Gait Analysis for Recognition and Classification

L. Lee W. E. L. Grimson
MIT Artificial Intelligence Lab
Cambridge, MA 02139

Abstract

This paper describes a representation of gait appearance for the purpose of person identification and classification. This gait representation is based on simple features such as moments extracted from orthogonal view video silhouettes of human walking motion. Despite its simplicity, the resulting feature vector contains enough information to perform well on human identification and gender classification tasks. We explore the recognition behaviors of two different methods to aggregate features over time under different recognition tasks. We demonstrate the accuracy of recognition using gait video sequences collected over different days and times and under varying lighting environments. In addition, we show results for gender classification based our gait appearance features using a support-vector machine.

1 Introduction

Gait is defined as "a manner of walking" in Webster's New Collegiate Dictionary. However, human gait is more than that: it is an idiosyncratic feature of a person that is determined by, among other things, an individual's weight, limb length, footwear, and posture combined with characteristic motion. Hence, gait can be used as a biometric measure to recognize known persons and classify unknown subjects. Moreover, we extend our definition of gait to include certain aspects of the appearance of the person, such as: the aspect ratio of the torso, the clothing, the amount of arm swing, the period and phase of a walking cycle, *etc*. To this end, we have designed a representation for the overall appearance of human gait that facilitates the recognition and classification of people by their gait.

Gait can be detected and measured at low resolution, and therefore it can be used in situations where face or iris information is not available in high enough resolution for recognition. Johansson [9] had shown in the 1970's that observers could recognize walking subjects familiar to them just by watching video sequences of lights affixed to joints of the walker. Hence, in theory, joint angles are sufficient for recognition of people by their gait. However, recovering joint angles from a video of walking person is an unsolved problem. In addition, using only joint angles ignores the appearance traits that are associated with individuals, such as heavy-set vs. slim, long hair vs. bald, or personal clothing and accessories. For these reasons, we have included appearance as part of our gait recognition features.

The challenges involved in gait recognition include imperfect foreground segmentation of the walking subject from the background scene, changes in clothing of the subject, variations in the camera viewing angle with respect to the walking subjects, and changes in gait as a result of mood or speed change, or as a result of carrying objects. The gait appearance features discussed below will tolerate some imperfection in segmentation and clothing changes, but not drastic style changes such as pants to skirts, nor is it impervious to changes in a person's gait. The view-dependent constraint of our gait appearance feature representation could be removed by synthesizing a walking sequence in a canonical view using the visual hull constructed from multiple cameras[17].

In the following sections, we summarize some of the previous work on the perception of gait and computational methods to measure gait, introduce our gait silhouette appearance representation along with two ways of aggregating the appearance representations over time, apply our representation to recognition and classification tasks, and suggest future work to improve the representation.

2 Previous Work

Several researchers have concluded that gait is indicative of a person's gender and identity. Johansson[9] used lights affixed to joints of a human body and showed observers motion sequences of the lights as a person walks. The observers were able to identify gender and, in cases where the observer was familiar with the walking subject, the identity of the walker. Cutting, *et al.* [5] studied human perception of gait using moving light displays (MLD) similar to that used by Johansson and showed human person identification results [4] and gender classification results [10]. They showed that human observers could identify gender with approximately 70% accuracy using only the visual cues from MLD.

Given the ability of humans to identify persons and classify gender by the joint angles of a walking subject, Goddard [7] developed a connectionist algorithm for gait recognition using joint locations obtained from MLD. However, computing joint angles from video sequence is still a difficult problem, though several attempts have been made on it [2, 6, 19]. The particular difficulties of joint angle computation from monocular video sequence are occlusion and joint angle singularities. Self occlusion of a limb from the camera view causes difficulties in tracking the hidden limb(s). Rehg and Morris [15] pointed out the singularity in motion along the optical axis of a camera.

