Gait Sequence Analysis using Frieze Patterns

Yanxi Liu, Robert T. Collins and Yanghai Tsin

CMU-RI-TR-01-38

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

©2001 Carnegie Mellon University

This research is supported in part by an ONR research grant N00014-00-1-0915 (HumanID), and in part by an NSF research grant IIS-0099597.

ABSTRACT

We analyze walking people using a gait sequence representation that bypasses the need for frame-to-frame tracking of body parts. The gait representation maps a video sequence of silhouettes into a pair of two-dimensional spatio-temporal patterns that are periodic along the time axis. Mathematically, such patterns are called “frieze” patterns and associated symmetry groups “frieze groups”. With the help of a walking humanoid avatar, we explore variation in gait frieze patterns with respect to viewing angle, and find that the frieze groups of the gait patterns and their canonical tiles enable us to estimate viewing direction. In addition, analysis of periodic patterns allows us to determine the dynamic time warping and affine scaling that aligns two gait sequences from similar viewpoints. We show how gait alignment can be used to perform human identification and model-based body part segmentation.

1 Motivation

Automated visual measurement of human body size and pose is difficult due to non-rigid articulation and occlusion of body parts from many viewpoints. The problem is simplified during gait analysis, since we observe people performing the same activity. Although individual gaits vary due to factors such as physical build, body weight, shoe heel height, clothing and the emotional state of the walker, at a coarse level the basic pattern of bipedal motion is the same across healthy adults, and each person's body passes through the same sequence of canonical poses while walking. We have experimented with a simple, viewpoint-specific spatio-temporal representation of gait. The representation collapses a temporal sequence of body silhouette images into a periodic two-dimensional pattern. This paper explores the use of these frieze patterns for viewing angle determination, human identification, and non-rigid gait sequence alignment.

2 Related Work

Many approaches to analyzing gait sequences are based on tracking the body as a kinematic linkage. Model-based kinematic tracking of a walking person was pioneered by Hogg [7], and other influential approaches in this area are [2, 3]. These approaches are often brittle, since the human body has many degrees of freedom that cannot be observed well in a 2D image sequence. Our work is more closely related to approaches based on pattern analysis of spatio-temporal representations. Niyogi and Adelson delineate a person's limbs by fitting deformable contours to patterns that emerge from taking spatio-temporal slices of the XYT volume formed from an image sequence [14]. Little and Boyd analyze temporal signals computed from optic flow to determine human identity from gait [10]. The key point is that analyzing features over a whole temporal sequence is a powerful method for overcoming noise in individual frames.

Liu and Picard [11] proposed to detect periodic motions by studying treating temporal changes of individual pixels as 1D signals whose frequencies can be extracted. Seitz and Dyer [15] replace the concept of period by the instantaneous period, the duration from the current time instant at which the same pattern reappears. Their representation is effective in studying varying speed cyclic motions and detecting irregularities. Cutler and Davis [4] also measure self-similarity over time to form an evolving 2D pattern. Time-frequency analysis of this pattern summarizes interesting properties of the motion, such as object class and number of objects.

3 A Spatio-Temporal Gait Representation

Consider a sequence of binary silhouette images $b(t) \equiv b(x, y, t)$, indexed spatially by pixel location (x, y) and temporally by time t. Form a new 2D image $F_C(x, t) = \sum_y b(x, y, t)$, where each column (indexed by time t) is the vertical projection (column sum) of silhouette image $b(t)$, as shown in Figure 1. Each value $F_C(x, t)$ is then a count of the number of silhouette pixels that are "on" in column x of silhouette image $b(t)$. The result is a 2D pattern, formed by stacking column projections together to form a spatio-temporal pattern. A second pattern $F_R(y, t) = \sum_x b(x, y, t)$ can be constructed by stacking row projections. Since a human gait is periodic with respect to time, F_C and F_R are also periodic along the time dimension. A two-dimensional pattern that repeats along one dimension is called a *frieze* pattern in the mathematics and geometry literature, a *tile* of a frieze pattern is the smallest rectangle region whose translated copies can cover the whole pattern without overlapping or gaps. Group theory provides a powerful tool for analyzing such patterns (Section 4.1).

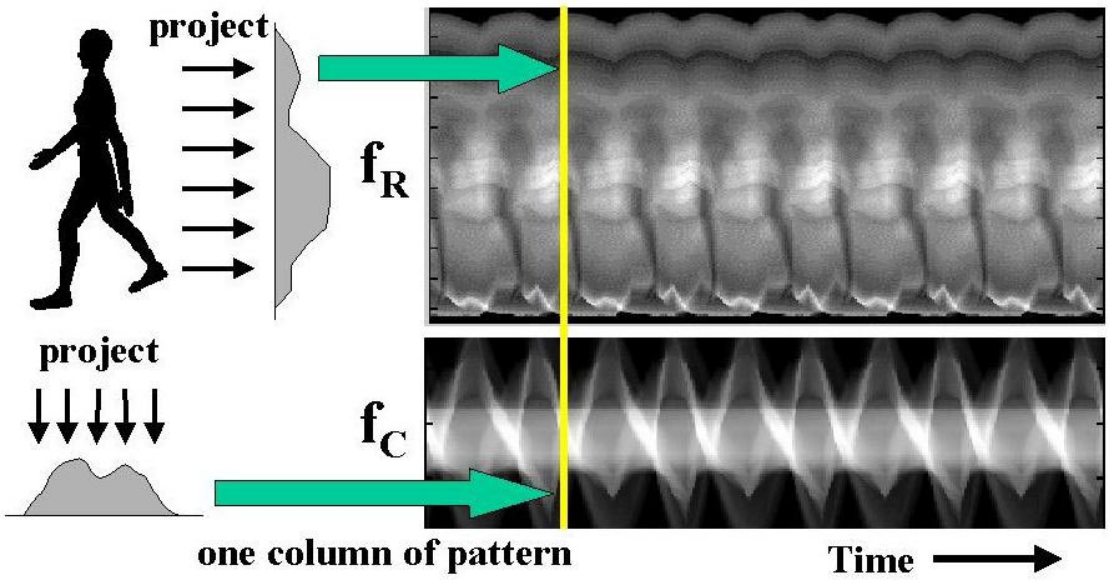

Figure 1: Spatio-temporal gait representations are generated by projecting the body silhouette along its columns and rows, then stacking these 1D projections over time to form 2D patterns that are periodic along the time dimension. A 2D pattern that repeats along one dimension is called a "frieze" pattern.

