



ELSEVIER

Available at  
[www.ComputerScienceWeb.com](http://www.ComputerScienceWeb.com)  
POWERED BY SCIENCE @ DIRECT®

Pattern Recognition Letters 24 (2003) 2489–2497

---

---

Pattern Recognition  
Letters

---

---

[www.elsevier.com/locate/patrec](http://www.elsevier.com/locate/patrec)

# Automatic gait recognition using area-based metrics

Jeff P. Foster, Mark S. Nixon<sup>\*</sup>, Adam Prügel-Bennett

*Department of Electronics and Computer Science, ISIS, University of Southampton, Southampton, SO17 1BJ, UK*

Received 12 July 2002; received in revised form 17 March 2003

---

## Abstract

A novel technique for analysing moving shapes is presented in an example application to automatic gait recognition. The technique uses masking functions to measure area as a time varying signal from a sequence of silhouettes of a walking subject. Essentially, this combines the simplicity of a baseline area measure with the specificity of the selected (masked) area. The dynamic temporal signal is used as a signature for automatic gait recognition. The approach is tested on the largest extant gait database, consisting of 114 subjects (filmed under laboratory conditions). Though individual masks have limited discriminatory ability, a correct classification rate of over 75% was achieved by combining information from different area masks. Knowledge of the leg with which the subject starts a gait cycle is shown to improve the recognition rate from individual masks, but has little influence on the recognition rate achieved from combining masks. Finally, this technique is used to attempt to discriminate between male and female subjects. The technique is presented in basic form: future work can improve implementation factors such as using better data fusion and classifiers with potential to increase discriminatory capability.

© 2003 Elsevier B.V. All rights reserved.

*Keywords:* Biometrics; Gait recognition; Moving object analysis; Moving object description

---

## 1. Introduction

In the last couple of years, gait has emerged as a potential practical biometric. It poses considerable challenges compared with traditional biometrics such as face, iris or speech in that it is a function of both space and time. The information available from a gait signature is thus tremendous, but this poses new problems due to the overwhelming amount of data. We shall consider gait as a se-

quence of images of a single subject walking across the plane of view of a stationary camera.

To reduce the dimensionality of the data we only compare a sequence of windowed silhouettes of the walker. The procedure for obtaining the silhouettes is discussed in the following section. The set of silhouettes obtained over the walking cycle provides more information (than single images), even if we do not consider them as part of a sequence (Huang et al., 1999a; Hayfron-Acquah et al., 2001).

Previous work on gait recognition can be divided into two distinct areas. Holistic approaches aim to extract statistical features from a subject's silhouette to differentiate between subjects

---

<sup>\*</sup> Corresponding author. Tel.: +44-2380-593542; fax: +44-2380-594498.

*E-mail address:* [msn@ecs.soton.ac.uk](mailto:msn@ecs.soton.ac.uk) (M.S. Nixon).

(Huang et al., 1999a; Hayfron-Acquah et al., 2001; Little and Boyd, 1998; Niyogi and Adelson, 1994; BenAbdelkader et al., 2002a; Collins et al., 2002; Phillips et al., 2002; Murase and Sakai, 1996) and model-based approaches which aim to model human gait explicitly (Yam et al., 2002; Cunado et al., 1999a). One early approach (Murase and Sakai, 1996) used PCA applied to silhouettes and used the trajectory information for classification. This was later extended using canonical analysis (Huang et al., 1999a), again showing encouraging results. Notwithstanding these results, classification becomes complex with large datasets and naturally depends on the exemplars used as training data since there is nothing specific to gait in the descriptive basis (further, the silhouette sequence could even be re-ordered without affecting the result). Recent baseline approaches have used matching of key gait events, this being the sole focus of specific relation to gait. Of these two ‘baseline’ approaches one (Collins et al., 2002) captures body as well as gait cues by comparing the silhouette at key gait events with other silhouettes stored in a database. Alternatively (Phillips et al., 2002), a correlation can be calculated by determining the number of intersecting pixels between the sequence of silhouettes of the test and of a reference subject. As with other emerging papers, the studies report encouraging results on a selection of databases ranging from around 30 to just over 70 subjects.