There have been a number of appearance based algorithms for gait and activity recognition. Cutler and Davis [3] used self-correlation of moving foreground objects to distinguish walking humans from other moving objects such as cars. Polana and Nelson [14] detected periodicity in optical flow and used these to recognize activities such as frogs jumping and human walking. Bobick [1] used a time delayed motion template to classify activities. Little and Boyd [11] used moment features and periodicity of foreground silhouettes and optical flow to identify walkers. Nixon, *et al.* [12] used principal component analysis of images of a walking person to identify the walker by gait. Shutler, *et al.* [18] used higher order moments summed over successive images of a walking sequence as features in the task of identifying persons by their gait.

The work described in this paper is closely related to that of Little and Boyd [11]. However, instead of using moment descriptions and periodicity of the entire silhouette and optical flow of a walker, we divide the silhouettes into regions and compute statistics on these regions. We also further study the capacity of our features in tasks beyond person identification, such as gender classification.

3 Gait appearance representation

We consider the canonical view of a walking person to be one which is perpendicular to the direction of walk. We also assume that the silhouette of the walker is segmented from the background (details to follow). We would like our gait dynamics feature vector to have the following properties: the ability to describe appearance at a level finer than whole body description without having to segment individual limbs; robustness to noise in video foreground segmentation; and simplicity of representation. A number of features intuitively come into mind that may measure the static aspects of gait and individual traits. One such feature is the height of an individual, which requires calibrating the camera to recover distances. Other features include the amount of bounce of the whole body in a full stride, the side-to-side sway of the torso, the maximum distance between the front and the back legs at the peak of the swing phase of a stride, the amount of arm and leg swing, etc. We do not use all of these features for various reasons such as inaccessibility (the side-to-side sway of torso) or difficulties in obtaining features, such as detecting the peaks of swing phase when foreground segmentation is noisy and includes shadows.

Our gait appearance feature vector is comprised of parameters of moment features in image regions containing the walking person aggregated over time either by averaging or spectral analysis. For each silhouette of a gait video sequence, we find the centroid and proportionally divide the silhouette into 7 parts as in Figure 1. The vertical line through the centroid is used to divide the silhouette into front and back sections, except for the top portion. The parts above and below the centroid are each equally divided in the horizontal direction, resulting in 7 regions that roughly correspond to: r_1,head/shoulder region; r_2, front of torso; r_3, back of torso; r_4, front thigh; r_5, back thigh; r_6, front calf/foot; and r_7, back calf/foot. These regions are by no means meant to segment the body parts precisely. We are only interested in a method to consistently divide the silhouette of a walking person into regions that will facilitate the person recognition task.

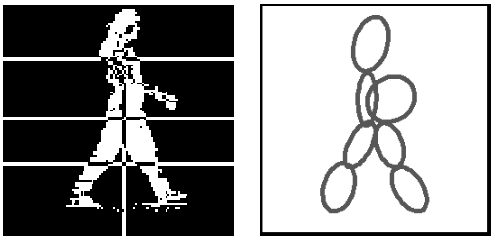

(a) (b)

Figure 1. The silhouette of a foreground walking person is divided into 7 regions, and ellipses are fitted to each region.

For each of the 7 regions from a silhouette, we fit an ellipse to the portion of foreground object visible in that region (Figure 1(b)). The features we extract from each of these regions are the centroid, aspect ratio (l) of major and minor axis of the ellipse, and the orientation (α) of major axis of the ellipse, i.e., the region feature vector $f(r_i)$ is,

$$f(r_i) = (\overline{x}_i, \overline{y}_i, l_i, \alpha_i), \text{ where } i = 1, \ldots, 7. \qquad (1)$$

These moment-based features are somewhat robust to noise in the silhouettes obtained from background subtraction, as long as the number of noise pixels is small and not systematically biased. The features extracted from each frame of a walking sequences consists of features from each of the 7 regions, *ie.* the frame feature vector F_j of the jth frame is,

$$F_j = (f(r_1), \ldots, f(r_7)). \qquad (2)$$

In addition to these 28 features, we use one additional feature, h, the height (relative to body length) of the centroid of the whole silhouette, to describe the proportions of the torso and legs. The intuition behind this measure is that an individual with longer torso will have a silhouette centroid that is positioned higher (relative to body length) on the silhouette than someone with a short torso.