Figure 2 shows the column projection frieze pattern F_C extracted from a roughly 30 second long sequence of a person walking along a test course. Note the changes in appearance of the frieze pattern as the walking direction changes. In our experiments, body silhouette extraction

is achieved by simple background subtraction and thresholding, followed by a 3x3 median filter operator to suppress spurious pixel values. Silhouettes across a gait sequence are automatically aligned by scaling and cropping based on bounding box measurements so that each silhouette is 80 pixels tall, centered within a template 80 pixels wide by 128 pixels high. Background subtraction

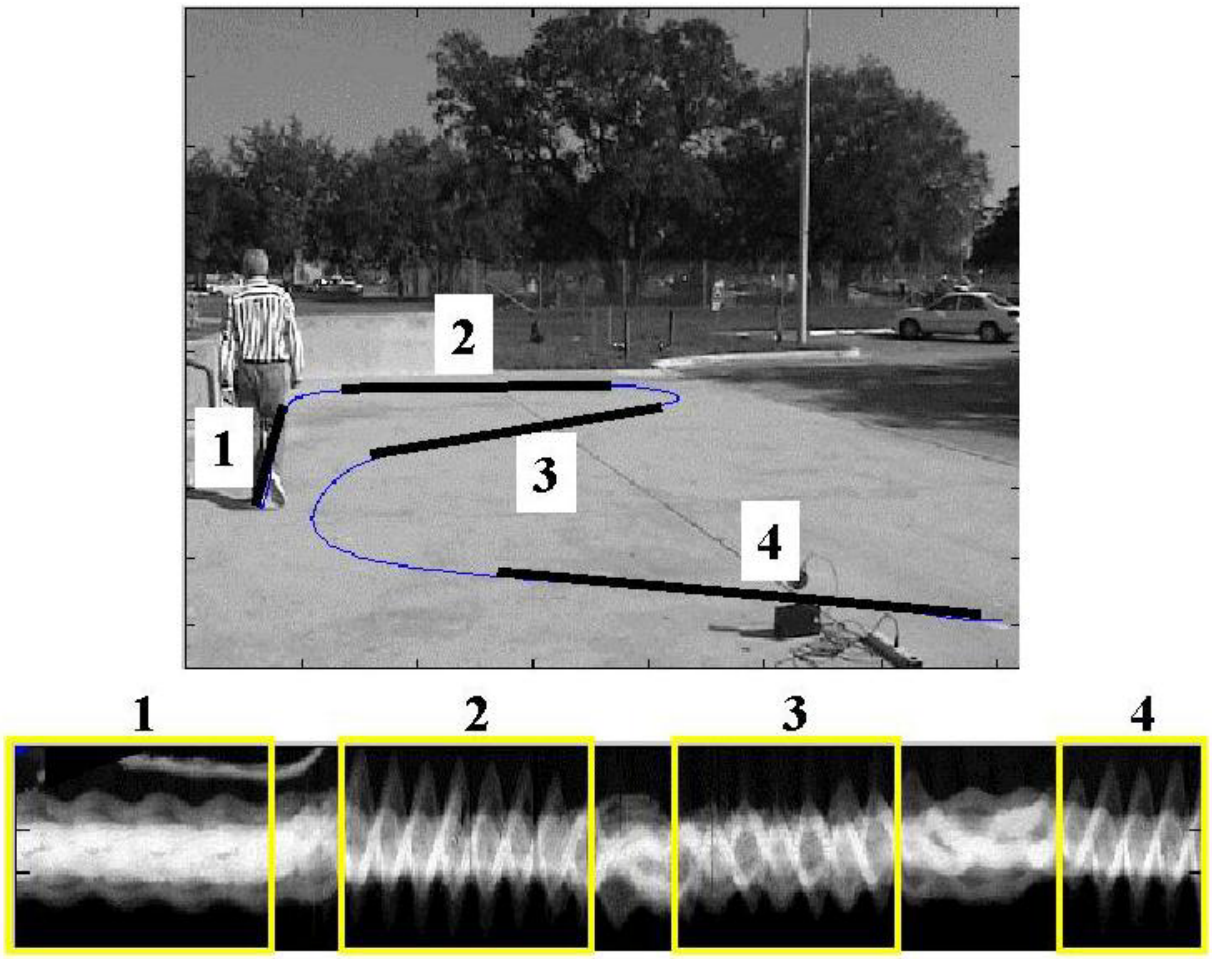

Figure 2: Frieze pattern extracted from a 30 second long walking sequence. Note the changes in appearance of the frieze pattern as the walking direction changes.

in real environments typically yields noisy silhouettes with holes, fragmented boundaries, and extra parts due to background clutter and shadows. It is difficult to automatically identify individual limb positions from such data. By distilling a sequence of silhouettes into a periodic pattern that can be smoothed and analyzed using robust signal analysis techniques, we no longer need to deal with noisy silhouette data.

