

EigenGait: Motion-based Recognition of People using Image Self-Similarity

Chiraz BenAbdelkader[†], Ross Cutler[‡],
Harsh Nanda[†] and Larry Davis[†]
[†] University of Maryland, College Park
{chiraz,rgc,nanda,lsd}@umiacs.umd.edu
[‡]Microsoft Research
rcutler@microsoft.com

Abstract

We present a novel technique for motion-based recognition of individual gaits in monocular sequences. Recent work has suggested that the image self-similarity plot of a moving person/object is a projection of its planar dynamics. Hence we expect that these plots encode much information about gait motion patterns, and that they can serve as good discriminants between gaits of different people. We propose a method for gait recognition that uses similarity plots the same way that face images are used in eigenface-based face recognition techniques. Specifically, we first apply Principal Component Analysis (PCA) to a set of training similarity plots, mapping them to a lower dimensional space that contains less unwanted variation and offers better separability of the data. Recognition of a new gait is then done via standard pattern classification of its corresponding similarity plot within this simpler space. We use the k -nearest neighbor rule and the Euclidean distance. We test this method on a data set of 40 sequences of six different walking subjects, at 30 FPS each. We use the leave-one-out cross-validation technique to obtain an unbiased estimate of the recognition rate of 93%.

1 Introduction

Human gait has long been an active subject of study in biomechanics, kinesiology, psychophysics, and physical medicine [25, 21, 28]. The reasons and applications for this interest include: detection of gait pathologies, rehabilitation of an injured person, improving athletic performance, and designing ergonomic-based athletic and office equipment. These gait studies typically analyze 3D temporal trajectories of marked points on the body in terms of their frequency and phase. The relationships among the component patterns of the gait, such as phase differences between the trajectories, are also a valuable source of information.

In machine vision, gait recognition has received growing interest due to its emergent importance as a biometric [9, 4]. The term *gait recognition* is used to signify recognizing individuals by the

way they walk in image sequences. Gait detection is the recognition of different types of human locomotion, such as running, limping, hopping, etc. Because human ambulation (gait) is one form of human movement, gait recognition is closely related to vision methods for detection, tracking and recognition of human movement in general (such as actions and gestures).

Gait recognition research has largely been motivated by Johansson’s experiments [21] and the ability of humans for motion perception from Moving Light Displays (MLDs). In these experiments, human subjects were able to recognize the type of movement of a person solely from observing the 2D motion pattern generated by light bulbs attached to the person. Similar experiments later showed some evidence that the identity of a familiar person (‘a friend’) [3], as well as the gender of the person [11] might be recognizable from MLDs, though in the latter case a recognition rate of 60% is hardly significantly better than chance (50%).

Despite the agreement that humans can perceive motion from MLDs, there is still no consensus on how humans interpret this MLD-type stimuli (i.e. how it is they use it to achieve motion recognition). Two main theories exist: The first maintains that people use motion information in the MLDs to recover the 3D structure of the moving object (person), and subsequently use the structure for recognition; and the second theory states that motion information is directly used to recognize a motion, without structure recovery. In machine vision, methods that subscribe to the former theory are known as *structure from motion* (SFM) [17], and those that favor the latter are known as motion-based recognition [8].

Consequently, there exist two main approaches for gait recognition each of which favors one of the two above theories. In SFM-based methods, a set of body points are tracked (as a result of body structure recovery), and their motion trajectories are used to characterize, and thereby recognize the motion or action performed by the body. Note that this approach emulates MLD-based motion perception in humans, since the body part trajectories are in fact identical to MLD-type stimuli. Furthermore, this approach is supported by biomedical gait research [28] which found that the dynamics of a certain number of body parts/points totally characterize gait. However, because tracking body parts in 3D over a long period of time remains a challenge in vision, the effectiveness of SFM-based methods remains limited.

Motion-based recognition methods, on the other hand, characterize the motion pattern of the body, without regard to its underlying structure. Two main approaches exist; one which represents human movement as a sequence (i.e. discrete number) of poses/configurations; and another which characterizes the spatiotemporal distribution generated by the motion in its continuum.

The method we describe in this paper takes a motion-based recognition approach. Our method makes the following assumptions:

- People walk with constant speed and direction for about 3-4 seconds.
- People walk approximately parallel to the image plane.
- The camera is sufficiently fast to capture dynamics of motion (we use 30Hz).