Given the region features across a gait sequence, we need a concise representation across time. We compute two types of features across time: 1) the mean and standard deviation of region features across time, and, 2) magnitudes and phases of each region feature related to the dominant walking frequency (of 1 step, or half of a stride).

Specifically, the gait average appearance feature vector of a sequence s is,

$$s = (\text{mean}_j(h_j), \text{mean}_j(F_j), \text{std}_j(F_j)), \qquad (3)$$

where $j = 1, \ldots,$ last frame, and s is 57-dimensional. This feature set is very simple to compute and robust to noisy foreground silhouettes. Intuitively, the mean features describe the average–looking ellipses for each of the 7 regions of the body; taken together, the 7 ellipses describe the average shape of the body. The standard deviation features roughly describe the changes in the shape of each region caused by the motion of the body, where the amount of change is affected by factors such as how much one swings one's arms and legs.

Our gait spectral component feature vector of a sequence is ,

$$t = (\Omega_d, |X_i(\Omega_d)|, \text{phase}(X_i(\Omega_d))), \qquad (4)$$

where

$$X_i = \text{FourierTransform}(F_{j=1\ldots last}(f(r_i))), \qquad (5)$$

where Ω_d is the dominant walking frequency of a given sequence. Intuitively, the magnitude feature components measure the amount of change in each of the 7 regions due to motion of the walking body, and the phase components measure the time delay between the different regions of the silhouette. Because of noise in silhouettes, the time series of region features is also noisy. The power spectra of many of the region features do not show an obvious dominant peak that is the dominant walking frequency. Even when the peak frequencies are found, they often do not agree between different region features. We use a normalized averaging of power spectra of all region features resulting in a much more dominant peak frequency that is also consistent across all signals. The magnitude of each region feature at the dominant frequency Ω_d can be directly used, but the phase cannot be directly used because each gait sequence is not predetermined to start at a particular point of a walking cycle. We compute the phases of all region features relative to one particular region feature that is "most stable," by which we use the 2nd moment of power spectrum about the peak frequency to measure. In our case, the standard phase is that of the x position of the front calf/foot region. The gait spectral component feature vector has 57-1(the standard phase)=56 dimensions.

The length of our gait appearance feature vector leads one to question whether all the features have equal discrimination capacity for the tasks at hand, which are person recognition and gender classification. It may be that for different regions of the silhouette a different set of features may be better suited for recognition. A good feature should minimize within class variance and maximize inter-class variance. We use the simplifying assumption that all features are independent from each other and apply analysis of variance to each feature to measure its discrimination capacity in person classification and gender classification. The features can then be ranked by the p-value computed by ANOVA, and a good subset of features for a classification task is taken to be the set of features with p-value smaller than some threshold. Strictly speaking, this assumes independence of features, but in practice it gives good results. This analysis is done for both the gait averaging features and the spectral component features.

4 Experimental methods and results

We apply our gait dynamics features to person identification and gender classification. We gathered gait data in indoor environments with different backgrounds on four separate days spanning two months. Twenty-four subjects, 10 women and 14 men, were asked to walk at their normal speed and stride, back and forth, twice, in front of a video camera that was placed perpendicular to their walking path. The walking gait sequences going the opposite direction were flipped so that all sequences are of one walking direction. In all, 194 walking sequences were collected, between 4 to 22 sequences for each subject, averaging 8 sequences per subject. A minimum of 3 complete walking cycles were captured at 15 frames per second for each walking sequence. We then use an adaptive background subtraction algorithm [20] to segment out the walking person as a moving foreground object and scale-normalize the silhouettes. An example of the foreground walking person is shown in Figure 2. The foreground segmentation is not perfect—the shadow on the ground and in some cases portions of the background are included. However, our gait representation tolerates this amount of noise in foreground segmentation.

