static Posture Classification System.
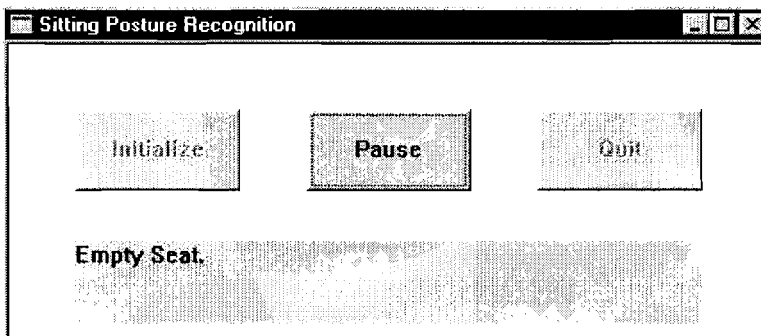
Fig. 2.20. Output when the sensing chair is empty.
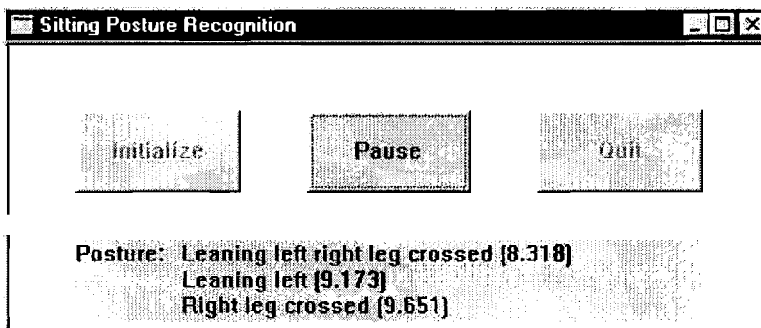
Fig. 2.21. This posture is correctly classified as LLRLC.

noticeable delay between the time of moving to a new posture and that of the display of the classification result. As the number of eigenvectors is decreased, say to 10, this delay is no longer observable.

Number of Eigenvectors	Average Classification Time (ms)
5	62.07
10	107.81
15	168.13
20	241.0

Table 2.4   Average classification time.

Second! extra pressure-map samples that did not get used in eigen decomposition were used to test the accuracy of the posture classification system, again, as a function of number of eigenvectors used. There were a total of 200 extra samples, 20 for each of the 10 postures. The results in Figure 2.22 shows an increase in accuracy as the number of eigenvectors used in classification is increased. Even when only 15 eigenvectors are used, the overall accuracy of the system is 96% correct. The reason the accuracy does not increase monotonically with dimension is due to the differences in the energy associated with the added eigenvector between posture spaces. For example, a test pressure distribution sample may be classified as posture P_1 for dimension $d = 1$. When the second eigenvector is added in the computation of the DF'FS, the same test sample may be classified as posture P_2. This is possible if the second eigenvector in posture space P_2 is more representative of the test sample than that of posture space P_1. Table 2.5 lists the accuracy as computed for each posture averaged over different numbers of eigenvectors used. These values range from 90.3% for posture LNB to 99.8% for posture SL.

For comparison, the test above was repeated on the training and testing data without the application of the smoothing operator. The mean-square error in the eigen-decomposition was greater and the overall accuracy was reduced. Figure 2.23 shows that the classification accuracy when the training and test data were smoothed was on average 7% higher than when training and test data were not smoothed.
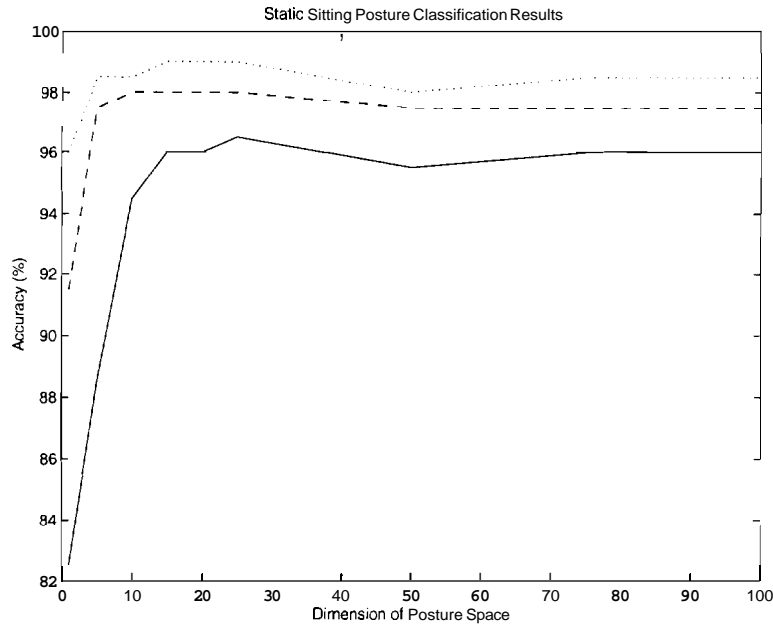
Fig. 2.22'. Accuracy vs. the number of eigenvectors used for classification. The solid line indicates the accuracy of the Static Posture Classification System. The dashed line shows the accuracy if either the closest or second closest posture space is the correct classification. The dotted line depicts the accuracy if the correct classification is one of the three nearest posture spaces to the test sample.

Posture Class	Classification Accuracy (%)
N	97.4832
LNF	95.6376
LNL	95.1678
LNR	98.2215
RLC	95.0671
LLC	94.7651
LLRLC	93.5235
LRLLC	96.5436
LNB	90.3020
SL	99.7987

Table 2.5 Classification accuracy for each posture class averaged over different numbers of eigenvectors used.
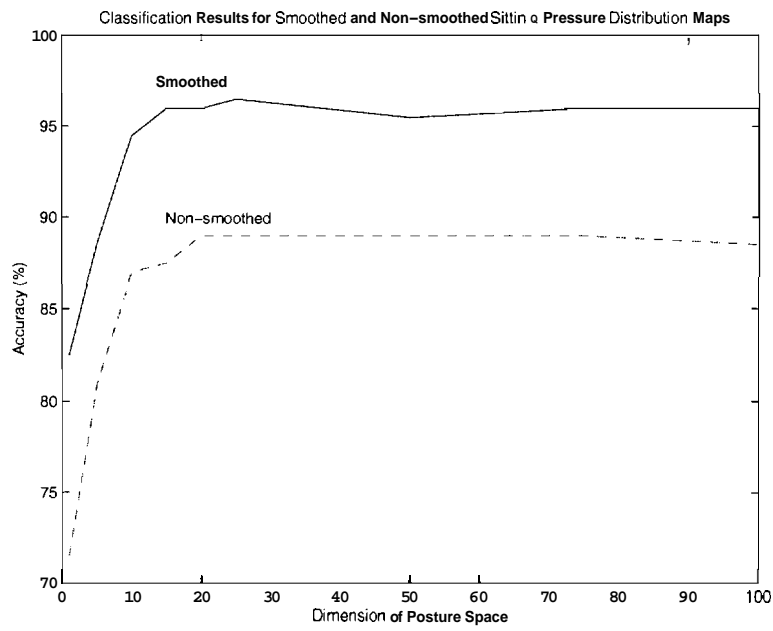
Fig. 2.23. Accuracy vs. the number of eigenvectors used for classification for both smoothed (solid line) and non-smoothed (dashed line) training and test data.

Third, real-time evaluation was conducted using four subjects who contributed pressure-distribution samples to the training data, and four others who did not. In general, the system correctly classified sitting postures for all subjects. Even using as few as five eigenvectors in classification, the system could still correctly classify static postures. The fact that our system performed equally well with the two subjects whom it had never "felt." before indicates that it can be used as a multi-user system as long as the user's build is represented in the Static Posture Database.

THIS PAGE INTENTIONALLY LEFT BLANK
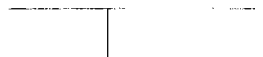
3. PROPOSED WORK

This chapter describes the proposed work for the Sensing Chair Project. As stated in the first chapter, the work can be divided into two parts:

- A real-time Static Posture Classification system

- A real-time Dynamic Posture Tracking system.