4 Model-Based Gait Analysis

With the aid of a 3D walking humanoid model, we have studied how the spatio-temporal frieze patterns described above vary with respect to camera viewpoint. Our model of human body shape and walking motion is encapsulated in a VRML/H-Anim 1.1 compliant avatar called "Nancy". [1] Nancy's 3D polyhedral body parts were generated by a graphics designer, and the gait motion, specified by temporal sequences of interpolated rotations at each joint, is based on motion studies from "The Human Figure in Motion" by Eadweard Muybridge. We have ported Nancy into an open-GL program that generates 2D perspective views of the avatar given a camera position and time step within the gait cycle. Gaits are sampled at a rate of 60 frames per stride (one stride is two steps, i.e. one complete cycle). Figure 4 illustrates variation of the column projection frieze

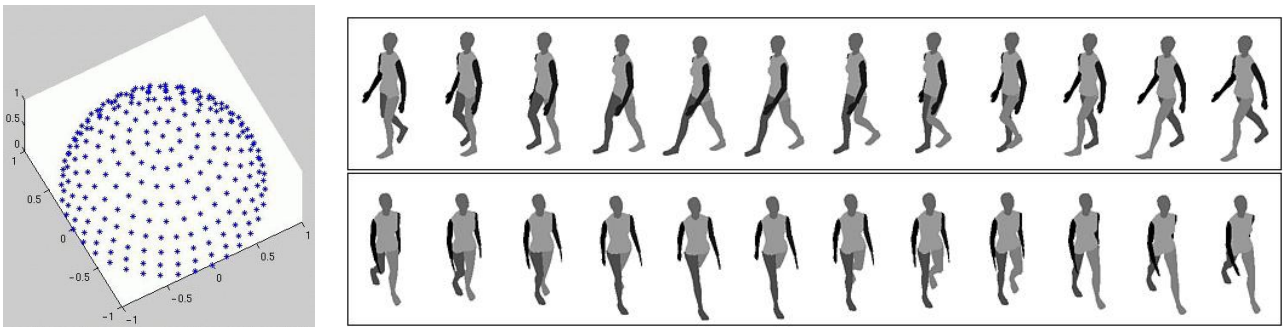

(a) (b)

Figure 3: (a) A database of gait sequences is generated from 241 sample viewpoints. The subject is a walking humanoid avatar. (b) Subsampled sequences from two viewpoints. Each body part of the avatar is color-coded with a different shade of grey.

patterns defined in Section 3 when Nancy's gait is seen from different viewing directions. The diversity inspires us to seek an encoding for these different types of frieze patterns in order to determine viewpoint from frieze group type. One natural candidate for categorizing frieze patterns is by their symmetry groups.

4.1 Frieze Symmetry Groups Classification

Any frieze pattern P_i in Euclidean space R^2 is associated with a unique symmetry group F_i, where $i = 1..7, \forall g \in F_i, g(P_i) = P_i$. These seven symmetry groups are called *frieze groups*, and their

[1] ©1997 Cindy Ballreich, 3Name3D / Yglesias, Wallock, Divekar, Inc. Available from `http://www.ballreich.net/vrml/h-anim/nancy_h-anim.wrl`

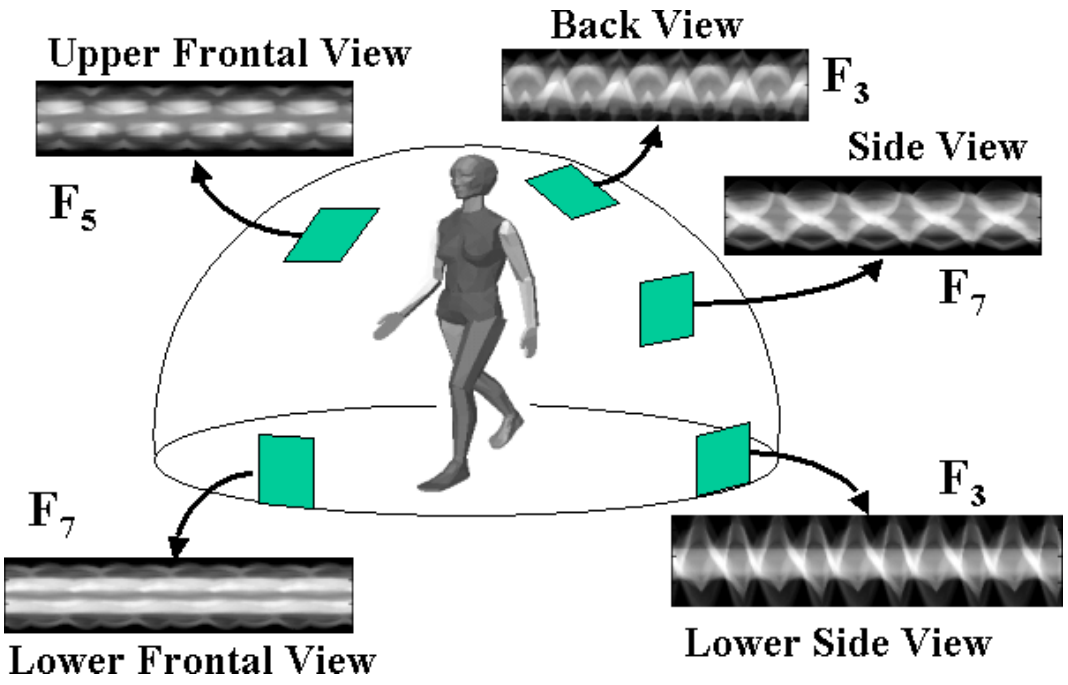

Figure 4: Different gait patterns of a computer avatar "Nancy", obtained from a real human being, viewed in different directions.

properties are summarized in Figure 5 and Table 1. Five different types of symmetries can exist for frieze patterns: (1) translation, (2) 2-fold rotation, (3) vertical reflection, (4) horizontal reflection and (5) glide-reflection. A frieze pattern can be classified into one of the 7 frieze groups based on what combination of these 5 primitive symmetries are present in the pattern. Computer algorithms for automatic frieze and wallpaper pattern symmetry group classification are proposed in [12, 13].