In addition to our gait database, we have extracted gait

(a) t=1 (b) t=2 (c) t=3 (d) t=4

Figure 2. A sample sequence of the silhouettes of a subject after background subtraction.

appearance features from the CMU gait database [8]. The setup for CMU gait data consists of subjects walking on a treadmill doing three different types of walk: holding a ball, fast walk, and slow walk. Each walking sequence is 10 seconds or more, recorded at 30 frames/second from several angles. We use foreground silhouettes provided by CMU and only the frontal-parallel view of the sequences. Each sequence is divided into 60-frame subsequences for comparisons between the subsequences. There are a total of 25 subjects, 23 men and 2 women, with 3 types of walk each (with the exception of one subject) and 5 subsequences for each walking type (after dividing into 60-frames subsequences), thus resulting in a total of 370 subsequences.

Using the gait averaged appearance and spectral components features that we described in the previous section, two feature vectors are computed for each walking sequence in our gait database and used in person identification and gender classification. We only compute the averaged appearance gait feature for the CMU database because the speed of the walker is determined by the treadmill. Only the recognition test is done on the CMU data because of gender imbalance in the database.

4.1 Person Identification

Assuming that all components in the feature vectors are equally important to gait recognition, using the Mahalanobis distance between feature vectors as a measure of the similarity between two sequences will remove the bias of parameters with large variance. We then use a nearest-neighbor approach to rank a set of gait sequences (and their corresponding features) by their distances to a query sequence. In addition, we applied ANOVA to rank the parameters of our gait features in the multi-person classification task and selected all features with $p < 10^{-9}$, which leaves us with the best 41 feature parameters in the case of gait averaged appearance feature set, and 32 parameters in the case of gait spectral components feature set.

We evaluated our gait recognition performance using the cumulative match score described by Phillips [13]. Under the closed-universe assumption, all query sequences contain a walker who has at least one other walking sequence in our collection of known walking sequences with known person ID's. We answer the question, "Is the correct person ID in the top n matches?" The cumulative match score has the rank of top matches on the x-axis and the percentage of correct identification on the y-axis.

Two gait recognition tests were performed using each of the two gait feature vectors. In one, the any-day test, a query walking sequence is compared against all sequences except itself, regardless of the day of recording. The second test, which we call the cross-day test, compares walking sequences taken on one day against sequences taken on other days only. This is a much more stringent test because sequences taken on different days have different lighting conditions, different backgrounds, and subjects may be wearing different style of clothing (*e.g.* pants vs. skirt).

4.1.1 Recognition using Averaged Appearance Feature

The cumulative match score using the best 41 features in the any-day test is 100% correct identification for the first match. In the case where all 57 features are used, it gives 97% correct identification for the first match and 100% by the 3rd match. The lower performance using the full set of 57 features suggest that the 16 features eliminated by a threshold on the p-value of ANOVA are unstable measures of gait appearance. Closer examination of the gait feature vector distances shows that the top matches are always those sequences of the same person taken on the same day, effectively making the any-day test into a same-day test, hence necessitating the second test—comparing sequences taken on different days.

The results of the second person identification test, the cross-day test, are shown in Figure 3. Each of these cumulative match score curves represents the comparison between walking sequences taken on one day (days A, B, C, or D) against sequences taken on all other days ((B,C,D), (A,C,D), (A,B,D), or (A,B,C), respectively). Table 1 shows the percentage of correct matchs using the 41 features selected using ANOVA vs. using all 57 features. The overall recognition performance is better when using the best 41 features than when using all 57 features.