Chapter 2 describes the development of a real-time classification system for labeling a set of preselected sitting postures. This involved the construction of a Static Posture Database and the development of an appearance-based classification system. The Static Posture Database contains 1.500 samples of sitting pressure distribution maps (five samples each of ten different postures from 15 female and 15 male subjects). The real-time sitting posture classification system has been implemented. This system interprets the sitting pressure distribution data from a sensing chair outfitted with pressure sensors to identify the posture which the data represents. A quantitative measurement on its accuracy was calculated on a set of sitting pressure distribution samples that were collected from the subjects but not used to train the system. These results show that an accuracy of 96 percent is achieved with a reduction in dimension D from 4032 to 10. Real-time tests have shown that the system is also capable of classifying sitting pressure distribution maps from people who did not contribute to the static posture database.

In this chapter, we will discuss further work on the Sensing Chair Project. This involves improving the Static Posture Classification System and the development of a real-time Dynamic Posture Tracking System. Section 3.1 describes how we wish to improve posture classification. Section 3.2 describes our ideas for the Dynamic Posture Tracking System. Finally, Section 3.3 will discuss some issues related to system evaluation.

3.1 Improving Classification

As mentioned in Section 2.3.4, our overall classification accuracy is 96% correct using only 10 eigenvectors for reconstruction. This is averaged over all 10 postures. The classification accuracy for individual postures varies from 90.3% for posture Leaning Back to 99.8% for posture Slouching (see Table 2.5 for all result's). While this indicates that the system most accurately identifies when a person is slouching, which is good for correcting the person's posture, improvements can be made to increase the accuracy of the other postures.

We would like to improve the classification accuracy of postures such as Leaning Back and Leaning Left with Right Leg Crossed. One way to do this is to incorporate additional information in the classification process. Figure 2.22 showed the classification scores if one of the first one, two or three choices for posture was the correct choice on our test samples. Knowing which postures are more likely to be misclassified can indicate when additional information is needed. One way to do this would be to look at the difference between the two smallest DFFS values. A small difference would indicate that a test map could be classified as either of the two postures. Additional information could either verify or reject the posture corresponding to the minimum DFFS value.

We will be looking at incorporating a mixture-model of the sitting pressure distribution data into the classification process. Each sitting pressure distribution map will be modeled with a set of lognormal densities that will be learned by using the EM algorithm [18], [53]. The distribution of the means and covariances of the densities for each of the postures will be learned from the training samples in the Static Posture Database and used for classification in a manner similar to the hierarchical model method developed in [9].

Additional evaluation of the system will be conducted on sample sitting pressure distribution maps from new sitters.

3.2 Dynamic Posture Tracking System

By its design, the static posture classification system does not always classify postures correctly when the subject is in transition. This section describes the two components necessary for the development of a Dynamic Posture Tracking System. Section 3.2.1 describes

a Dynamic Posture Database that will be used to collect sitting pressure distribution map sequences for training and for testing. Section 3.2.2 describes the framework for the dynamic posture tracking.

3.2.1 Dynamic Posture Database

A database containing dynamic sitting pressure distribution data is needed to train and test a sitting posture tracking system. To be consistent with the Static Classification System and to restrict the seemingly unconstrained domain of dynamic sitting posture, we will collect a set of *movies* that contain a single transition between postures in the Static Posture Database. By movies, we mean a sequence of sitting pressure distribution rnaps that start at one *static* posture and end at a different *static* posture and are recorded at a fixed sampling rate. We will collect only those sequences that move directly from one posture to another (e.g., *leaning left → upright*). An *indirect* sequence would move through an alternate posture before transitioning to the end posture (e.g., *leaning left → upright → leaning right*). We are using the single transition sequences because they are representative of how a person working in an office environment would move.

3.2.2 Dynamic Posture Tracking System Development

The posture classification system determines the posture represented by a sitting pressure distribution map. In a posture tracking system, interpreting the posture represented by the current sitting pressure distribution map is more difficult because the tracking system will have to be able to handle a sitting pressure distribution map when it lies near to the edge of the posture's class distribution. This section discusses some of the aspects in formalizing a real-time sitting posture tracking system.

In moving from one posture to another the sitting pressure distribution undergoes changes. For example, a pressure peak can appear and then disappear in going from one posture to another. The occurrence of this peak and, in general, other key transition points such as this can provide a tracking system with a road map or story board to follow from one posture to the next.

There are two modalities of tracking to investigate. Those that use an appearance-based framework and those that use a 3-D human model to recover the pose of the person sitting in the chair. In an appearance-based framework, action is interpreted from a sequence of 2-D images. Representations can be view-based or feature-based. Drawbacks to these approaches include dealing with complex background and correct extraction of features. In our data, we do not, have the problem of complex backgrounds. It is straightforward to segment out the various parts of the body in the sitting pressure distribution maps. Systems that utilize a 3-D model tend to be more robust but still rely on feature extraction. These systems can also use a. Kalman filtering approach to predict future poses of the object. We believe the method best suited for posture tracking is one that uses a 3-D object model.

Here, our model is a model of the human body. What is different in our problem, as compared to tracking a walking human, is that the subject deforms against the surfaces of the chair, as compared to a projection onto an image plane. A simple skeleton model or cylindrical model is inadequate. To compensate for this model must be deformable, such as with a Finite Element Model. This would involve extending such work as [59] and [19] which deal with modeling deformation of the soft tissue in human thighs in seated postures.

With the recovered pose of the 3-D model, numerous features can be extracted to give specific information regarding the posture of the person sitting in the chair. These features include the orientation of the spine, the pose of the legs, and the locations of the shoulders.

Key issues of the system include:

- The type of 3-d model

  The parts of the body that come in contact with the chair include the shoulders, back, buttocks, and thighs. The model will need to incorporate these itenns as well as take into account the head, arms, and lower legs because they affect the pressure distribution on the chair.

- Extracting local features (e.g., orientation of the spine)

  With the current pressure distribution and 3-D model in registration, various features can be extracted for use in a more specific description of posture. For example, the

orientation of the pelvis can indicate in which direction a user is looking or is interested in.

- Robustness: Dealing with the variation in subject size

  We want a multi-user system and do not want the system to be extensively trained for each new user. Having an initial bootstrapping procedure when a user first sits in the chair will permit the system to appropriately scale the model. This can be done by having the user sit in the upright position while the system determines such features as the total force, the size and weight distribution in the thigh area, and the distribution of pressure on the seat back.

- Real-time tracking

  One of the key facets to our work is that our system work in real-time.. This will involve determining the minimal amount of pressure information needed for tracking.

3.3 Performance Evaluation

Testing, as with the Static posture Classification System, will be conducted on pre-recorded posture sequences not used in training and on real-time testing of subjects who did and did not contribute samples used in training. The tests will determine how well the system performs on subjects, whose anthropometry is represented in the training data, and will indicate the system's robustness.

An area in which to evaluate a tracking system is in the reduction of pressure data in the sitting pressure distribution map. The reasoning behind this is two-fold.

First, we are interested the development of a real-time posture tracking system. Each of the two pressure mats contains 2016 pressure sensing elements. This gives a sitting pressure distribution map that contains 4032 data points. Tests of tracking performance and execution time will be conducted on down-sampled sitting posture presure maps.

Second, if a real-time posture tracking system were to be implemented in a real-world application, it would be desired to be economical. While this is not an immediate goal, reducing the number of sensing elements is one way to accomplish this. This can be simulated by down-sampling the sitting pressure distribution maps.