The disadvantage of traditional holistic approaches is that they ignore the temporal component of gait. Subjects’ gait could be affected by carried items, but there is no specificity that allows for any analysis of this (except, naturally, construction of the analysis database). In contrast, model-based approaches (Yam et al., 2002, Cunado et al., 1999b) aim to model accurately how a subject walks by analysis of the motion of the legs. The disadvantage of model-based approaches is typically the computational complexity. The focus of this paper will be to present a new holistic technique which considers temporal information that is intimately related to the nature of gait, with low computational cost.

This new technique uses a masking function to determine area and hence captures a low dimen-

sional temporal signature of the subject. We use fixed templates which allows us to test for particular invariances (e.g. using information from only the legs). There are also speed advantages to our new approach to gait-signature generation, over the speed of established statistical techniques. Standard pattern recognition techniques are used to classify subjects based on these signatures. We present results using the new SOTON database that consists of 114 subjects with a minimum of eight samples for each subject (Shutler et al., 2002). This database is much larger than those available previously. The database consists of sequences of images collected under controlled conditions. Subjects are filmed walking normally in front of a green background and chroma-key techniques are used to isolate the moving subject. Following this, background subtraction and a connected components algorithm are used to extract a binary silhouette of the walking subject. Finally, the silhouette is windowed and resized to a  $64 \times 64$  silhouette. Detailed information about the SOTON database is available in (Shutler et al., 2002). This data is aimed specifically to allow us to answer the question of gaits’ potency as a biometric, rather than in a “real-world”/computer vision scenario.

We show that there is sufficient information in these dynamical signatures for recognition, with a 75% correct classification rate on the SOTON database. To explore in more detail where this performance arises, we present results for experiments where we have removed all DC components from our dynamical signatures so the remaining information is purely temporal. We also report on studies where we have separated the samples of walkers moving left to right from those moving in the opposite direction. Furthermore, we consider learning from sequences which have been synchronised, so that all walks start on the same foot. Using this additional knowledge improves the recognition performance from a single mask, but does not improve performance when masks are combined. Finally, we report on an attempt to differentiate between male and female subjects based on these dynamic signatures. Preliminary results using this approach have been reported elsewhere (Foster et al., 2001a,b) but this paper

contains the first complete analysis of this algorithm on a large database.

## 2. Extracting an area dynamic signature

Before a dynamic signature can be obtained, the silhouette of the subject must be separated from the background. In part of the SOTON database, the gait sequences are derived in the laboratory, aiming for near-perfect conditions to obtain the best possible silhouette. Subjects are viewed fronto-parallel walking along a track in either direction. The camcorder's digital video is turned into individual frames. Fig. 1(a) shows the original image from which the silhouette is extracted using chroma-key subtraction in conjunction with a connected components algorithm. The silhouette is then windowed to remove those parts of the background away from the chroma-key, such as the roof, Fig. 1(b), and the resized to a  $64 \times 64$  image, Fig. 1(c). Examples of silhouettes of different subjects are shown in Fig. 2. Subjects are constrained to walking normal to the camera's plane of view as we are interested in analysing the basic properties of gait as a biometric. Future work can concentrate on the consequences of a subject walking at an angle to the camera.

Binary silhouettes remove information about colouring and interior shape but give invariance to clothing colour and most lighting conditions. Furthermore, by normalising the size of the silhouettes we remove variation due to distance from the camera, although we lose information about the subjects' height.

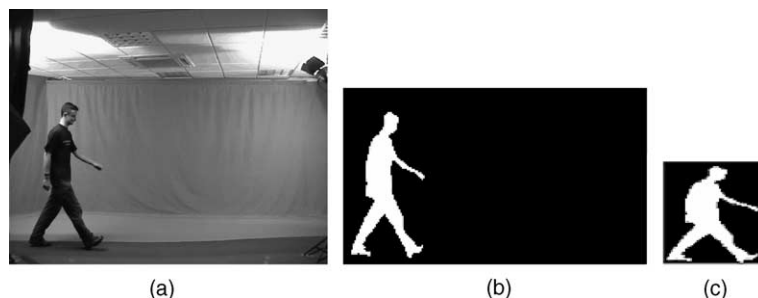


Fig. 1. Silhouette extraction process: (a) original image; (b) silhouette and (c) windowed silhouette.



Fig. 2. Sample windowed silhouettes.