We are interested in classifying imperfect and noise-contaminated frieze patterns generated from avatar and human gaits. There are two important and intertwined computational issues for frieze symmetry group classification: 1) given an imperfect frieze pattern, how to decide whether or not it has certain types of symmetries; and 2) given the symmetry measures for a pattern, how to give each of the seven frieze groups an equal chance to be chosen as the symmetry group of the pattern, since these groups are not disjoint. The first issue is addressed by defining a distance measure between a pattern and a family of patterns have the same frieze group. The second issue is addressed by using geometric AIC for symmetry group model selection.

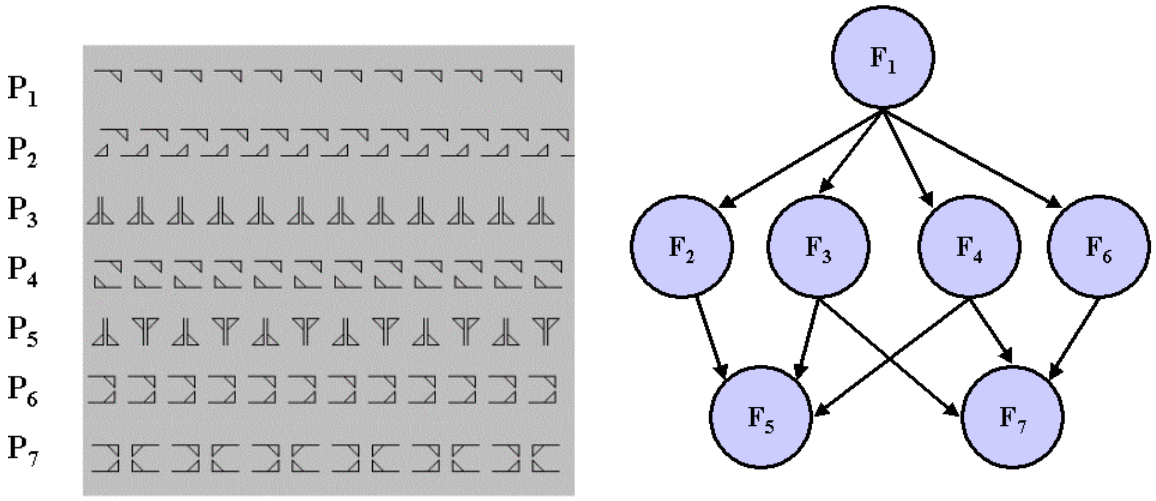

(A) (B)

Figure 5: (A) The seven frieze patterns ($P_1 ... P_7$) in Euclidean space R^2. (B) The subgroup relationship among the seven frieze symmetry groups ($F_1 ... F_7$ in Table 1). $F_i \rightarrow F_j$ means F_i is a subgroup of F_j.

4.1.1 Distance to the Nearest Frieze Patterns

We define the symmetry distance (SD) of an approximately periodic pattern P to frieze patterns $\{P_n\}$ with frieze group F_n as

$$SD_n(P) = \min_{Q \in \{P_n\}} \{ \sum_{i=1}^{tN} \left(\frac{p_i - q_i}{s_i} \right)^2 \} \tag{1}$$

where N is the number of pixels in a tile (smallest 2D repeating region), t is the number of tiles being studied, p_i and q_i are intensity values of corresponding pixels of pattern P and $Q \in \{P_n\}$ respectively, and s_i is the standard deviation of the frieze pattern at pixel i. For independent Gaussian noise, the distance SD_n has a χ^2 distribution with tN degrees of freedom.

The symmetry distance measure is defined with respect to a frieze pattern $Q \in \{P_n\}$ that has the minimal distance to P. We can show that this pattern Q can be constructed as follows: (1) For $t > 1$ and $n = 1$, Q is the pixel-wise average of all the tiles in P. (2) For $t = 1$ and $n > 1$, Q = $\frac{(\mathcal{O}_\setminus(\mathcal{P})+\mathcal{P})}{2}$, where $\mathcal{O}_\setminus(\mathcal{P})$ is the pattern obtained by applying the set of symmetry operations in F_n to P. (3) For $t > 1$ and $n > 1$, Q is the pixel-wise average of each Q obtained above. Our definition of frieze pattern symmetry distance in pixel intensity space is analogous to that of Zabrodsky et.al. [16, 9] for polygon distance in vertex location space.

Table 1: Symmetries of Frieze Patterns (N is number of pixels in one tile)

Symmetry Group	translation	180^0 rotation	Horizontal reflection	Vertical reflection	Nontrivial glide-reflection	Degrees of Freedom
F1	yes	no	no	no	no	N
F2	yes	no	no	no	yes	N/2
F3	yes	no	no	yes	no	N/2
F4	yes	yes	no	no	no	N/2
F5	yes	yes	no	yes	yes	N/4
F6	yes	no	yes	no	no	N/2
F7	yes	yes	yes	yes	no	N/4

4.1.2 Geometric AIC for Frieze Group classification

The frieze symmetry groups form a hierarchical structure (Figure 5)(B) where frieze group F_1 is a subgroup of all the other groups and so on. For example, a frieze pattern P_3 (with vertical reflection symmetry) is a more general pattern type than P_5 or P_7, since any P_5 or P_7 frieze with more complicated symmetries also has vertical reflection symmetry. But this implies that the distance of a pattern P to P_3 is always no greater than the distance to P_5, since the set of P_5 patterns is a subset of the P_3 patterns. If no care is taken, a symmetry group classification algorithm based on raw symmetry distance scores will always favor P_3 over P_5. To address this problem, we adopt the concept of Geometric-AIC (G-AIC) proposed by Kanatani [8, 9]. Given two possible frieze patterns whose symmetry groups have a subgroup relationship, G-AIC states that we should prefer F_m over F_n if

$$\frac{SD_m}{SD_n} < 1 + \frac{2(d_n - d_m)}{r(tN) - d_n} \tag{2}$$

where d_m and d_n are the degrees of freedom for frieze patterns of F_m and F_n respectively, and r is the co-dimension. Since the data space (the intensity space) is dimension one, and our model space (point in multidimensional intensity space) dimension is 0, the co-dimension $r = 1 - 0 = 1$.