What we seek is a graph showing system performance (e.g., accuracy, speed, cost, etc.) vs. down sampling and dimension D. Again, D means the dimension of our data (e.g., the number of eigenvectors or number of features). For example, Figure 3.1 shows possible curves for the performance indices of speed and accuracy as a function of down sampling. The selection of how much to down sample and to what to set the dimension are application dependent. Knowing the shapes of the curves describing how these parameters affect system performance can guide a person to set them to appropriate values. For example, if the application was automatic control of airbag deployment force, down sampling would be limited in order to maintain a high accuracy rate and speed. If the application was the medical analysis of posture for a person suffering from lower-back pain, then such an analysis could be carried out off-line and speed would not be a factor. In this case, dimension would mainly be determined by the desired accuracy. By systematically characterizing the tradeoffs among parameters such as resolution, accuracy, speed, and cost, we hope to provide a foundation for the applications of our Posture Tracking System.
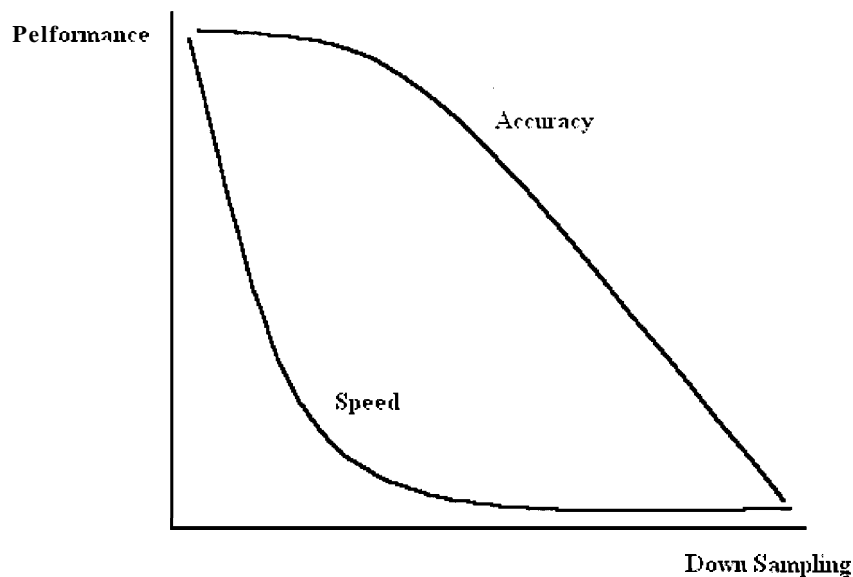
Fig. 3.1. System performance as a function of down sampling.

LIST OF REFERENCES

[1] K. Akita. Image sequence analysis of real world human motion. *Pattern Recognition,* 17, 1984.

[2] S. Baker and S. K. Nayar. A theory of pattern rejection. Technical Report CUCS-013-95, Department of Computer Science, Columbia University, 1995.

[3] Richard I. Barnett and Frederick E. Shelton. Measurement of support surface efficacy: Pressure. *Advances In Wound Care,* pages 21–29, Nov/Dec 1997.

[4] Peter N. Belhumeur, Joao P. Hespanha, and David J. Kriegman. Eigenfaces vs. Fisherfaces Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 19(7):711–720, September 1997.

[5] James A. Birke, James G. Foto, Sunil Deepak, and Jean Watson. Measurement of pressure walking in footwear used in leprosy. *Lepr Rev,* 65:262–271, 1994.

[6] R. Brooks. A robust layered control system for a mobile robot. Technical Report AI Lab Memo 864, Massachusetts Institute of Technology, 1985.

[7] R. Brooks. Intelligence without representation. *Artificail Intelligence,* 47:1–3, 1991

[8] R. A. Brooks. Model-based three-dimensional interpretations of two-dimensional images. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 5(2):140–150, March 1983.

[9] Igor V. Cadez, Christine E. McLaren, Padhraic Smyth, and Geoffrey J. McLachlan. Hierarchical Models for Screening Iron Deficiency Anemia. In *in Proceedings of the 1999 International Conference on Machine Learning.* June 1999.

[10] Chu-Yin Chang. *Eigenspace Methods for Correlated Images.* PhD thesis, Purdue University, 1999.

[11] Michael H. Coen. Building brains for rooms: Designing distributed software agents. In *Proceedings of the 9th Conference on Innovative Applications of Artificial Intelligence (AAAI-97),* Providence, RI, 1997.

[12] Michael H. Coen. The future of human-computer interaction, or how i learned to stop worrying and love my intelligent room. *IEEE Intelligent Systems,* pages 8–10, March/April 1999.

[13] David Cohen. An objective measure of seat comfort. Aviation, Space: and Environmental *Medicine*, 69(4):410–414, April 1998.

[14] Stephen F. Conti, Rob L. Martin, Ruth Chaytor, Christopher Hughes, and Leslie Luttrell. Plantar pressure measurement during ambulation in weightbearing conventional short leg casts and total contact casts. Foot and *Ankle International*, 17(8):464–469, August 1996.

[15] Peter Coy. Ready for its maiden flight: a new airline seat. Business Week, page 199, November 4 1996.

[16] T. Darrell and A. Pentland. Space-time gestures. In proceedings of the IEEE international conference on computer vision and pattern recognition, 1993.

[17] James W. Davis and Aaron F. Bobick. The representation of action using temporal templates. In Proceedings of the IEEE Conference *on* Computer Vision and *pattern* Recognition, 1997.

[18] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. Journal of the Royal Statistical Society, 39(1):1–22, 1977.

[19] Bing Deng and Robert Hubbard. Measuring and Modeling Force-Deflection Responses of Human Thighs in Seated Postures. Advances *in* Bioengineering, *ASME* Bioengineering *Division* (BED), 28:101–102, 1994.

[20] R. O. Duda and P. E. Hart. Pattern *Classification* and Scene Analysis. John Wiley & Sons, 1973.

[21] K. Etemad and R. Chellapa. Discriminant analysis for recognition of human face images. *Journal* of Optics of America, 14(8):1724–1733, 1997.

[22] James L. Flanagan. Autodirective sound capture: towards smarter conference rooms. IEEE Intelligent Systems, pages 16–19, March/April 1999.

[23] David Franklin. The intelligent classroom. IEEE *Intelligent* Systems, pages 2–5, September/October 1999.

[24] K. Fukunaga. Introduction to Statistical Pattern Recognition. Academic Press, New York, second edition, 1990.

[25] D. Gavrila and L. Davis. Tracking of humans in actions: a 3-d model-based approach. In *A*RPA Image Understanding Workshop, February 1996.

[26] L. Goncalves, E. Dibernardo, E. ursella, and P. Perona. Monocular tracking of the human arm. In ICCV, June 1995.

[27] Kent M. Van De Graaff and Stuart Ira Fox. Concepts of *Human anatomy* and physiology. Wm. C. Brown Publishers, Boston, fourth edition, 1995.

[28] L. Grewe and A. C. Kak. interactive learning of a multiple-attribute hash table classifier for fast object recognition. Computer *Vision* and Image Understanding, 61(3):387–416, May 1995.

[29] Clifford M. Gross, Ravindra S. Goonetilleke, Krishna K. Menon, Jose Carlos N. Banaag, and Chandra M. Nair. The biomechanical assessment and prediction of seat comfort. In Rani Lueder and Kageyu, editors, Hard facts about *soft* machines: *the Ergonomics* of Seating. Taylor & Francis, Bristol, PA, 1994.

[30] Lisa Guernsey. Magic carpet chairs of virtual reality. The New York Times, page G13, May 6 1999.

[31] Peter J. B. Hancock, A. Mike Bruton, and Vicki Bruce. Face processing: Human perception and principal components analysis. Memory and Cognition, 24(1):26–40, 1996.

[32] D. Hogg. Model-based vision: a paradigm to see a walking person. *Image* and *Vision* Computing, 1(1), 1983.

[33] Eric C. Hughes, Wenqi Shen, and Alicia Vertiz. The effects of regional compliance and instantaneous stiffness on seat back comfort. Society of *Automotive* Engineers, (980658), 1998.

[34] Wesley R. Iversen. Tactile sensing, 1990s style. *Assembly*, pages 23–26, Feb-Mar 1993.