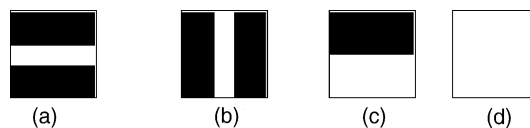


Fig. 3. Sample area masks: (a) horizontal line; (b) vertical line; (c) bottom half and (d) full.

Following this process, we are left with a sequence of normalised and centralised silhouettes for each subject. We use multiple sequences containing a single walking cycle (two steps) to perform recognition. This typically comprises of approximately filmed at 25 frames/s. We denote the images in the sequence by a vector  $x(t)$  where  $t$  labels the time when the image was taken.

To obtain the dynamic signature we compute the area of the silhouette within a mask. These masks are chosen intuitively to isolate a particular part of the image. Examples of area masks are shown in Fig. 3. Here, a horizontal line (Fig. 3(a)) isolates those parts in the region of the waist whereas a vertical one, Fig. 3(b), selects the thorax and these parts of the legs intersecting with the vertical window. Many alternatives are possible such as a combination of the upper body and legs, Fig. 3(c) merely measures the entire area change of

the image Fig. 3(d) simply analyses the legs of the subject.

Each mask,  $m_j$ , represents an area mask shown in Fig. 3 as a binary vector, with 1 representing the white parts of the image. For each mask, we obtain a signature by determining the area between the mask and the images thus

$$S_j(t) = m_j^T x(t)$$

We obtain an area history by applying each mask. Note that all the masks used are left right symmetric masks and we therefore lose all information corresponding to the asymmetric part of gait. However, this also gives invariance to the direction of walk. That is, we could mirror any image within our sequence and we would obtain the same gait signature. We could use an asymmetric masks which might give us additional information, but we would then need to reflect all walkers to walk in the same direction.

Examples of the gait sequences  $S_j(t)$  for three different masks are shown in Fig. 4 (the full area mask gives a similar output, but is not shown as the details are harder to see when plotted on a scale showing all four functions). As can be seen, the dynamic signature is intimately related to gait. The peaks in the graph represent when the selected area of the subject is at a maximum (i.e. when the legs are furthest apart) and the dips represent when the selected area is at a minimum (i.e. when the

legs are closest together). The bottom half mask selects the lower portion of the image so has a much greater number of points than the vertical line, but the shape is naturally very similar. The horizontal line mask selects those points near the region of the waist which varies less with thorax motion than the associated movement of the subject's arms. Clearly, the masks emphasise bilateral symmetry: one leg swings whilst the other is in stance after which the limbs swap in function. Limping is an inherently asymmetrical gait, and this would be evidenced in a clear disparity in measured area between the two halves of the gait cycle.

To compare sequences that start at different points in the gait cycle we have to align the sequences. We can do this by determining the minima in a sequence and using this as a starting point. To eliminate variations due to the sampling time and speed at which the subject walks we select a part of the sequence corresponding to the full walking cycle and resample using cubic splines to interpolate between the observations  $S_j(t)$ . A cubic spline curve is fitted for the whole gait cycle, and 30 evenly spaced samples are taken from the whole curve, giving a 30 element vector for each area mask used. Note that this does lose information about the subject's speed. BenAbdelkader et al. (2002b) has shown that this information can be used for recognition, however, our technique is aimed at extracting the dynamics of area change and thus the speed of the walker can be removed.

Each mask yields a vector describing the dynamics of area change within that mask. However, there are many degrees of freedom in a walking body, so we would expect that there would be a considerable degree of independent information from different masks. Therefore, we have combined the information from multiple masks to provide a more complete dynamics signature. We have chosen the simplest way to do this, namely to concatenate the 30 element vector from each area mask to form a single long vector of  $n \times 30$ , where  $n$  denotes the number of area masks used. We denote this dynamic signature by  $v$ . We will have a different vector  $v$  for each sequence of images. We use these vectors to perform the recognition of subjects in our database. Note that the vectors still contain some information about the static shape of

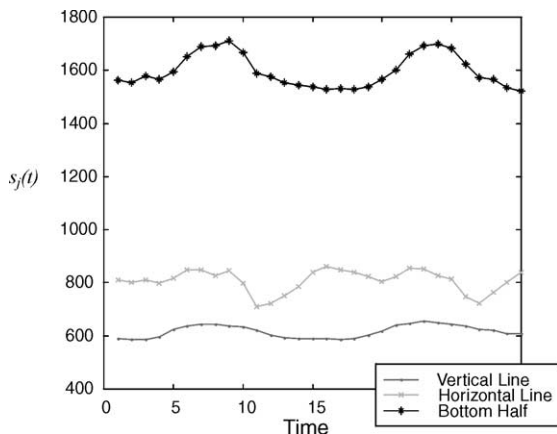


Fig. 4. Sample output for three different area masks from a single subject.

the silhouette. This information is contained in the average value of sequence  $S_i(t)$ . We can remove this information by subtracting the average of the sequence, although this reduces the recognition performance. We return to this issue in the next section.