The degrees of freedom (DOF) of a frieze pattern depends on how the intensity of each pixel on the pattern is constrained. For frieze patterns with translation symmetry only, the only constraint for each of the tN pixels is to have the same intensity value as the pixel t units to the left. Thus its DOF is N. On the other hand, pixels on a P_3 pattern have to satisfy a vertical reflection symmetry constraint, and thus half of the pixel intensities need to be the same as the other half. So the DOF of a P_3 pattern is $N/2$. The last column of Table 1 and Figure 6 explain the DOFs of the seven

frieze groups. In summary, we would prefer classify a pattern P as frieze group F_m rather than

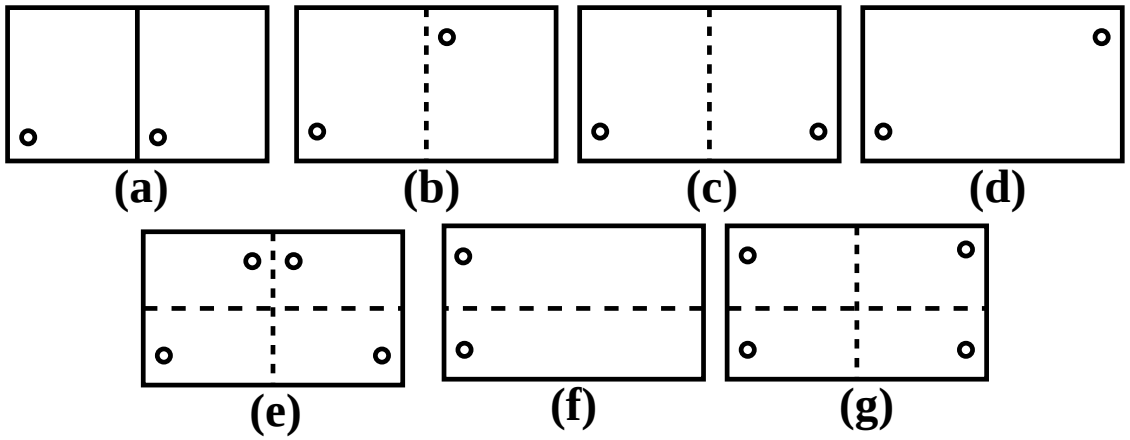

Figure 6: Determining the degrees of freedom of frieze patterns by how many constraints a pixel intensity has to satisfy. The figure shows the corresponding pixels that must have the same intensity values in (a) two tiles of a P_1 pattern; (b)-(g) a tile from frieze pattern $P_2 ... P_7$ respectively.

F_n if

$$\frac{SD_m(P)}{SD_n(P)} < \frac{t}{t-1}, \text{ for } m = 2, 3, 4, 6 \text{ and } n = 1 \quad (3)$$

$$\frac{SD_m(P)}{SD_n(P)} < \frac{2t}{2t-1}, \text{ for } m = 5, 7 \text{ and } n = 2, 3, 4, 6 \quad (4)$$

$$\frac{SD_m(P)}{SD_n(P)} < \frac{2t+1}{2t-2}, \text{ for } m = 5, 7 \text{and } n = 1 \quad (5)$$

4.2 View Direction Estimation

We have generated a database of 241 walk sequences, indexed by viewing direction azimuth and elevation, by sampling the view sphere at roughly every 10 degrees. Figure 7 shows the associated frieze groups with the 241 frieze patters from Nancy's gaits viewed from the 241 directions.

The underlying assumption in current approach is that the distribution of symmetry groups of the gait patterns from different views of the computer model can provide guidance for determining the viewing angle of an observed human subject's gait. Table 2 lists the data of two human subjects observed from six cameras and the frieze symmetry group associated with each frieze pattern. Figure 8 shows a comparison among corresponding avatar and subject frieze patterns, viewed from six different viewing angles. One can observe that frieze patterns from the same view point share the same frieze symmetry group, and the tiles have a similar appearance.