[35] Barry H. Kantowitz and Robert D. Sorkin. *Human* Factors: Understaning People-System Relationships. John Wiley & Sons, New York, 1983.

[36] Klaus Kompaβ and Michel Witte. The BMW Seat Occupancy Monitoring System: A Step Towards "Situation Appropriate Airbag Deployment". Topics in *Vehicle* Safety *(SP-1139)*, (960226), February 1996.

[37] J. Little and J. Boyd. Describing motion for recognition. In *Proceedings* of ihe IEEE International *Symposium* on Computer *Vision*, Coral Gables, FL, November 1995.

[38] Rani Lueder and Kageyu Noro, editors. Hard *Facts* About *Soft* Machines. Taylor & Francis, London, 1994.

[39] Steve Mann. Humanistic intelligence: Wearcomp as a new framework for intelligent signal processing. Proceedings of the IEEE, 86(11):2123–2125, November 1998.

[40] Aleix Martinez. Face images retrieval using HMMs. In IEEE *Workshop* on Content-*Based* Access of Images and Video Libraries, 1999.

[41] Marvin Minsky. The Society of *Mind*. Simon & Schuster, New York, 1986.

[42] Michael C. Mozer. An intelligent environment must be adaptive. IEEE *Intelligent* Systems, pages 11–13, March/April 1999.

[43] M. J. Mueller and M. J. Strube. Generalizability of in-shoe peak pressure measures using the f-scan system. Clinical Biomechanics, 11(3):159–164, 1996.

[44] Hiroyasu Murakami and B. V. K. Vijaya Kumar. Efficient calculation of primary images from a set of images. IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-4(5):511–515, 1982.

[45] H. Murase and S. K. Nayar. Visual learning and recognition of 3D objects from appearance. International Journal of Computer Vision, 14(1):5–24, 1995.

[46] H. Murase and S. K. Nayar. Detection of 3D objects in cluttered scenes using hierarchical eigenspace. Pattern Recognition Letters, 18(4):375–384, 1997.

[47] Hiroshi Murase and Shree K. Nayar. Learning models for appearance. In proceedings of *AAAI*, Washington, D.C., July 1993.

[48] Shree K. Nayar, Sameer A. Nene, and Hiroshi Murase. Real-time 100 object recognition system. In Proceedings of the 1996 IEEE International Conference on Robotics and *Automation*, pages 2321–2325. IEEE Robotics and Automation Society, April 1996.

[49] S. A. Nene and S. K. Nayar. Algorithm and architecture for high dimensional search. Technical Report CUCS-030-95, Department of Computer Science, Columbia University, 1995.

[50] Alex Pentland, Baback Moghaddam, and Thad Starner. View-based and modular eigenspaces for face recognition. In IEEE International Conference on Computer Vision and *Pattern* Recognition, 1994.

[51] Alex P. Pentland. Smart rooms. Scientific American, 274(4):68–76, April 1996.

[52] R. Polana and R. Nelson. Low level recognition of human motion. In IEEE Workshop on Non-rigid and articulated motion, 1994.

[53] Richard A. Redner and Homer F. Walker. Mixture Densities, Maximum Likelihood and the EM Algorithm. Society for Industrial and Applied *Mathematics* Review, 26(2):195–239, April 1984.

[54] J. Rehg and T. Kanade. Model-based tracking of self-occluding articulated objects. In *ICCV*, June 1995.

[55] Mary Roach. Hot seat. Discover, pages 74–77, March 1998.

[56] K. Rohr. Towards model-based recognition of human movements in images sequences. CVGIP, Image *Understanding*, 59(1), 1994.

[57] A. Rosenfeld and A. C. Kak. Digital Picture Processing, volume 2. Academic Press, San Fransisco, second edition, 1982.

[58] Joseph A. Sember III. The biomechanical relationship of seat design to the human anatomy. In Rani Lueder and Kageyu, editors, *Hard facts about soft machines: the ergonomics of seating.* Taylor & Francis, Bristol, PA, 1994.

[59] Richard H. Setyabudhy, Akram Ali, Robert P. Hubbard, Clifford Beckett, and Ronald C. Averill. Measuring and Modeling of Human Soft Tissue and Seat Interaction. *Progress with Human Factors in Automotive Design: Seating Comfort, Visibility, and Safety SAE Special Publications,* 1242(970593):135–142, 1997.

[60] E. Shavit and A. Jepson. Motion understanding using phase portraits. In *IJCAI Workshop: Looking at People,* 1995.

[61] D. L. Swets and J. J. Weng. Using discriminant eigenfeatures for image retrieval. *IEEE Transactions on Pattern Analysis and machine Intelligence,* PAMI-18(8):831–836, 1996.

[62] Hong Z. Tan. A Sensing Chair. In *Proceedings of the American Society of Mechanical Engineers: Dynamic Systems and Control,* Nashville, TN, November 15 1999. To appear.

[63] Tekscan, Inc., South Boston, MA. *Tekscan Body Pressure Measurement System User's Manual.* 1998.

[64] Kuntal Thakurta, Daniel Koester, Neil Bush, and Susan Bachle. Evaluating short and long term seating comfort. *Society of Automotive Engineers,* (950144), 1995.

[65] Kuntal Thakurta and Cindy Linder. Seat design gets comfortable. *Automotive Body Interior and Safety Systems, IBEC,* pages 7–11, 1995.

[66] Mark C. Torrance. Advances in human-computer interaction: The intelligent room. In *CHI '95 Research Symposium,* Denver, CO, May 6-7 1995.

[67] Matthew Turk and Alex Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience,* 3(1):71–86, 1991.

[68] J. J. Weng. *Early Visual Learning,* chapter Crescepton and SHOSLIF: Towards comprehensive visual learning. Oxford university Press, 1996.

[69] A. Wilson and A. Bobick. Learning visual behavior for gesture analysis. In *Proceedings of the IEEE International Symposium on Computer Vision,* Coral Gables, FL, November 1995.

[70] Christopher R. Wren. Understaning expressive action. Technical Report 498, Massachusetts Institute of Technology, 1999.