Only subjects who have gait data collected from different days are used in this recognition test. The gait recognition performance is higher when comparing data collected on day A against all other sequences than matching sequences between those recorded on other days against the rest of the gait data. Closer examinations of the ranked matches show that the worse performances on the cross-day B, C, and D tests are the results of lack of similarity between the types of clothing that some of the subjects wore for the day B, C, or D gait sequence data and what they each wore on the other days. In other words, the query sequences show sub-

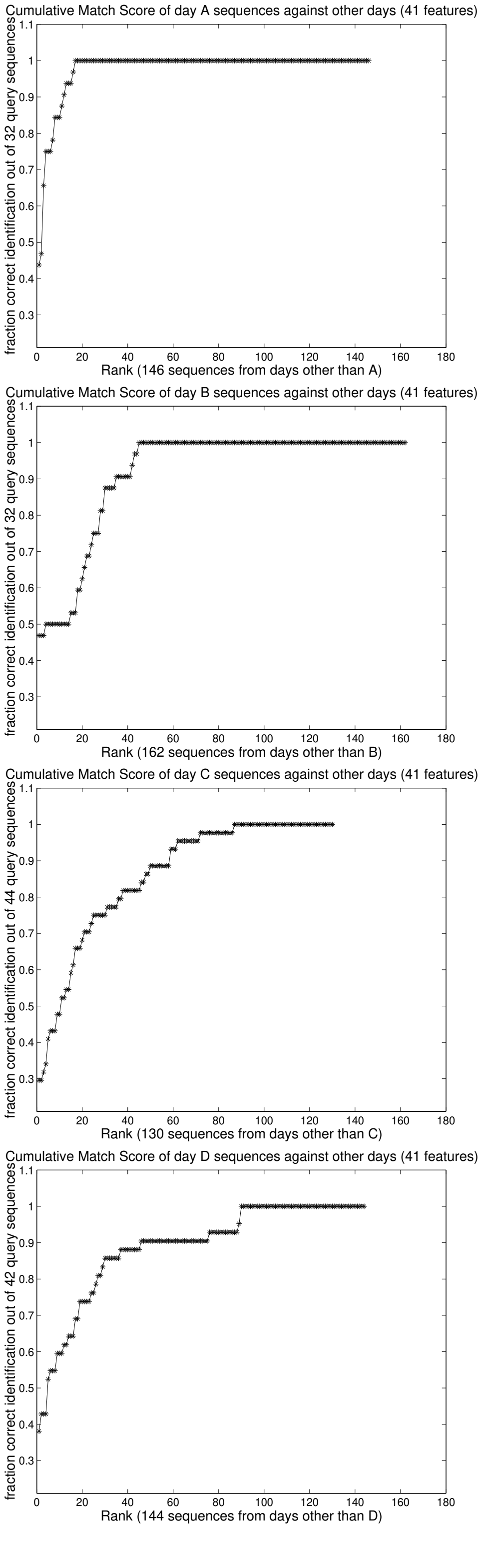

Figure 3. Cumulative match score for comparing sequences taken on one day against sequences taken on all other days using the 41 best averaged appearance features.

	41 features (% correct)				57 features (% correct)			
	1st	10%	20%	50%	1st	10%	20%	50%
A	44	94	100	100	47	91	100	100
B	47	53	88	100	25	56	69	91
C	30	55	75	95	25	55	80	91
D	38	64	83	90	26	69	81	90

Table 1. Percentage of correct identification using 41 and 57 averaged appearance features at the best match, the top 10%, 20%, and 50% recalls.

jects wearing clothing that is substantially different from that in the library sequences. For example, the 7 query sequences with the worst match scores from day B are all from one subject who wore baggy pants and whose only other sequences were collected on day D when he wore shorts. For the same reason, recognition results of matching day D gait sequences against all other sequences suffer because of lack of a similar appearance model in day B for the same subject. Day C contains sequences of one subject wearing a short dress whose only other sequences in the database showing her wearing pants.

We performed recognition test one on the 370 subsequences of CMU data set using the averaged appearance gait feature set. With the exception of one subsequence, all other subsequences had as their best match the correct person identity and walk type. The only mismatch was of a subsequence of a subject doing his slow walk in a style very much like his fast walk and hence the best matches were his fast walk sequences. We are unable to perform recognition test two on CMU data because sequences of the same subject were collected on the same day.