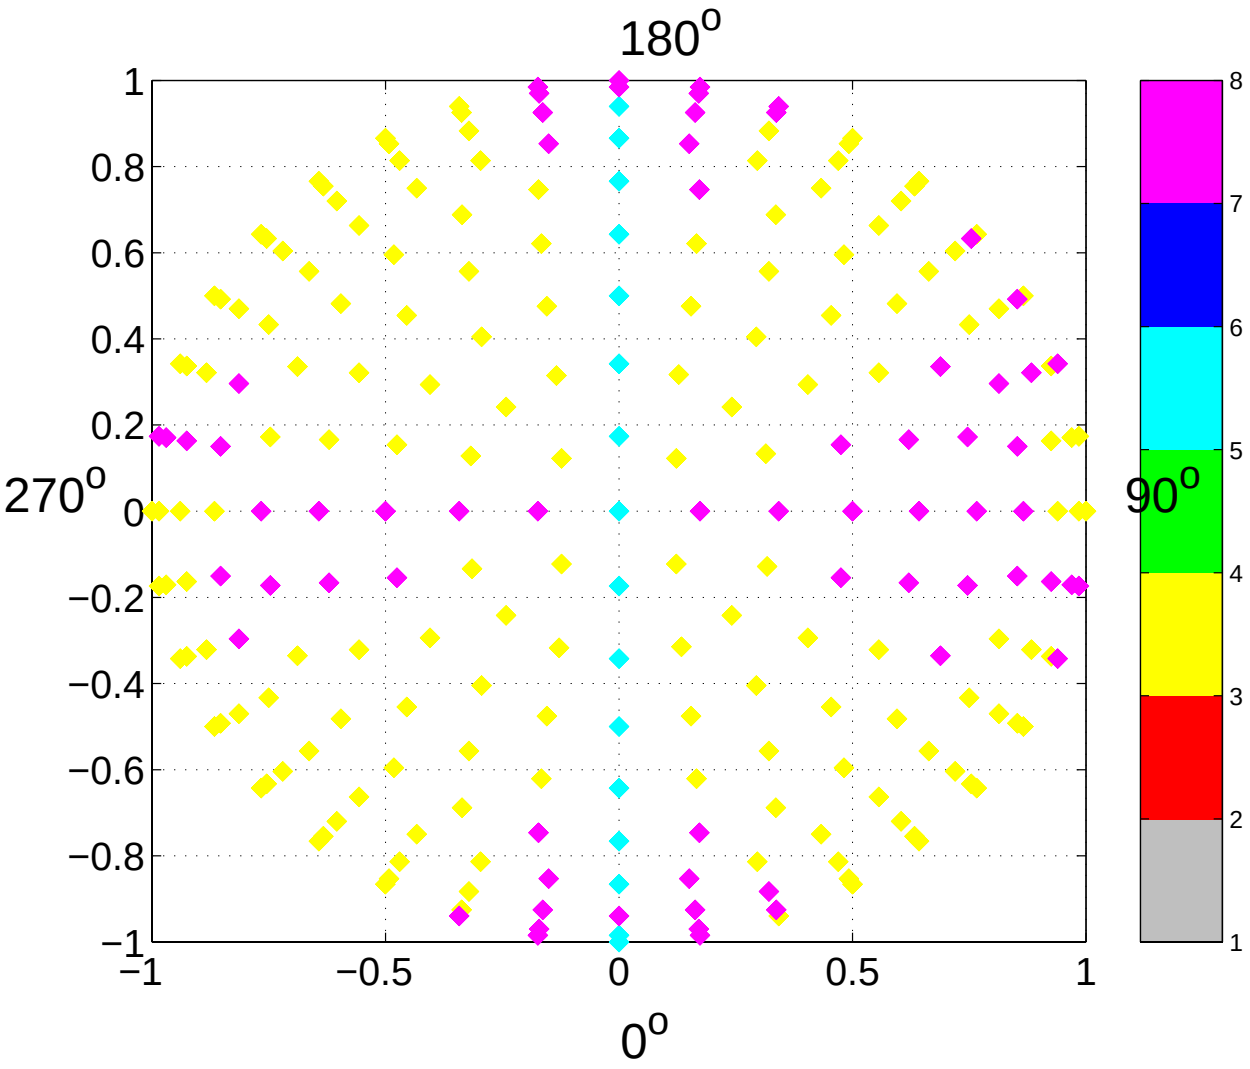

Figure 7: Symmetry group labelings of the frieze patterns of "Nancy" when viewed from all different orientations. Each sample point on the hemisphere is projected to the plane of elevation 0 degrees. From this pattern, one can see slight bilateral left-right asymmetry of the symmetry groups distribution. This reflects the possible gait asymmetry of the person this computer model is coming from. In general, the distribution of the gait frieze patterns symmetry groups is bilaterally symmetrical viewed from side to side. 0^o is the frontal view. This figure is best viewed in color.

Table 2: Frieze Groups of Human Subjects Gaits Viewed from Six Different Cameras

Camera ID	Sym. Group	View Dir.	elevation angle	azimuth angle	Vertical FOV	Horizontal FOV	subj # estimated	subj #2 estimated
3	F7	L. side	15.4	83.2	42.5	33.1	e:50 a:75	e:10 a:80
5	F3	L. front	12.0	37.4	37.6	29.0	e:30 a:160	e:80 a:45
7	F5	Frontal	25.0	359.8	43.7	33.9	e:20 a:0	e:20 a:0
13	F3	R. back	11.4	234.9	39.0	30.1	e:20: a:200	e:20 a:240
16	F3	R. front	11.9	314.5	32.1	24.7	e:40 a:334	e:10 a:20
17	F5	Back	26.5	181.4	42.3	32.8	e:20 a:180	e:20 a:180

Given an observed human subject gait pattern $P = F_C$ (Section 3), we use a moment-based method (Section 5.1) to align the model friezes P_i from each of the 241 viewing directions to the subject frieze. Applying PCA to a typical tile from P and taking the non-dominant PCA components that are most sensitive to discriminate pattern variations, the closest K nearest neighbors are found in this subspace. We used a dual elimination method to decide which angle values from these K neighbors we can count on. The first condition is that P and P_i have the same symmetry group. The second condition is that corresponding pixels of tiles from P and P_i must have similar intensities. Results for classifying viewing direction for two human subjects is listed in Table 2.

In this framework we have assumed affine camera models, thus only one viewing direction (specified by azimuth and elevation) is estimated for a ray directed towards the center of the person. However, the data used in the experiment comes from a perspective camera, and due to the proximity of the subject, the difference in viewing ray elevation between their head and feet is roughly 28 degrees. Much more accurate estimations of viewing angles are achieved for frieze patterns with non-F_3 groups (Table 2). This can be explained by the multiple possible angle ranges for P_3 patterns (Figure 7). We have dealt with this using majority votes and robust median estimators.

5 Spatio-Temporal Gait Alignment

Consider two gait sequences, represented by two pairs of frieze patterns $F_C(x,t)$, $F_R(y,t)$ and $F'_C(x',t')$, $F'_R(y',t')$. We seek to align the patterns temporally and spatially, as a precursor for further correspondence-based analysis. Temporal alignment of gait sequences amounts to align-