For each walking sequence, we have obtained an  $n \times 30$  element vector  $\mathbf{v}$  that provides a dynamic signature describing the gait. To classify the signatures we used canonical analysis (Fisher, 1938; Fukunaga, 1990; Swets and Weng, 1996) to select a low dimensional subspace where the differences between classes are maximised and the difference within classes is minimised. For completeness we describe the canonical analysis procedure in the appendix. The centroids of each class are calculated in this new space and a  $k$ -nearest neighbour classifier is used to decide which subject a sequence of images belong to.

### 3. Results

#### 3.1. Person identification

Table 1 shows the recognition rates for single area masks on the SOTON database. Six samples of each subject (three left to right, three right to left) were used to train the database and two samples (one left to right, one right to left) were used as the test data.

As can be seen from the results in Table 1, the performance of the system using only a single area mask is unimpressive. It should be noted however, that the results are significantly better than the chance recognition rate which is less than 1% on a database of this size.

All results are reported in terms of recognition rate with an estimate of the expected error in that

rate. The tests consisted of  $N$  independent measurements leading to a binary outcome, i.e. the results follow a distribution for Benoulli trials. The probability of obtaining  $M$  correct results is given by the standard binomial relationship (DeGroot and Schervish, 2001)

$$P(M) = \binom{N}{M} p^M (1-p)^{N-M}$$

where  $p$  is the true recognition rate. The mean and variance for  $M$  are  $Np$  and  $Np(1-p)$ , respectively. The natural unbiased estimate for the recognition rate is  $M/N$  and the expected error in the unbiased estimate is  $\sqrt{p(1-p)/N}$ . The observed recognition rate can be used as an approximation to  $p$ . For our database of 114 subjects with two test samples of each,  $N = 228$  so the expected error in the mean is around 3.0%, as shown in Table 1.

From Table 1, there is no significant difference in the horizontal line, bottom half and full masks. However, the vertical line mask appears to have significantly less discriminatory power than either of the others, and its estimated error is then proportionally greater than for the other masks.

The advantage of this approach is that the execution speed is very fast, especially compared with other extant approaches such as Huang et al. (1999b). Since these results are substantially higher than chance, we must assume that there is potential to using area masks. As such we looked at increasing the recognition performance of the system whilst preserving the advantage of fast computation time.

#### 3.2. Combining area masks to increase performance

By combining the results from area masks as discussed previously and using canonical analysis in the same manner as before the recognition rate increases markedly. Results are combined by simply concatenating vectors together to produce a vector of size  $n \times 30$  where  $n$  denotes the number of area masks used (Table 2). By combining masks we greatly increase the information available for discriminant analysis and thus the results improve. The achieved recognition rate of over 75% is encouraging and significantly greater than chance alone.

Table 1  
Recognition rates from single masks

Area mask	Recognition rate (%)	Expected error, $\sigma_e$ (%)
Horizontal line mask	29.9	3.03
Vertical line mask	16.1	2.43
Bottom half mask	27.2	2.95
Full mask	29.2	3.01

Table 2  
Results from combining multiple masks

Area masks combined	Recognition results (%)	Expected error, $\sigma_e$ (%)
Horizontal line mask + vertical line mask	49.6	3.31
Bottom half + full mask	50.0	3.31
All above masks	75.4	2.85

### 3.3. Using only the temporal components

Gait differs from traditional biometrics because it contains a temporal as well as a spatial component. The spatial component represents the shape of the person, whilst the temporal component represents the movement of the body over time. Psychological research (Johansson, 1973) would seem to suggest that humans can recognise movement patterns merely from the temporal component.

Using only the DC component of the area masks' output, produces poor recognition results. Using the DC component of the vertical line mask produces a recognition rate of 7.5%.