4.1.2 Recognition using Spectral Components Features

The cumulative match scores using the gait spectral component features are shown in Figure 4 for the query sequence compared with all other sequences regardless of day of recording and in Figure 5 for the query against sequences taken on different days.

As is evident in the recognition results, while the spectral component gait features do not perform as well as the average appearance gait features in the any-day test, its recognition performance is much better in the cross-day tests. The spectral components features perform only marginally better in the day A against day B, C, D test, but significantly better in the other three cross–day tests. Those three cross-day tests are exactly the ones that contain query sequences showing a subject wearing a style of clothing that is not present in the library sequences. This behavior re-

test	1st	5%	10%	20%	30%
any-day	82	97	98	100	100
cross-day A	63	88	97	100	100
cross-day B	50	84	97	100	100
cross-day C	45	70	84	91	95
cross-day D	31	81	93	98	100

Table 2. Percentage of correct identification using gait spectral component features at the best match, top 5%, 10%, 20%, and 30% recalls

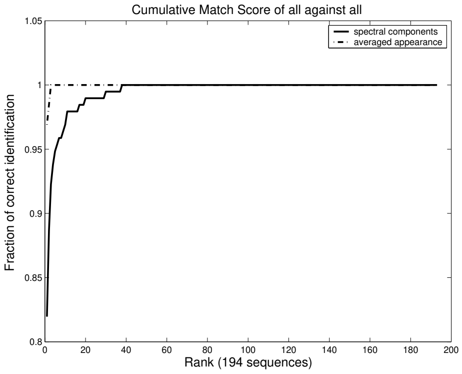

Figure 4. Cumulative match score for the all sequences against all sequences test using all 56 spectral component features and all 57 averaged appearance features.

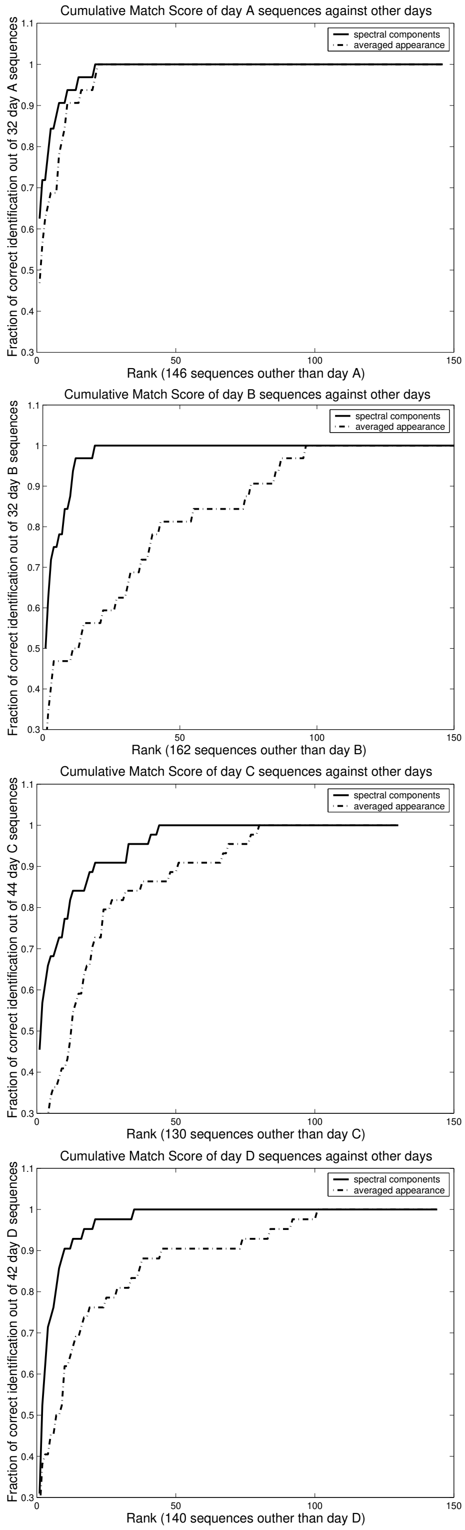

Figure 5. Cumulative match score for the sequences taken on one day against sequences taken on all other days tests using the spectral component features and the averaged appearance features.

sults from the characters of the features themselves. The average appearance gait features contain the mean shape of the 7 regions of the silhouettes. The shapes of these regions are very much affected by drastic changes in clothing style, such as pants vs. dress, shorts vs. pants. The gait spectral component features only contain the changes in shape and the time delay between the different regions, hence it is less affected by the change of clothing style.