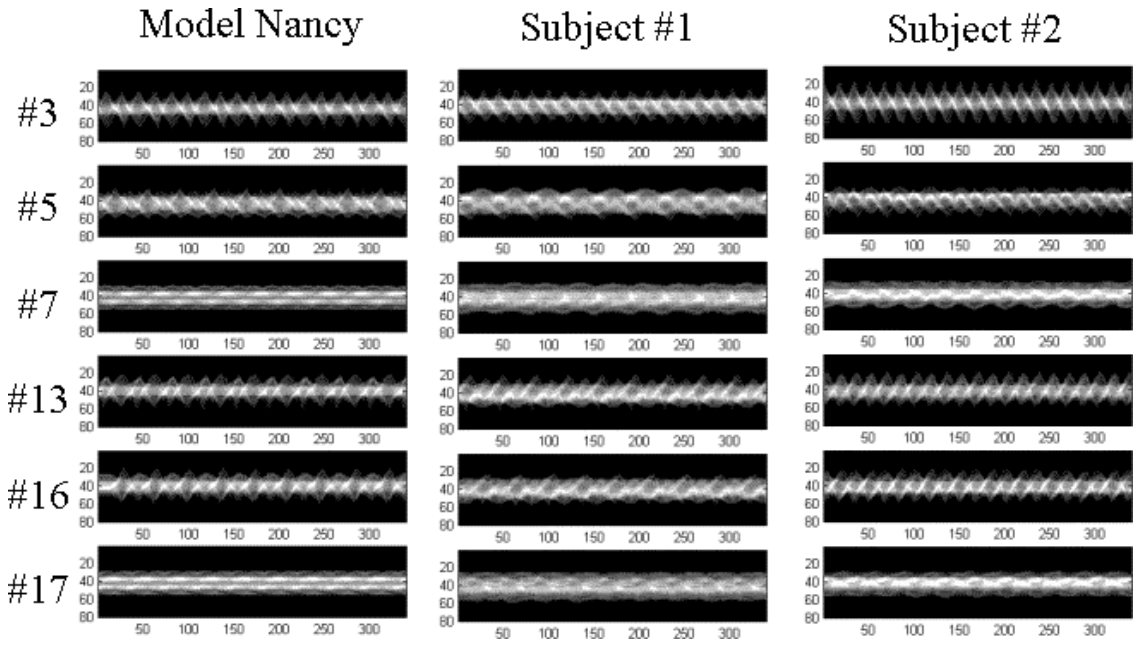

Figure 8: The view of the frieze patterns from six different angels. Left: Avatar Nancy. Middle: subject # 1. Right: subject # 2.

ing frieze patterns horizontally, thereby determining a mapping between time variables t and t'. Temporal alignment of frieze patterns is greatly simplified by the periodic nature of the patterns, allowing us to use simple periodic signal analysis in place of expensive dynamic time warping procedures. Spatial alignment means finding a mapping between pixel locations (x, y) in sequence 1 and (x', y') in sequence 2. We restrict this to a four parameter affine mapping, and show that it can be found by aligning the corresponding row and column friezes along their vertical dimensions.

5.1 Moment-Based Gait Alignment

It is well known that the first and second moments of two binary silhouettes can be used to determine an affine transformation that coarsely aligns them, and that some of the moments of a silhouette image can be computed from its row and column projections [1]. This forms the basis of our gait alignment method.

First, we generalize the concept of moments of a binary image to cover a time series of moments computed from a sequence of binary images. Define a *moment sequence* as

$$m_{ij}(t) = \sum_x \sum_y x^i y^j b(x, y, t)$$

which is a sequence of single-frame binary silhouette moments, indexed by time. Note that $m_{00}(t)$

is just the area of the binary silhouette over time, while $\bar{x}(t) \equiv m_{10}(t)/m_{00}(t)$ and $\bar{y}(t) \equiv m_{01}(t)/m_{00}(t)$ are the coordinates of the silhouette centroid over time. Similarly, define a *central moment sequence* as

$$\mu_{ij}(t) = \sum_x \sum_y (x - \bar{x}(t))^i (y - \bar{y}(t))^j b(x, y, t)$$

which is a sequence of moments measured after translating each silhouette so that its centroid is at the origin. The second central moments measure the spread of silhouette pixels about the centroid, and can be used to derive the principal axis of the silhouette shape.

Since we have summarized each sequence of silhouettes as frieze patterns, we are concerned only with moments that can be computed from row and column projections. For example, consider silhouette area

$$m_{00}(t) = \sum_x \sum_y b(x, y, t) = \sum_x \left(\sum_y b(x, y, t) \right) = \sum_x f_C(x, t)$$

which can thus be computed from the frieze pattern as well as the original silhouette sequence. Any moment sequence $m_{ij}(t)$ or central moment sequence $\mu_{ij}(t)$ with either i or j (or both) equal to zero can be computed from frieze patterns $f_C(t)$ and $f_R(t)$. In the present case, we will use $m_{00}(t)$, $m_{10}(t)$, $m_{10}(t)$, μ_{20}(t), and $\mu_{02}(t)$. Note that the second central moment $\mu_{11}(t)$ can not be determined from the two frieze patterns, and we will therefore not be able to adjust skew or principle axis rotation when aligning silhouette shapes using friezes alone. Difference in skew angle is not an issue when matching silhouettes seen from nearby viewpoints. Requiring the principle axes to be the same means that we are assuming the roll angles of the camera are also similar. Should an application require these extra degrees of freedom, they could be computed by included a third frieze formed by silhouette projection along a 45 degree diagonal axis, from which the cross-moments $m_{11}(t)$ and $\mu_{11}(t)$ can be measured [1].

We now present an algorithm for moment-based gait alignment. To a first approximation, the temporal alignment between the two periodic gait sequences can be represented as $t' = \rho\, t + \phi$, where ρ corrects for the relative stride frequency and ϕ corrects for the relative phase difference (position within a stride). The average stride frequency of each gait sequence is found by taking signal $m_{00}(t)$, "whitening" it by subtracting its mean and dividing by its standard deviation, then autocorrelating to find peaks occurring at a fundamental frequency. From some viewpoints this is the stride frequency, and from others it is half the stride frequency (e.g. a bipedal gait viewed from the side looks self-similar halfway through a full stride). Whether the autocorrelation of $m_{00}(t)$ yields peaks at half the stride frequency is viewpoint dependent, and can be calibrated using the walking avatar model. Let f and f' denote the average frequencies of the two gait sequences,

computed from m_{00} of sequence 1 and m'_{00} of sequence 2. Then $\rho = f'/f$. To determine the relative phase, we crop a subsequence of temporal length f from $m_{00}(t)$, expand or contract it by ρ, then correlate with m'_{00}. The average lag of prominent peaks of the correlation result determines the relative phase. There may be a two-fold ambiguity in the phase from those viewpoints for which the autocorrelation of m_{00} yields peaks at half the stride frequency. For people close to the camera, the perspective effects are usually enough to uniquely determine the phase. For people far away, however, it can be difficult to distinguish between left foot forward or right foot forward on the basis of silhouette moment information alone.