The full mask component represents the spatial component of the gait sequence. As can be seen from Table 1, the recognition rate using only the spatial component is approximately 30%.

To determine where the recognition capability was derived, we repeated the tests but this time removing the DC component of gait sequence. This was accomplished by subtracting the mean of each sequence from the data. Thus, the sequence only contains temporal information.

As can be seen from Table 3 the results are similar in levels of performance to those in Table 1. However, we can see that the recognition rates using the horizontal line mask has dropped by almost 20% to less than 10%. As expected, this shows that the horizontal mask contains very little temporal information which can be used for recognition. By combining the masks in the manner described previously, it is possible to increase this recognition rate. Combining the four masks produced a recognition rate of 52.7%. This suggests that approximately half of the recognition rate come from the temporal components of gait, with the remaining coming from the shape of the silhouettes.

Table 3  
Results when DC component is removed

Area mask	Recognition rate (%)	Expected error, $\sigma_e$ (%)
Horizontal line mask	8.9	1.89
Vertical line mask	26.3	2.92
Bottom half mask	23.7	2.81
Full mask	24.1	2.83

### 3.4. Is gait symmetric?

The psychologists' view is that gait is a symmetrical pattern of motion (Cutting et al., 1978; Sadeghi et al., 2000). We have assumed this to be true, and not taken into account the foot on which the subject starts, or their direction of travel. Does taking this information into account result in a significant difference in recognition rate?

Table 4 shows the recognition rates for each combinations of starting leg and walk direction. Four area masks were combined (vertical line, bottom half, horizontal line and full) and leave one out cross-validation was used to produce the recognition rates shown. As can be seen recognition rates are approximately constant for each of the combinations. This was to be expected as all subjects are able-bodied and none showed a distinct impediment in their walk, to the human eye.

Analysing the results in more detail indicates that if knowledge of the starting leg is used, then performance can be improved. Table 5 shows the recognition rates of individual masks on a data set where all the subjects are walking right to left and starting on the left foot.

If we compare the results in Table 5 to those in Table 1 we can see there is a marked increase in the recognition rate when the starting leg and direc-

Table 4  
Recognition rates for different starting legs and directions using four area masks

Direction of walk	Starting left	Recognition rate (%)	Expected error, $\sigma_e$ (%)
Left	Left	79.8	3.76
Left	Right	76.6	3.96
Right	Left	76.6	3.96
Right	Right	76.5	3.97

Table 5  
Recognition rates using various masks when direction and starting leg is the same

Mask used	Recognition rate (%)	Expected error, $\sigma_e$ (%)
Horizontal line mask	52.2	4.67
Vertical line mask	37.7	4.54
Bottom half mask	49.6	4.68
Full mask	45.0	4.66

Table 6  
Recognition rates by combining area masks when direction and starting leg is the same

Area masks combined	Recognition results (%)	S.D. of expected error, $\sigma_e$ (%)
Horizontal line mask + vertical line mask	67.7	4.40
Bottom half + full mask	61.9	4.55
All masks	76.6	3.97

tion of travel is taken into account. Table 6 shows the results on the same data set when the masks are combined.

From the results in Table 6 we can see that whilst knowledge of the starting leg and direction of walk can increase performance when single masks are used, it does not provide significant advantages when masks are combined. Further research is necessary to determine if knowing the starting leg is a substantial advantage to the potency of gait as a biometric. Essentially this poses a question concerning the latent symmetry of gait which is yet to be determined by automatic analysis.

### 3.5. Gender classification

The SOTON database consists of 114 subjects, of which 20 are female. Gender discrimination could be a practical first stage in gait classification if the number of subjects is large, as it would reduce the number of subjects needed to be searched. To avoid bias in the test sample, 20 people of each sex were used, giving a total database size of 40. Fifteen people of each sex were used as training, with the remainder as test subjects. In this case, eight sequences were used for each. This gave 120 training sequences and 40 test sequences.