2 Related Work

We review vision methods used in detection, tracking and recognition of human movement in general, as they are closely related to gait recognition ([8, 1, 14] are good surveys on this topic). These methods can be divided into two main categories: methods that recover high-level structure of the body and use this structure for motion recognition, and those that directly model how the person moves. We shall describe the latter in more detail as it is more relevant to the gait recognition approach proposed in this paper.

2.1 Structural Methods

A 2D or 3D structural model of the human body is assumed, and body pose is recovered by extracting image features and mapping them to the structural components of the model (i.e. body labelling). Hence a human is detected in the image if there exists a labelling that fits the model well enough (based on some measure of goodness of fit) [17, 18, 32, 15, 33]. Once a person has been detected and tracked in several images, motion recognition is done based on the temporal trajectories of the body parts, typically by mapping them to some low-dimensional feature vector and then applying standard pattern classification techniques [2, 34, 7, 26].

2.2 Structure-free Methods

To recognize a moving object (or person), these methods characterize its motion pattern, without regard to its underlying structure. They can be further divided into two main classes. The first class of methods consider the human action or gait to be comprised of a sequence of poses of the moving person, and recognize it by recognizing a sequence of static configurations of the body in each pose [27, 20, 16]. The second class of methods characterizes the spatiotemporal distribution generated by the motion in its continuum, and hence analyze the spatial and temporal dimensions simultaneously [29, 30, 12, 24, 23, 10]. Our method is closely related to the work of [10], in that both use similarity plots to characterize human motion, though we use them for gait recognition, and Cutler and Davis use them mainly for human detection.

State-space Methods These methods represent human movement as a sequence of static configurations. Each configuration is recognized by learning the appearance of the body (as a function of its color/texture, shape or motion flow) in the corresponding pose.

Murase and Sakai [27] describe a template matching method which uses the parametric eigenspace representation as applied in face recognition [35]. Specifically, they use PCA (Principal Component Analysis) to compute a 16-dimensional manifold for all the possible grey-scale images of a walking person. An input sequence of images (after normalization) is hence mapped to a trajectory in this 16-dimensional feature space, and gait recognition is achieved by computing the distance between the trajectories of the input image sequence and a reference sequence.

Huang et al. [20] use a similar technique, as they apply PCA to map the binary silhouette of the moving figure to a low dimensional feature space. The gait of an individual person is represented

as a cluster (of silhouettes) in this space, and gait recognition is done by determining if all the input silhouettes belong to this cluster.

He and Debrunner [16] recognize individual gaits via an HMM that uses the quantized vector of Hu moments of a moving person's silhouette as input.

Spatiotemporal Methods Here, the action or motion is characterized via the entire 3D spatiotemporal (XYT) data volume spanned by the moving person in the image. It could for example consist of the sequence of grey-scale images, optical flow images, or binary silhouettes of the person. This volume is hence treated as a 'large' vector, and motion recognition is typically done by mapping this vector to a low-dimensional feature vector, and applying standard pattern classification technique in this space. The following methods describe different ways of doing this.

Of particular interest is the recent work by Cutler and Davis [10], in which they show that human motion is characterized by specific periodic patterns in the *similarity plot* (a 2D matrix of all pairwise image matching correlations), and describe a method for human detection by recognizing such patterns. They also use similarity plots to estimate the stride of a walking and running person, assuming a calibrated camera. Here, a person is tracked over a ground plane, and their distance travelled, D , is estimated. The number of steps N is also automatically estimated using periodic motion, which can be a non-integer number. The stride is D/N , which could be used as a biometric, though they have not conducted any study showing how useful it is as a biometric.

3 Motivation

In this section, we give the motivation for the methodology used in this paper. One method of motion-based recognition is to first explicitly extract the dynamics of points on a moving object. Consider a point $\mathbf{P}(t) = (x(t), y(t), z(t))$ on a moving object as a function of time t (see Figure 1). The dynamics of the point can be represented by the phase plot $(\mathbf{P}(t), d\mathbf{P}/dt(t), \dots)$. Since we wish to recognize different types of motions (viz. gaits), it is important to know what can be determined from the projection T of $\mathbf{P}(t)$ onto an image plane, $(u, v) = T(\mathbf{P})$. Under orthographic projection, and if $\mathbf{P}(t)$ is constrained to planar motion, the object dynamics are completely preserved up to a scalar factor. That is, the phase space for the point constructed from (u, v) is identical (up to a scalar factor) to the phase space constructed from $\mathbf{P}(t)$. However, if the motion is not constrained to a plane, then the dynamics are not preserved. Under perspective projection, the dynamics of planar and arbitrary motion are in general not preserved.