[71] J. Yamato, J. Ohya, and K. Ishii. Recognizing human action in time sequential images. In *proceedings of the IEEE international conference on computer vision and pattern recognition,* 1992.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A: RESULTS OF K-MEANS
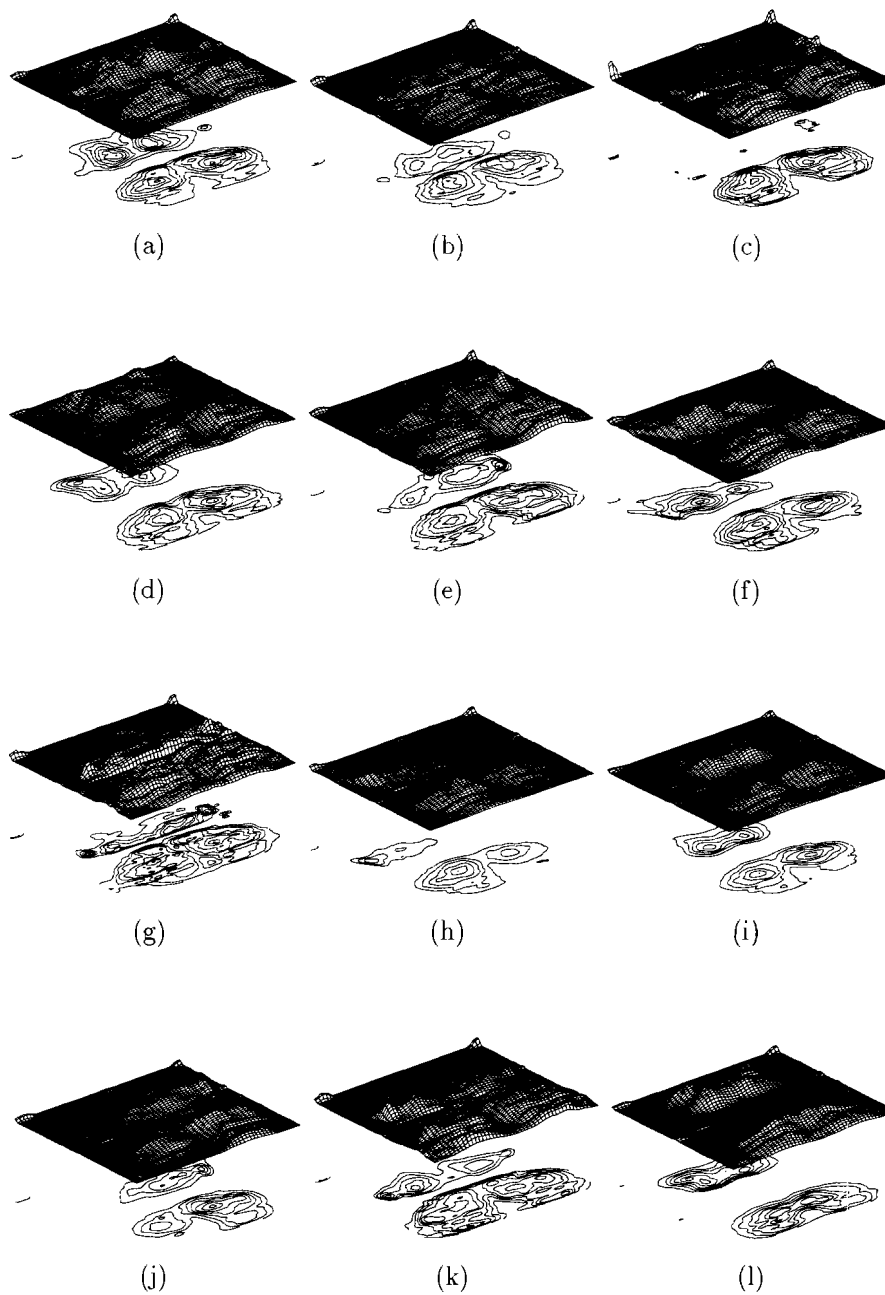