By using canonical analysis, the maximal gender discrimination rate was 64% which is better than chance, but not by much. Whilst results on gender classification are disappointing, they are justifiable. With the subject walking normal to the camera's plane of view there is very little gender discrimination information available. In mitigation, the clothing of males is far more uniform than that of females, suggesting that diversity of apparel could increase the potential recognition rate for women. Previous medical work has actually indicated that male and female walkers differ in terms of lateral body sway, with males tending to swing their shoulders from side to side more than their hips, and females tending to swing their hips more than their shoulders (Mather and Murdoch, 1994; Kozlowski and Cutting, 1977). This sort of information might be better discerned from a fronto-normal or an overhead view. Further research is needed to determine whether area masks, using different views of the subject, will be able to fulfil the task of gender discrimination.

## 4. Discussion and conclusions

This paper demonstrates for the first time that there is sufficient information in the dynamical time signature of gait to be able to recognise people with a high accuracy (75%). It is clear that we have discarded much information that might be useful in recognition so that, in this respect, we have under estimated the potential for gait. On the other hand, the information that we have discarded gives as a high degree of invariance that makes our system robust to superficial changes (e.g. lighting conditions, distance from camera, clothing and walking speed). Of course, the tests used a highly controlled database and we would not expect the same performance on real world data without spending much more effort in pre-processing the images.

The area-based metrics are very simple to compute and consequently very fast, which gives them a significant advantage over many other gait recognition techniques such as Huang et al. (1999a) and Murase and Sakai (1996). As such it is a fast baseline metric for analysing moving shapes

with specificity to the target application. This is advantageous, as the amount of data that needs to be processed is typically huge. It would be feasible to build a real time system using area-based metrics. The most demanding step would be silhouette extraction. Given the silhouettes the other steps required to perform recognition using area-based metrics are easily automated.

There still remain a large number of open questions. These include how the recognition rate behaves when the angle of the walker relative to the camera varies and when the walker changes clothes. New databases are currently being developed to address these issues.

### Acknowledgements

We gratefully acknowledge partial support by the European Research Office of the US Army under Contract No. N68171-01-C-9002.

### Appendix A. Canonical analysis

Canonical analysis (also known as linear discriminant analysis) is a procedure for finding a low-dimensional subspace which gives a high discrimination between classes of data. It differs from principal component analysis in that it uses knowledge of the correct classification of the data points. It is therefore supervised as opposed to an unsupervised learning algorithm. We assume that we have a set of classes  $C$  and for each class  $c \in C$  we have a set of instances  $y_c$  consisting of vectors  $y_c(i)$ , where  $i$  labels the instances. The class means are given by

$$\mu_c = \frac{1}{|y_c|} \sum_{i \in y_c} y_c(i)$$

While the mean of all the data is

$$\mu = \frac{1}{|C|} \sum_{c \in C} \mu_c$$

The covariance within a class  $c$  is given by

$$V_c = \frac{1}{|y_c| - 1} \sum_{i \in y_c} (y_c(i) - \mu_c)(y_c(i) - \mu_c)^T$$

where we divide by  $|y_c| - 1$  to obtain an unbiased estimate for the covariance. The average within-class covariance is equal to

$$W = \frac{1}{|C|} \sum_{c \in C} V_c$$

The between-class covariance is given by

$$B = \frac{1}{|C| - 1} \sum_{c \in C} (\mu_c - \mu)(\mu_c - \mu)^T$$

We now wish to find set of orthogonal directions,  $z_i$ , that give a high variance between classes and a low within class variance. To achieve this we define the discriminant function

$$J(z) = \frac{z^T B z}{z^T W z}$$

The set of vectors which extremise  $J(z)$  are found by setting its gradient to zero

$$\nabla J(z) = \frac{1}{z^T W z} (B z - J(z) W z) = 0$$

The extremal vectors,  $z_i$ , therefore satisfy the generalised eigenvalue equation

$$B z_i = \lambda_i W z_i$$

where  $\lambda_i = J(z_i)$ . The solution of the eigenvalue equation with the largest eigenvalues (i.e. the largest  $J(z_i)$ ) are those that separate the classes most effectively. For a complete description on the practicalities of implementing canonical analysis see (Swets and Weng, 1996).

Canonical analysis reduces the dimensionality of the data. This can improve the generalisation performance of a classifier system if we have too little data.