Fortunately, planar motion is an important class of motion, and includes "biological motion" [17]. In addition, if the object is sufficiently far from the camera, the camera projection becomes approximately orthographic (with scaling). In this case, and assuming we can accurately track a point $\mathbf{P}(t)$ in the image plane, then we can completely reconstruct the phase space of the dynamic system (up to a scalar factor). The phase space can then be used directly to classify the object motion (e.g., [7]).




Fig. 1. Point $P(t)$ on an object moving in R^3 , imaged onto a 2D plane.

In general, point correspondence is not always possible in realistic image sequences (without the use of special markers), due to occlusion boundaries, lighting changes, insufficient texture, image noise, etc. However, for classifying motions, we do not necessarily have to extract the complete dynamics of the system; qualitative measures may suffice to distinguish a class of motions from each other. In this paper, we use correspondence-free, qualitative measures for motion-based gait recognition.

4 Method

4.1 Self-similarity Plots

Computation from an Image Sequence Given a sequence of grey-scale images obtained from a static camera, we detect and track the moving person, extract an image template corresponding to the person’s motion blob in each frame, then compute the image self-similarity plot from the obtained sequence of templates. For this, we use the method described in [10], except that we use background modeling and subtraction [13, 19] for foreground detection, since the camera is assumed to be stationary.

Moving objects are tracked in each frame based on spatial and temporal image coherence. An image template at time t , denoted by O_t , is extracted for each tracked object, consisting of the image region enclosed within the bounding box of its motion blob in the current frame. Deciding whether a moving object corresponds to a walking person is currently done based on simple shape (such as aspect ratio of the bounding box and blob size) and periodicity cues.

Once a person has been tracked for N consecutive frames, its N image templates are scaled to the same dimensions $H \times W$, as their sizes may vary due to change in camera viewpoint and segmentation errors. The image self-similarity, S , of the person is then computed as follows:

$$S(t_1, t_2) = \sum_{(x,y) \in B_{t_1}} |O_{t_1}(x, y) - O_{t_2}(x, y)|,$$

where $1 \leq t_1, t_2 \leq N$, B_{t_1} is the bounding box of the person in frame t_1 , and $O_{t_1}, O_{t_2}, \dots, O_{t_N}$ are the scaled image templates of

the person. In order to account for tracking errors, we compute the minimal S' by translating over a small search radius r :

$$S'(t_1, t_2) = \min_{|dx, dy| < r} \sum_{(x,y) \in B_{t_1}} |O_{t_1}(x + dx, y + dy) - O_{t_2}(x, y)|.$$

Figure 2(b) shows a plot of S' for all combinations of t_1 and t_2 , for the two walking sequences (80 frames each) shown in Figure 2(a) (note the similarity values have been linearly scaled for visualization to the grayscale intensity range $[0, 255]$, where dark regions show more similarity).

Properties The similarity plot, S' , of a walking person has the following properties:

1. $S'(t, t) = 0$, i.e. it has a dark main diagonal.
2. $S'(t_1, t_2) = S'(t_2, t_1)$, i.e. it is symmetric along the main diagonal.
3. $S'(t_1, kp/2 + t_1) \simeq 0$, i.e. it has dark lines *parallel* to the main diagonal.
4. $S'(t_1, kp/2 - t_1) \simeq 0$, i.e. it has dark lines *perpendicular* to the main diagonal.

where $t_1, t_2 \in [1, N]$, p is the period of walking, and k is an integer. The first two properties are generally true for any similarity function (though if substantial image scaling is required, the second property may not hold). The latter two, however, are a direct consequence of the periodicity and the bilateral symmetry, respectively, of the human gait.

Figure 2(a) shows the 5 key poses over one walking cycle for two different persons; poses A and C correspond to when the person is in maximum swing (i.e. the two legs are furthest apart), and Pose B corresponds to when the two legs are together. One can easily see that the intersections of the dark lines in S' (i.e. the diagonals and cross diagonals) correspond to pose combinations AA, BB and CC, AC , and CA . Thus these intersections, which are the local minima of S' , can be used to determine the frequency and phase of walking [10].