Fig. A.1. Output from running K-means on the data in the Static Posture Database with K = 12.
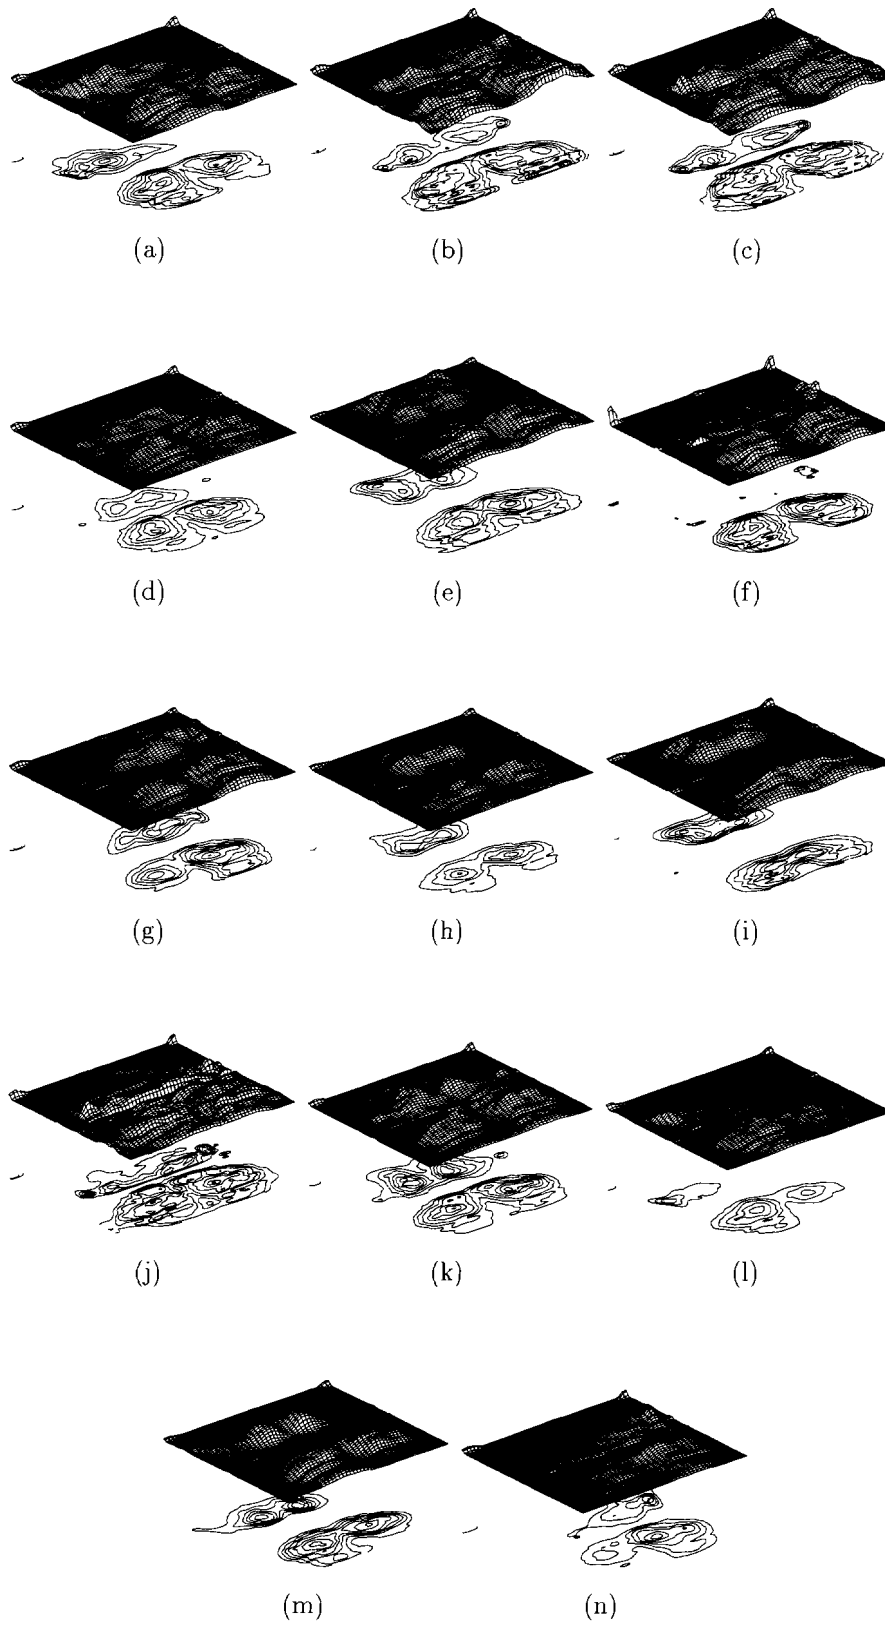
Fig. A.2. Output from running K-means on the data in the Static Posture Database with K = 14.
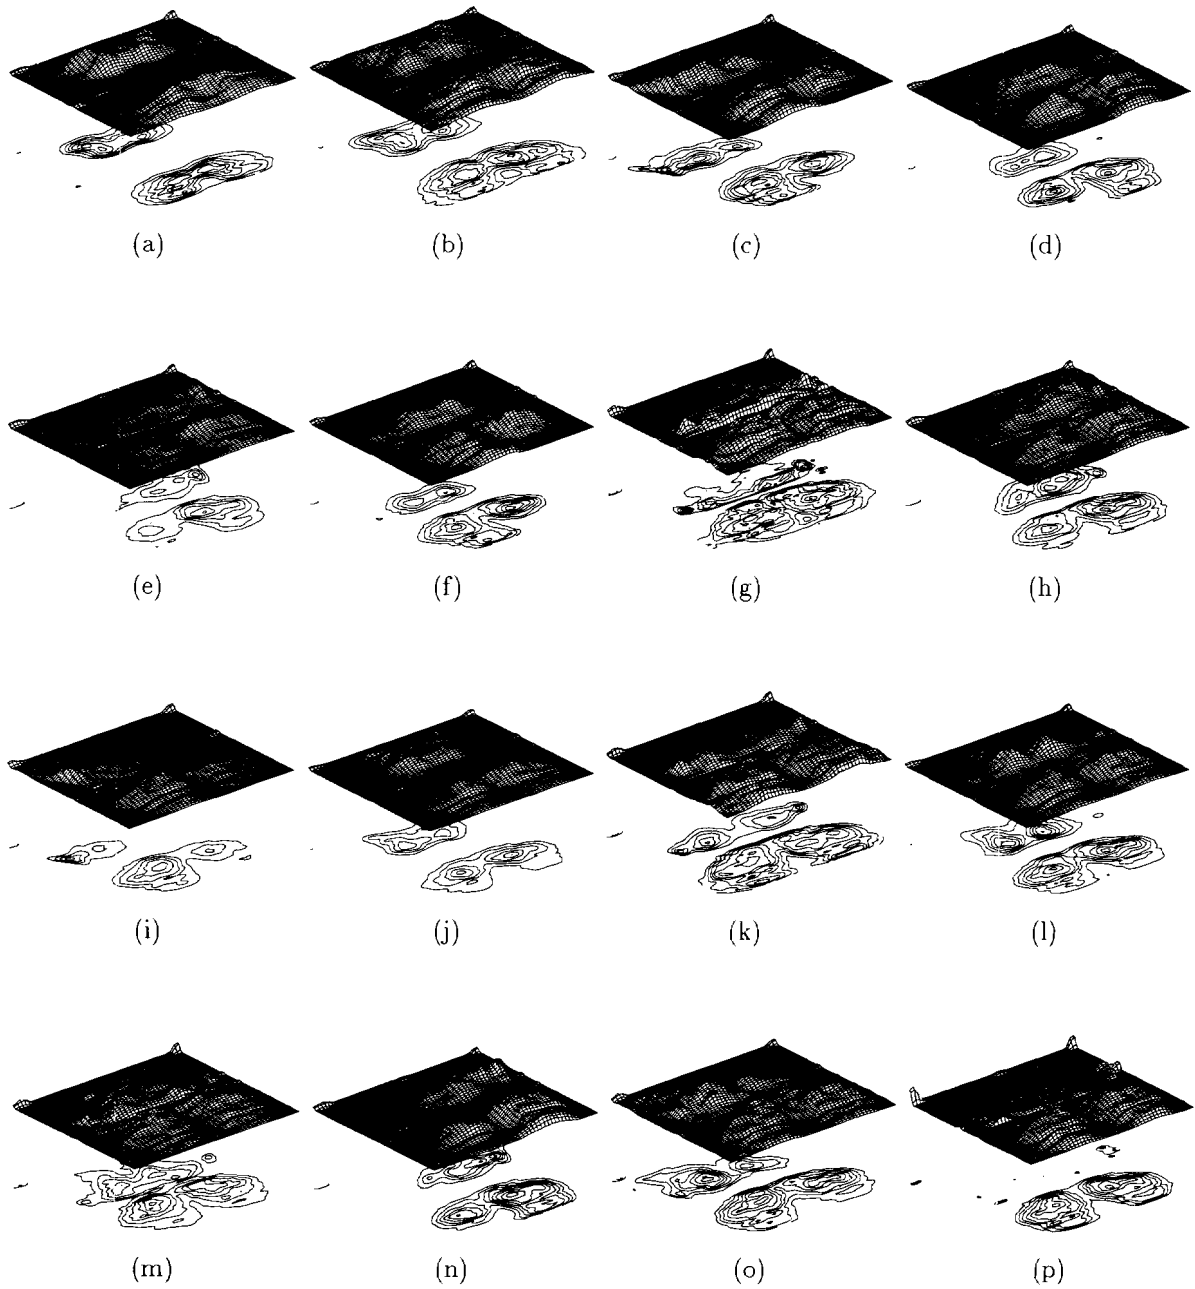
Fig. **A.3.** Output from running K-means on the data in the Static Posture Database with K = 16.