We used the best 32 spectral components features in the above two recognition tests. Our results show that the reduced feature set has similar recognition rate with a slight degradation.

4.1.3 Discussion

The different behaviors of the average appearance features and the spectral component features point to choices that need to be made in a gait recognition system. If a gait recognition system is allowed to collect library sequences of each individual in all the distinct classes of clothing that he/she wears, then the average appearance is a very useful feature for recognition. If the system does not have the capabilities to collect extensive library, then the spectral components are much more reliable measures. It remains to be seen if there is a method to effectively combine these two types of features. We have done preliminary studies of combining the averaged gait appearance features and the gait spectral components features. The results need further investigation.

4.2 Gender classification

We apply the gait average appearance features to the task of gender classification. Specifically, we used the full 57 dimensional average appearance features as described in Section 3, as well as a smaller set of features selected using ANOVA. We ranked each of the 57 features based on the p-value of ANOVA in separating the genders and set a threshold of $p < 10^{-9}$, which resulted in the best 6 features (Table 3 for gender classification.

rank	region	feature type
1	hips/back thigh	std of orientation
2	chest/front arm	std of x of centroid
3	back foot/calf	std of orientation
4	back foot/calf	mean aspect ratio
5	hips/back thigh	mean orientation
6	front foot/calf	std of x of centroid

Table 3. Top 6 features for gender classification

We trained and tested support vector machines on our gait appearance features under two conditions. Under the random-sequence test, we randomly selected gait feature vectors of approximately half of the sequences, without regard to the identity of the walking subject, and tested on the gait features of the remaining sequences. Under the random-person test, we randomly selected approximately half of our walking subjects, trained the SVM on all sequences from these walkers, and tested on all sequences of the remaining walking subjects.

We used an implementation of support-vector machine by Rifkin [16] and experimented with the linear, gaussian, and the 2nd degree polynomial kernels. The SVM's are trained using the 57 and the 6 gender features and under the random-person vs. random-sequence conditions. The reuslts for these tests conditions are listed in Table 4. Overall, we found that the random-sequence test is easier because sequences from the same person, though not the same sequences, are in both the training and the testing set.

The random-person test condition is a more accurate representation of how a gender classifier would be used in a real application. The performance of the three kernels in the random-person case show that the linear kernel performed at least as well as the gaussian and the polynomial kernels. This leads us to believe that the boundary between the genders may be approximately linear.

	Random Sequence	
Kernel type	57 features	6 features
polynomial(d=2)	91%	94%
gaussian	93.5%	90%
linear	94%	88%
	Random Person	
Kernel type	57 features	6 features
polynomial(d=2)	79%	84.5%
gaussian	66%	83.5%
linear	80%	84.5%

Table 4. SVM gender classification results.

5 Discussion and Future Work

Our gait appearance features are view and appearance based. The view dependent nature of this representation can be resolved by synthesizing a view of the walking in the canonical view used in this paper. The dependence of our gait feature on the appearance of a subject is difficult to remove, unless the joint movements of a walker can be accurately detected or appearance insensitive features are used. However, as we argued earlier, the appearance of a walking subject actually contains information about the identity of the person. To allow accurate recognition of a person with drastic clothing style change then requires having multiple representations of the person. Recognition modalities such as face recognition can be very useful here because face is less susceptible to appearance change and can be easily combined with gait appearance recognition to improve recognition performance [17].