That S' encodes the frequency and phase of gait can be explained by the fact that the similarity plot of a walking person is (approximately) a projection of the planar dynamics of the walking person when viewed sufficiently far from the camera, as previously suggested in [10]. Intuitively, this is because S' is obtained via a sequence of transformations (image projection and template matching) applied to the set of 3D points constituting the person’s body. It can be shown that these transformations approximately preserve the dynamics of these points (and hence the dynamics of the gait) under certain assumptions.

4.2 Gait Classifier

As mentioned in the previous section, the similarity plot is a projection of the dynamics of the walking person that preserves the frequency and phase of the gait. The question then arises as to whether this projection preserves more detailed (higher-dimensional) aspects of gait dynamics, that capture the unique way a person walks. In other words, does a similarity plot contain




Fig. 2. (a) A few frames of the key poses in a walking person sequence, and the corresponding (b) image self-similarity plots.

sufficient information to distinguish (not necessarily uniquely) the walking gaits of different people?

To evaluate the usefulness of the self-similarity plot in characterizing and recognizing individual gaits, we propose to build a gait pattern classifier that takes an SP (self-similarity plot) as the input feature vector. For this, we take an ‘eigenface’ approach [35], in which we treat a similarity plot the same way that a face image is used in a face recognizer. The gist of this approach is that it extracts ‘relevant information’ from input feature vectors (face images or SPs) by finding the principal components of the distribution of the feature space, then applies standard pattern classification of new feature vectors in the lower-dimensional space spanned by the principal components. We use a simple non-parametric pattern classification technique for recognition. In the following, we explain the details of the proposed gait classifier.

Normalizing the Input In order to account for different walking paces and starting poses, we need to normalize the self-similarity plots so that they are phase-aligned, have the same frequency, and contain the same number of cycles. To this end, we compute the frequency and phase of each similarity plot using the method in [10]. We choose the pose corresponding to when the legs are maximally apart, i.e. poses A or C in Figure 2(a), for phase alignment.

Training the Classifier Let S'_1, S'_1, \dots, S'_M be a given training set of M labelled (i.e. corresponding to a known person) normalized similarity plots, of size $N \times N$ each, and let s'_i be the vector of length N^2 corresponding to the i th similarity plot S'_i (obtained by concatenating all its rows). We compute the principal components [22] of the space spanned by s'_1, \dots, s'_M by computing the eigenvalue decomposition (also called Karhunen-Loeve expansion) of their covariance matrix:

$$C_s = \frac{1}{M} \sum_{i=1}^M (s'_i - \bar{s}') (s'_i - \bar{s}')^T$$

where \bar{s}' is the simple mean of all training vectors s'_1, \dots, s'_M . This can be efficiently computed in $O(M)$ time (instead of the brute force $O(N^2)$) [35].

We then consider the space spanned by the n most significant eigenvectors, u_1, \dots, u_n , that account for 90% of the variation in the training SPs ¹. We denote this space the *Eigengait*. Hence each training vector s'_i can be sufficiently approximated by a n -dimensional vector w_i obtained by projecting it onto the Eigengait, i.e. $w_i \equiv \sum_{j=1}^n u_j^T \cdot s'_i$. Furthermore, assuming that the training vectors are representative of the variation in the entire feature space, then any new feature vector can be similarly approximated by a point in Eigengait space.

Classification Gait recognition now reduces to a standard pattern classification in a n -dimensional Eigengait space. The advantage of doing pattern classification in this space is not only that n is typically much smaller than N^2 and M , but also that it contains less unwanted variation (i.e. random noise)² and hence provides better separability of the feature vectors, or SPs.

Given a new SP (corresponding to an unknown person), the procedure for recognizing it is to first convert it to a N^2 -vector, map it to a point in Eigengait, find the k closest training points to it, then decide its class (or label) via the *k-nearest neighbor rule* [5, 31].

5 Experiments

To evaluate our method, we build the gait classifier described above using k -nearest neighbor classification and the gait data set of Little and Boyd [23]. We use the leave-one-out cross-validation to obtain a statistically accurate estimate of the recognition rate [36, 31].

The data set consists of 40 image sequences and six different subjects (7 sequences per person except for the 5th person). Figure 3 shows all six subjects overlaid at once on the background image.