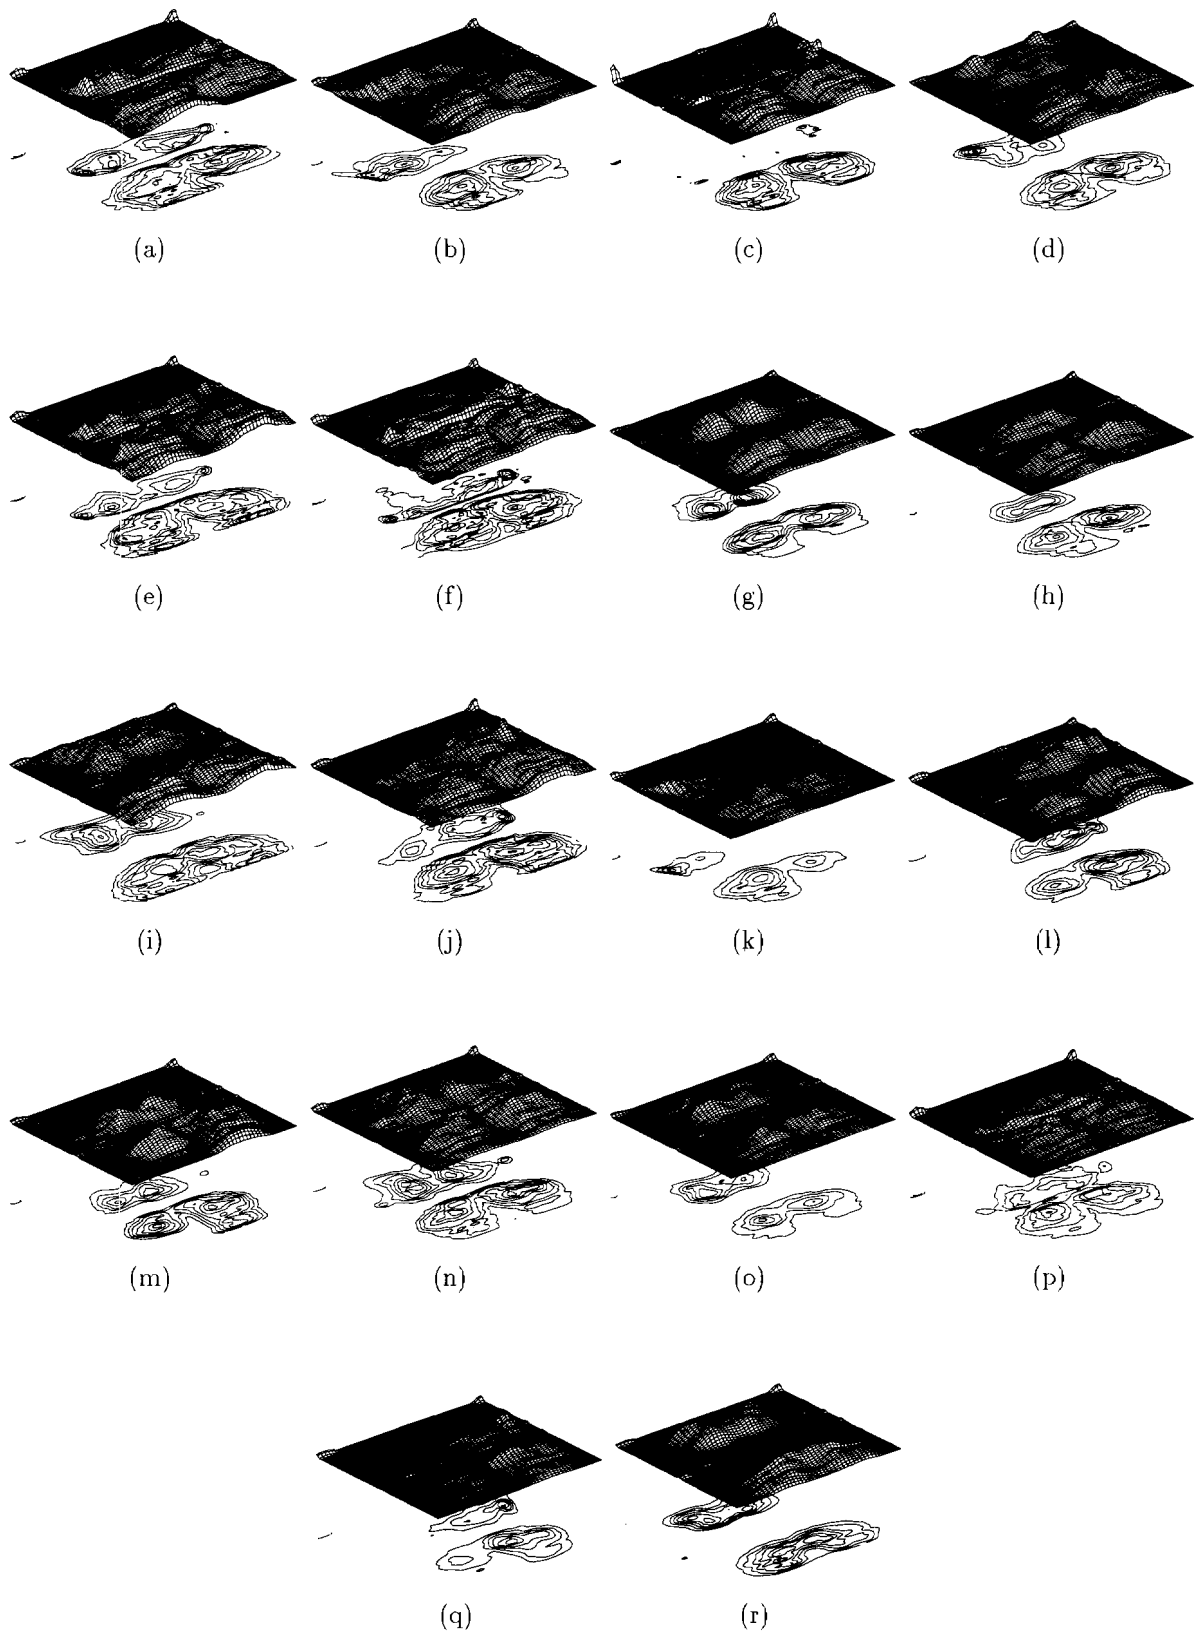
Fig. A.4. Output from running K-means on the data in the Static Posture Database with
K = 18.
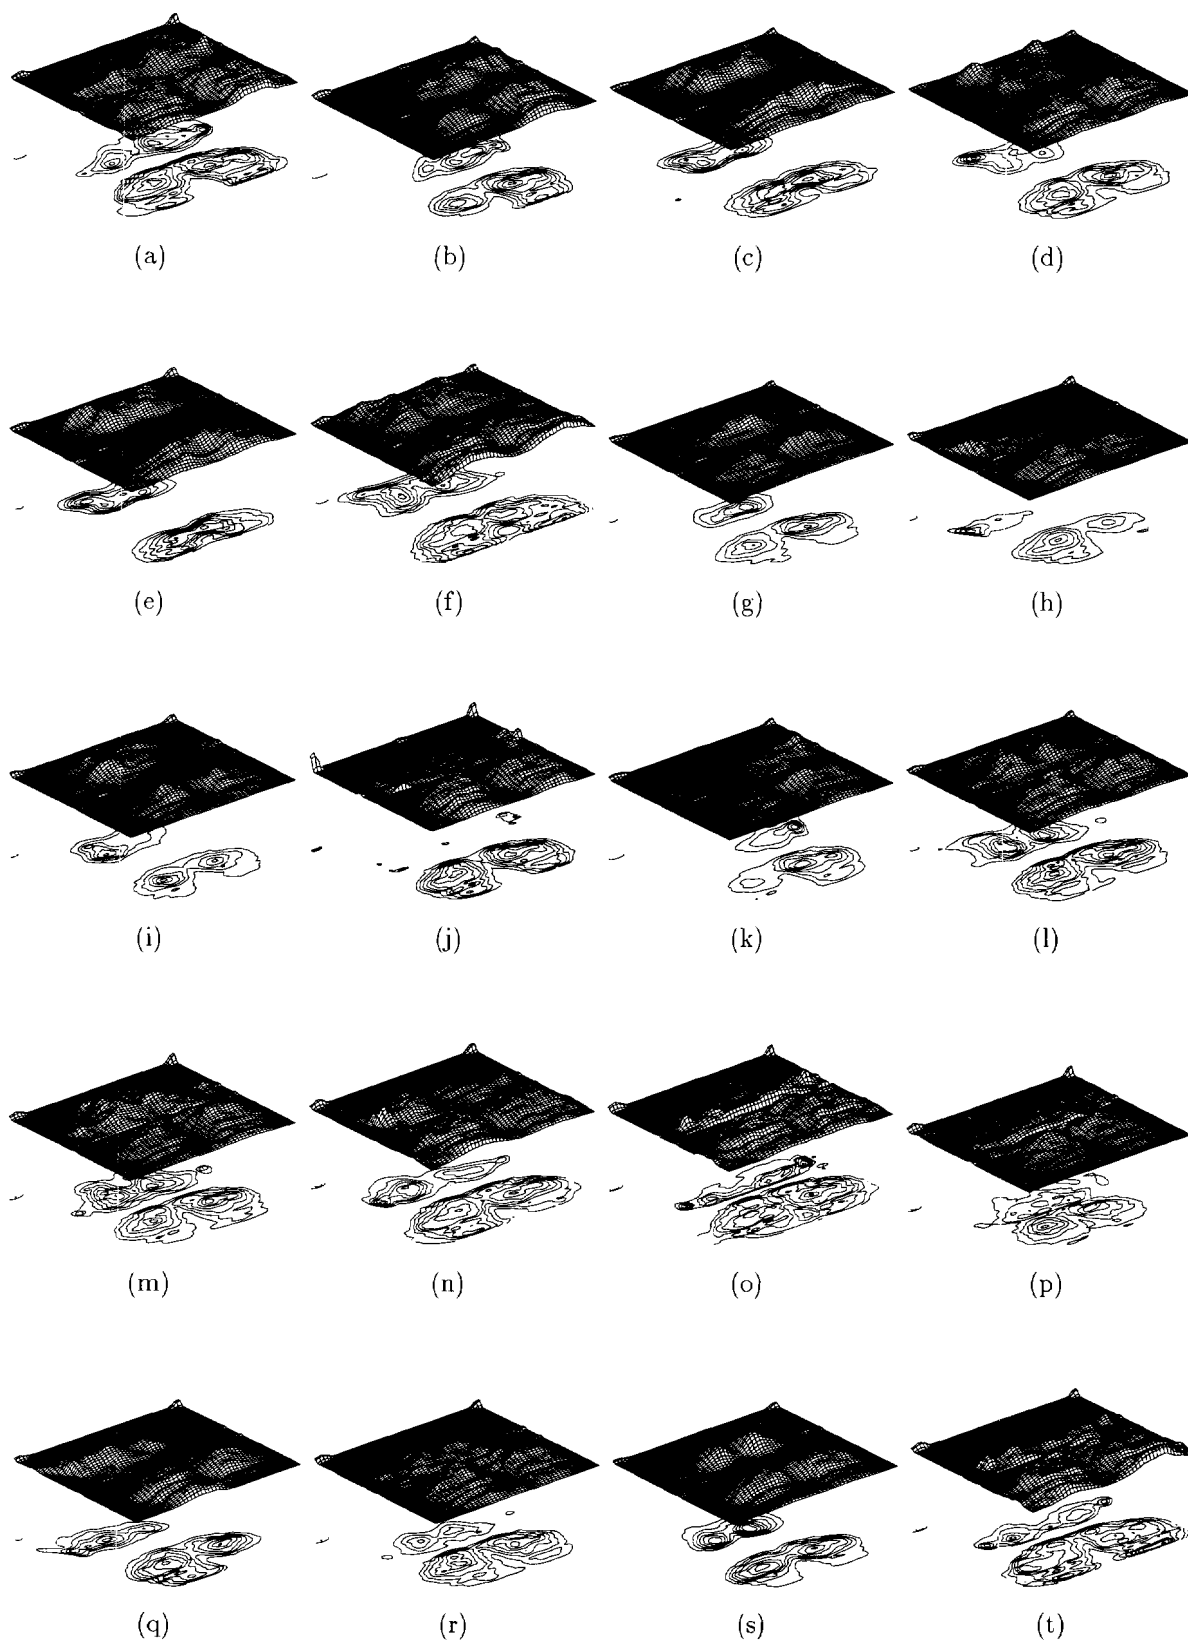
Fig. A.5. Output from running K-means on the data in the Static Posture Database with K = 20.