We have done experiments comparing gait appearance features from each individual against the rest of the population. Our preliminary results show that individuals range from having a number of very distinguishing features (from the population at large) to having just nominally distinguishing features. This leads us to believe that we can use the top-ranked features to build context-sensitive gait recognition engines to look for a particular individual or to decide that this person does not have a distinct gait appearance and some other measure should be used to recognize that individual.

6 Summary and Conclusions

We have presented a simple representation of human gait appearance based on moments computed from the silhouette of the walking person for the purpose of person identification and gender classification. Our representation is rich enough to show promising results in these tasks. We have described the characteristic behaviors of two types of gait features, one based on average appearance of gait, and one based on spectral components. We also show that the distinction between genders is consistent with a linear foundary in our averaged appearance representation. This view and appearance dependent model of gait can be further extended to accommodate a multiple appearance model of a person and in conjunction with other recognition modalities.

Acknowledgments

We would like to thank the anonymous reviewers for their suggestions. This research is supported by DARPA contract N00014-00-1-0907.

References

[1] A. Bobick and J. Davis. The recognition of human movement using temporal templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(3), March 2001.

[2] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *Proc. of IEEE CVPR*, 1998.

[3] R. Cutler and L. Davis. Robust real-time periodic motion detection, analysis, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):781 –796, 2000.

[4] J. Cutting and L. Kozlowski. Recognizing friends by their walk: gait perception without familiarity cues. *Bull. Psychonometric Soc.*, (9):353–356, 1977.

[5] J. Cutting, D. Proffitt, and L. Kozlowski. A biomechanical invariant for gait perception. *Journal of Experimental Psychology: Human Perception and Performance*, 4(3):357–372, 1978.

[6] D. DiFranco, T. Cham, and J. Rehg. Recovery of 3D articulated motion from 2d correspondences. Technical report, Compaq Cambridge Research Lab, 1999.

[7] N. Goddard. *The Perception of Articulated Motion: Recognizing Moving Light Displays*. PhD thesis, University of Rochester, 1992.

[8] R. Gross and J. Shi. The CMU motion of body (mobo) database. Technical Report CMU-RI-TR-01-18, Carnegie Mellon University, 2001.

[9] G. Johansson. Visual motion perception. *Scientific American*, (232):76–88, 1975.

[10] L. Kozlowski and J. Cutting. Recognizing the sex of a walker from a dynamic point-light display. *Perception and Psychophysics*, 21(6):575–580, 1977.

[11] J. Little and J. Boyd. Recognizing people by their gait: The shape of motion. *Videre*, 1(2), 1998.

[12] M. Nixon, J. Carter, J. Nash, P. Huang, D. Cunado, and S. Stevenage. Automatic gait recognition. In *Motion Analysis and Tracking (Ref. No. 1999/103), IEE Colloquium on*, pages 3/1–3/6, 1999.

[13] P. Phillips, H. Moon, P. Rauss, and S. Rizvi. The feret evaluation methodology for face-recognition algorithms. In *CVPR*, pages 137–143, 1997.

[14] R. Polana and R. Nelson. Detecting activities. In *CVPR*, 1993.

[15] J. Rehg and D. Morris. Singularities in articulated objtect tracking with 2-d and 3-d models. Technical Report CRL 98/8, Compaq Cambridge Research Lab, 1997.

[16] R. Rifkin. *SvmFu*. http://five-percent-nation.mit.edu/SvmFu.

[17] G. Shakhnarovich, L. Lee, and T. Darrell. Integrated face and gait recognition from multiple views. In *CVPR*, 2001.

[18] J. Shutler, M. Nixon, and C. Harris. Statistical gait recognition via velocity moments. In *Visual Biometrics, IEEE*, pages 10/1–10/5, 2000.

[19] H. Sidenbladh, M. Black, and D. Fleet. Stochastic tracking of 3d human figures using 2D image motion. In *European Conference on Computer Vision (ECCV)*, volume 2, pages 702–718, 2000.

[20] C. Stauffer and W. Grimson. Adaptive background mixture models for real-time tracking. In *CVPR*, 1999.