Since the camera is static we used median filtering to recover the background image. Templates of the moving person were extracted from each image by computing the difference of the image and the background and subsequently applying a threshold as well as morphological operations to clean up noise. This simple method for person detection and tracking was sufficient because the background is static and each sequence only contains one moving person. Absolute correlation was used to compute the self-similarity plot for each of the 40 template sequences. For temporal normalization (since the image sequences were of varying lengths and the

¹ According to the theory of PCA, if $\lambda_1, \dots, \lambda_n$ are the n largest eigenvalues, then the space spanned by their corresponding eigenvectors account for $\sum_{i=1}^n \lambda_i / \text{trace}(C_s)$ of the total variation in the original feature vectors.

² Assuming data variation is much larger than noise variation.

persons had different gait cycle lengths), the similarity plots were cropped and scaled so that they all contained 4 gait cycles starting on the same phase, and are of size 64x64. Figure 4 shows examples of these normalized similarity plots, where each column of three plots corresponds to one person.

Since our data set is relatively small, we use the leave-one-out cross-validation method. The leave-one-out error rate estimator is known to be an (almost) unbiased estimator of the true error rate of the classifier. Hence, out of the 40 similarity plots, we build (or train) our classifier on all but one of the samples, test the classifier on the sample missed (or left out), and record the classification result. This is repeated 40 times, leaving out each of the 40 samples in turn. The recognition rate is then obtained as the ratio of the number of correctly classified test samples out of the total 40.

The classifier is built simply by storing the training vectors as points in Eigengait space, and the test sample is classified by determining its k -nearest neighbor with $k = 5$, using the Euclidian distance as a distance metric and simple majority as a decision rule. The recognition rate thus obtained is 0.93 (37 out of 40).

To visualize how well separated the gaits of the 6 people are, we applied PCA to all 40 SPs. Figure 5 shows all 40 SPs projected onto the 3 most significant eigenvectors thus obtained. Each closed contour encloses the samples (points) corresponding to one person.



Fig. 3. The six people contained in the test sequences, overlaid on the background image.




Fig. 4. Normalized self-similarity plots (columns correspond to a single person).




Fig. 5. Similarity plots projected onto the space spanned by the three most dominant eigenvectors.

6 Conclusion and future work

In this paper, we have used a correspondence-free motion-based method to recognize the gaits of a small population (6) of people. While the results are promising, more evaluation of the method needs to be done. Future studies include larger test populations (20-100 people) and images taken from multiple view points (not just parallel to the image plane)³. In addition, image sequences of the same individual need to be acquired in different lighting conditions, and with various types of clothing.

Finally, we are working to combine the results of this correspondence-free gait recognition method with more feature-oriented methods, such as stride and height estimation.

Acknowledgment

We would like to thank Dr. Jeffrey E. Boyd of the Department of Computer Science at the University of Calgary, Canada, for providing the gait data we used in testing. The support of DARPA (Human ID project, grant No. 5-28944) is also gratefully acknowledged.

References

1. J. K. Aggarwal and Q. Cai, "Human motion analysis: a review," in *Proc. of IEEE Computer Society Workshop on Motion of Non-Rigid and Articulated Objects*, 1997.
2. K. Akita, "Image Sequence Analysis of Real World Human Motion," *Pattern Recognition*, Vol. 17, No. 1, pp. 73–8, 1984.
3. C. Barclay, J. Cutting, and L. Kozlowski, "Temporal and Spatial Factors in Gait Perception that Influence Gender Recognition," *Perception and Psychophysics*, Vol. 23, No. 2, pp. 145–152, 1978.

³ The Keck Laboratory [6] will be utilized to acquire multiview (up to 64) images sequences of a person walking.