### References

- BenAbdelkader, C., Cutting, J., Davis, L., 2002a. Motion-based recognition of people in eigen gait space. In: Proc. 5th IEEE Face Gesture Recognition, pp. 378–384.
- BenAbdelkader, C., Cutting, J., Davis, L., 2002b. Stride and cadence as a biometric in automatic person identification. In: Proc. IEEE Face Gesture Recognition, pp. 372–378.
- Collins, R., Gross, R., Shi, J., 2002. Silhouette-based human identification from body shape and gait. In: Proc. 5th IEEE Face Gesture Recognition, pp. 366–371.
- Canado, D., Nixon, M.S., Carter, J.N., 1999a. Automatic gait recognition via model-based evidence gathering. In: Proc.



- AutoID99: IEEE Workshop on Automated ID Technol., pp. 27–30.
- Cunado, D., Nash, J.M., Nixon, M.S., Carter, J.N. 1999b. Gait extraction and description by evidence-gathering. In: Proc. Audio Video-based Biometric Person Authentication, pp. 43–48.
- Cutting, J.T., Proffitt, D.R., Kozlowski, L.T., 1978. A biomechanical invariant for gait perception. *J. Exper. Psychol.: Human Percept. Perform.*, 357–372.
- DeGroot, M., Schervish, M., 2001. *Probability and Statistics*, third ed. Addison Wesley, MA.
- Fisher, R.A., 1938. The statistical utilization of multiple measurements. *Ann. Eugen.* 8, 376–386.
- Foster, J.P., Nixon, M.S., Prugel-Bennett A., 2001a. New area-based measures for gait recognition. In: Proc. Audio–video-based Biometric Person Authentication, pp. 312–317.
- Foster, J., Nixon, M., Prugel-Bennett, A. 2001b. New area based metrics for automatic gait recognition. In: Proc. British Machine Vision Conf. 2001, pp. 233–242.
- Fukunaga, K., 1990. *Introduction to Statistical Pattern Recognition*, second ed. Academic Press, New York.
- Hayfron-Acquah, J.B., Nixon, M.S., Carter, J.N. 2001. Recognising human and animal movement by symmetry. In: Proc. IEEE Internat. Conf. on Image Process., pp. 290–293.
- Huang, P.S., Harris, C.J., Nixon, M.S., 1999a. Recognising humans by gait via parametric canonical space. *Artif. Intell. Eng.* 13, 359–366.
- Huang, P.S., Harris, C.J., Nixon, M.S., 1999b. Human gait recognition in canonical space using temporal template. *IEEE Proc. Vision Image Signal Process.* 146, 93–100.
- Johansson, G., 1973. Visual perception of biological motion and a model for its analysis. *Percept. Psychophys.* 14, 201–211.
- Kozlowski, L.T., Cutting, J.E., 1977. Recognizing the sex of a walker from a dynamic point-light display. *Percept. Psychophys.* 21, 575–580.
- Little, J.J., Boyd, J.E., 1998. Recognising people by their gait: The shape of motion. *Videre, Internat. J. Comput. Vision* 14, 83–105.
- Mather, G., Murdoch, L., 1994. Gender discrimination in biological motion displays based on dynamic cues. *Proc. Roy. Soc.*, 273–279.
- Murase, H., Sakai, R., 1996. Moving object recognition in eigenspace representation: Gait analysis and lip reading. *Pattern Recognition Lett.* 17, 155–162.
- Niyogi, S.A., Adelson, E.H., 1994. Analyzing and recognizing walking figures in XYT. *Proc. Comput. Vision Pattern Recognition*, 469–474.
- Phillips, P.J., Sarkar, S., Robledo, I., Grother, P., Bowyer, K. 2002. Baseline results for the challenge problem of human ID using gait analysis. In: Proc. 5th Automat. Face Gesture Recognition, pp. 137–143.
- Sadeghi, H., Allard, P., Prince, F., Labelle, H., 2000. Symmetry and limb dominance in able bodied gait: A review. *Gait Posture* 12, 34–45.
- Shutler, J.D., Grant, M.G., Nixon, M.S., Carter, J.N., 2002. On a large sequence-based human gait database. In: Proc. 4th Internat. Conf. on Recent Advances in Soft Computing, Nottingham, UK (accepted).
- Swets, D.L., Weng, J.J., 1996. Using discriminant eigenfeatures for image retrieval. *IEEE Trans. Pattern Anal. Machine Intell.* 18, 831–836.
- Yam, C.-Y., Nixon, M.S., Carter, J.N., 2002. Gait recognition by walking and running: A model-based approach. In: Proc. 5th Asian Conf. Comput. Vision, pp. 1–6.