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Table
a	40	0	4	2	42	25	14	1	10	1	139
b	35	20	13	12	32	27	5	3	5	0	152
c	7	82	2	2	1	0	1	0	0	0	95
d	6	0	0	0	0	2	0	0	82	24	114
e	0	0	0	40	5	0	13	0	0	0	58
f	0	0	33	0	2	2	0	59	0	0	96
g	5	17	6	9	5	5	5	5	2	0	59
h	5	14	69	0	8	8	1	57	12	0	174
j	33	2	4	5	42	71	6	19	29	5	216
i	0	0	0	79	9	0	104	0	0	0	192
k	19	15	19	1	4	9	0	6	0	0	73
l	0	0	0	0	0	1	1	0	10	120	132

Table A.1 Number of samples from each posture that belong to each cluster from k-means with K=12. The cluster labels in the table correspond to the image of the cluster mean of Figure A.1.

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Table
a	0	0	41	0	0	1	0	39	0	0	81
b	9	7	8	13	4	0	5	0	0	0	46
c	11	8	11	7	6	10	0	7	0	0	60
d	39	21	12	6	45	35	0	2	10	0	170
e	5	0	0	0	0	1	0	0	79	31	116
f	7	81	1	1	0	0	0	0	0	0	90
g	0	0	0	33	11	3	72	0	0	0	119
h	27	4	0	10	17	37	10	6	29	7	147
i	0	0	0	0	0	0	1	0	8	111	120
j	5	17	9	10	5	5	5	5	2	0	63
k	41	0	5	6	42	28	13	5	10	1	151
l	4	12	61	0	1	5	0	49	11	0	143
m	2	0	2	0	17	25	0	37	1	0	84
n	0	0	0	64	2	0	44	0	0	0	110

Table **A.2** Number of samples from each posture that belong to each cluster from k-means with K=14. The cluster labels in the table correspond to the image of the cluster mean of Figure A.2.

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Table
a	0	0	0	0	0	1	0	0	6	92	99
b	2	0	0	0	0	1	0	0	48	48	99
c	0	0	7	0	2	3	0	52	0	0	64
d	11	0	3	4	62	0	11	0	0	0	91
e	0	9	0	70	2	1	47	0	0	0	129
f	7	0	2	0	0	62	0	25	0	0	96
g	7	18	6	12	6	5	5	5	5	0	69
h	0	0	0	26	2	0	22	0	0	0	50
i	4	8	73	0	2	1	0	46	2	0	136
j	21	4	0	4	15	18	6	3	53	9	133
k	19	15	18	20	9	9	4	5	0	0	99
l	36	0	0	0	28	37	6	3	28	1	139
m	33	7	6	2	18	10	0	0	4	0	80
n	0	0	0	11	0	1	49	0	0	0	61
o	1	0	31	0	4	1	0	11	2	0	50
p	9	89	4	1	0	0	0	0	2	0	105

Table **A.3** Number of samples from each posture that belong to each cluster from k-means with K=16. The cluster labels in the table correspond to the image of the cluster mean of Figure **A.3.**

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Table
a	1	0	3	0	0	15	0	15	0	0	34
b	0	0	33	0	1	1	0	53	0	0	88
c	7	86	1	3	0	0	0	0	0	0	97
d	5	0	0	0	0	1	0	0	80	5	91
e	16	14	15	6	5	0	5	0	0	0	61
f	10	18	9	7	5	0	5	0	4	0	58
g	3	0	1	0	30	27	9	14	0	2	86
h	22	5	6	6	15	63	0	17	0	0	134
i	0	0	0	0	0	0	0	0	17	22	39
j	0	0	0	23	5	0	7	0	0	0	35
k	4	9	62	0	1	0	0	44	5	0	125
l	0	0	0	30	1	0	61	0	0	0	92
m	0	0	0	0	54	0	7	0	0	0	61
n	40	0	10	4	10	18	3	5	13	0	103
o	16	0	0	3	10	11	6	1	23	13	83
p	26	14	10	11	13	13	5	1	0	0	93
q	0	4	0	57	0	0	42	0	0	0	103
r	0	0	0	0	0	1	0	0	8	108	117

Table A.4 Number of samples from each posture that belong to each cluster from k-means with I<=18. The cluster labels in the table correspond to the image of the cluster mean of Figure A.4.

Cluster	N	LNF	LNL	LNR	RLC	LLC	LLRLC	LRLLC	LNB	SL	Table
a	5	0	0	19	5	0	7	0	1	0	37
b	0	0	0	22	7	0	65	0	0	0	94
c	0	0	0	0	0	1	1	0	11	66	79
d	6	0	0	0	0	1	0	0	75	4	86
e	0	0	0	0	0	0	0	0	1	56	57
f	0	0	0	0	0	0	0	0	15	13	28
g	19	9	0	17	27	32	6	4	5	0	119
h	4	7	65	0	2	4	0	45	4	0	131
i	10	3	1	3	10	8	7	0	17	10	69
j	7	76	3	1	0	0	0	0	0	0	87
k	0	0	0	55	1	0	38	0	0	0	94
l	13	0	7	0	7	12	0	8	17	1	65
m	32	0	2	0	25	14	0	0	0	0	73
n	3	0	21	0	0	10	0	10	0	0	44
o	8	18	5	18	6	5	5	5	1	0	71
p	5	8	6	7	8	6	5	3	2	0	50
q	0	0	10	0	1	2	0	65	0	0	78
r	25	15	21	3	15	23	0	3	0	0	105
s	4	0	1	0	31	32	11	7	1	0	87
t	9	14	8	5	5	0	5	0	0	0	46

Table **A.5** Number of samples from each posture that belong to each cluster from k-means with K=20. The cluster labels in the table correspond to the image of the cluster mean of Figure **A.5.**

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B: GLOSSARY

Anterior: Toward the front.

Anthropometry: The measurement of the human body with a view to determine its average dimensions, and the proportions of its parts, at different ages and in different races or classes.

Biomechanics: The study of the mechanical laws relating to the movement or structure of living organisms.

Coronal Plane: The plane that passes through the length of the body and divides the front from the back.

Cummulative Trauma Disorders: Disorders of the muscles, tendons, and/or nerves caused, precipitated or aggravated by repetitive motion activity that applies stress to the body.

Ergonomics: The scientific study of the efficiancy of man in his working environment.

Inferior: Away from the head.

Ischial Tuberosities: The sitting bones of the pelvis.

Kinematics: The science of pure motion, considered without reference to the matter or objects moved, or to the force producing or changing the motion.

Lateral: Toward the side of the body.

Medial: Toward the midline of the body.

Midsagittal Plane: The plane that passes through the midline of the body dividing it into left and right halves.

Popliteal fossa: The back of the knee.

Popliteal height: The vertical distance between the bottom of the foot and the crease just behind the knee of a seated person.

Posterior: Toward the back.

Seat Depth: The horizontal distance from the most posterior part of the buttocks to the crease just behind the knee of a seated person.

Superior: Toward the head.

Trochanter: A large knobby projection at the end of the femur near the pelvis. The greater trochanter is on the lateral side of the femur and is larger in size to the lesser trochanter on the medial side.

Transverse Plane: The plane that divides the body into superior (upper) and inferior (lower) regions.