4. J. Bigun, G. Chollet, and G. Borgefors, *Audio- and Video-based Biometric Person Authentication*, Springer, 1997.
5. C. Bishop, *Neural Networks for Pattern Recognition*, Oxford: Clarendon Press, 1995.
6. E. Borovikov, R. Cutler, T. Horprasert, and L. Davis, "Multi-perspective Analysis of Human Actions," in *Third International Workshop on Cooperative Distributed Vision*, 1999.
7. L. W. Campbell and A. Bobick, "Recognition of Human Body Motion Using Phase Space Constraints," in *International Conference on Computer Vision*, 1995.
8. C. Cedras and M. Shah, "A survey of motion analysis from moving light displays," in *Proceedings of the Computer Vision and Pattern Recognition*, pp. 214–221, 1994.
9. D. Cunado, M. Nixon, and J. Carter, "Using Gait as a Biometric, via Phase-Weighted Magnitude Spectra," in *Proceedings of 1st Int. Conf. on Audio- and Video-Based Biometric Person Authentication*, pp. 95–102, 1997.
10. R. Cutler and L. Davis, "Robust Real-time Periodic Motion Detection, Analysis and Applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 2, pp. 129–155, 2000.
11. J. Cutting and L. Kozlowski, "Recognizing Friends by Their Walk: Gait Perception Without Familiarity Cues," *Bulletin Psychonomic Soc.*, Vol. 9, No. 5, pp. 353–356, 1977.
12. J. W. Davis and A. F. Bobick, "The representation and recognition of action using temporal templates," in *Proceedings of the Computer Vision and Pattern Recognition*, pp. 928–934, 1997.
13. A. Elgammal, D. Harwood, and L. Davis, "Non-parametric model for background subtraction," in *Proceedings of International Conference on Computer Vision*, 2000.
14. D. Gavrilu, "The visual analysis of human movement: a survey," *Computer Vision and Image Understanding*, Vol. 73, pp. 82–98, January 1999.
15. D. M. Gavrilu and L. Davis, "Towards 3-D Model-based Tracking and Recognition of Human Movement: a Multi-View Approach," in *IEEE International Conference on Automatic Face and Gesture Recognition*, (Zurich, Switzerland), 1995.
16. Q. He and C. Debrunner, "Individual Recognition from Periodic Activity Using Hidden Markov Models," in *IEEE Workshop on Human Motion*, 2000.
17. D. Hoffman and B. Flinchbaugh, "The interpretation of biological motion," *Biological Cybernetics*, 1982.
18. D. Hogg, "Model-based vision: a program to see a walking person," *Image and Vision Computing*, Vol. 1, No. 1, 1983.
19. T. Horprasert, D. Harwood, and L. Davis, "A Robust Background Subtraction and Shadow Detection," in *Asian Conference on Computer Vision*, 2000.
20. P. S. Huang, C. J. Harris, and M. S. Nixon, "Comparing Different Template Features for Recognizing People by their Gait," in *BMVC*, 1998.
21. G. Johansson, "Visual Motion Perception," *Scientific American*, pp. 75–88, June 1975.
22. I. T. Joliffe, *Principal Component Analysis*, Springer-Verlag, 1986.
23. J. Little and J. Boyd, "Recognizing people by their gait: the shape of motion," *Videre*, Vol. 1, No. 2, 1998.
24. F. Liu and R. Picard, "Finding periodicity in space and time," *International Conference on Computer Vision*, pp. 376–383, January 1998.
25. K. Luttgens and K. Wells, *Kinesiology: Scientific Basis of Human Motion*, Saunders College Publishing, 7th ed., 1982.
26. D. Meyer, J. Psl, and H. Niemann, "Gait Classification with HMMs for Trajectories of Body Parts Extracted by Mixture Densities," in *BMVC*, pp. 459–468, 1998.
27. H. Murase and R. Sakai, "Moving object recognition in eigenspace representation: gait analysis and lip reading," *Pattern Recognition Letters*, Vol. 17, pp. 155–162, 1996.
28. M. Murray, "Gait as a total pattern of movement," *American Journal of Physical Medicine*, Vol. 46, No. 1, pp. 290–332, 1967.
29. S. Niyogi and E. Adelson, "Analyzing and recognizing walking figures in XYT," in *Proceedings of the Computer Vision and Pattern Recognition*, pp. 469–474, 1994.
30. R. Polana and R. Nelson, "Detection and Recognition of Periodic, Non-rigid Motion," *International Journal of Computer Vision*, Vol. 23, pp. 261–282, June/July 1997.
31. B. Ripley, *Pattern Recognition and Neural Networks*, Cambridge: Cambridge University Press, 1996.
32. K. Rohr, "Towards model-based recognition of human movements in image sequences," in *CVGIP*, vol. 59, 1994.
33. Y. Song, X. Feng, and P. Perona, "Towards Detection of Human Motion," in *Proceedings of the Computer Vision and Pattern Recognition*, 2000.
34. P. Tsai, M. Shah, K. Keiter, and T. Kasparis, "Cyclic Motion Detection for Motion-based Recognition," *Pattern Recognition*, Vol. 27, No. 12, pp. 1591–1603, 1994.
35. M. Turk and A. Pentland, "Face Recognition using Eigenfaces," in *Proceedings of the Computer Vision and Pattern Recognition*, 1991.
36. S. Weiss and C. Kulikowski, *Computer Systems that Learn*, Morgan Kaufman, 1991.