After determining the temporal mapping between t and t', we now align the frieze patterns spatially. Given the moments that we can compute from frieze patterns, we determine the two translations and two scale factors that relate (x, y) and (x', y') for corresponding time steps in the two sequences. Dropping the time variables from the notation, this affine transformation is found to be

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \sqrt{\frac{\mu'_{20}\, m_{00}}{\mu_{20}\, m'_{00}}} & 0 \\ 0 & \sqrt{\frac{\mu'_{02}\, m_{00}}{\mu_{02}\, m'_{00}}} \end{bmatrix} \begin{bmatrix} x - m_{10}/m_{00} \\ y - m_{01}/m_{00} \end{bmatrix} + \begin{bmatrix} m'_{10}/m'_{00} \\ m'_{01}/m'_{00} \end{bmatrix}$$

Whether to allow both scale factors to vary independently for each time step, to enforce their ratio to be constant, to compute a temporal average for each, or other variations depends on the application and on the amount of noise one can expect in the underlying silhouette data.

5.2 Applications of Gait Alignment

We illustrate the utility of moment-based frieze alignment with two applications. The first involves comparing frieze tiles to classify a walking person's identity given a prior training set of gait data. The second application concerns matching a walking humanoid model to gait silhouette data from a real person, in order to locate specific body parts in each frame.

5.2.1 Human Identification

Given a dataset of gait sequences collected from one camera viewpoint, we want to analyze a new sequence to determine which person it is. Our approach is to create row and column silhouette projection friezes for each sequence, warp them all temporally to a canonical frequency and phase using the first half of the above alignment procedure, then cut out several tiles corresponding to individual strides from each sequence. These aligned frieze tiles are compared using normalized correlation, and subject classification is performed by nearest neighbor matching on correlation scores. This approach implicitly captures biometric shape cues such as body height/width ratio,

body-part proportions, stride length and amount of arm swing.

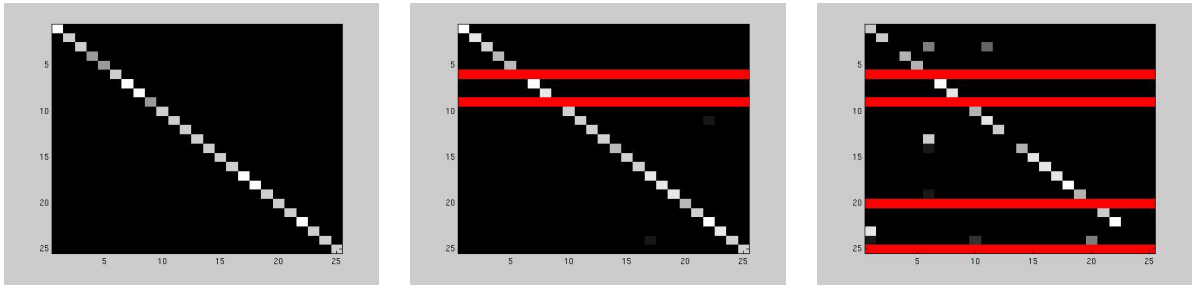

(a) (b) (c)

Figure 9: Confusion matrices for nearest neighbor classification of human identities (25 subjects) using gait frieze pattern tiles. (a) Result from training and testing on non-overlapping slow walking gait sequences. Classification rate is 100%. (b) Training on slow walk, testing on fast walk. Classification rate is 100%. (c) Training on slow walk, testing on walking carrying a ball (to inhibit arm swing). Classification rate is 81%.

To test this approach, we use the CMU MoBo database [5], which contains motion sequences of 25 subjects walking on a treadmill. Each subject is recorded performing four different types of walking: slow walk, fast walk, inclined walk, and slow walk holding a ball (to inhibit arm swing). Figure 9 shows results achieved for side views, for gait combinations slow-slow, slow-fast and slow-ball. For the slow-slow experiment, the gallery consisted of tiles from the first five seconds of each subject's slow walk gait sequence, and the probe set consists of tiles from the last five seconds of the same sequences. For both slow-fast and slow-ball, the classification algorithm is trained on all tiles from the slow walk sequences, and tested on all tiles from the other two gait sequences. We see that the results are quite good, even across different gait types. Although the match similarity metric is simple normalized correlation, each tile succinctly represents both spatial and temporal information from an entire stride subsequence.

5.2.2 Model-Based Body Part Analysis

Assume that we know the camera viewpoint, and have rendered a walking humanoid model from that viewpoint. We now have a sequence of model body silhouettes that can be matched against a real gait sequence. After spatio-temporal gait alignment, the temporal pairing of each frame of the data sequence with a corresponding model frame is known, along with a four parameter affine transformation that aligns those two binary silhouettes. Thus, for each frame, we can project the model silhouette contour onto the data silhouette image. A sample frame showing an overlayed

model contour found through automatic gait sequence alignment is shown in Figure 10A. The aligned model contour does not exactly coincide with the person's body outline due to a variety of factors, including differences in body shape and joint angle kinematics between the avatar and the human being (e.g. body proportions and amount of arm swing), as well as small differences in camera perspective between the model and data viewpoints. However, note that the overall temporal and spatial alignment is quite good, in the sense that the aligned model tells us what body parts should be visible, and roughly where they should appear in the image. More importantly, we know which body parts are occluded and should not be considered for further analysis in this frame.

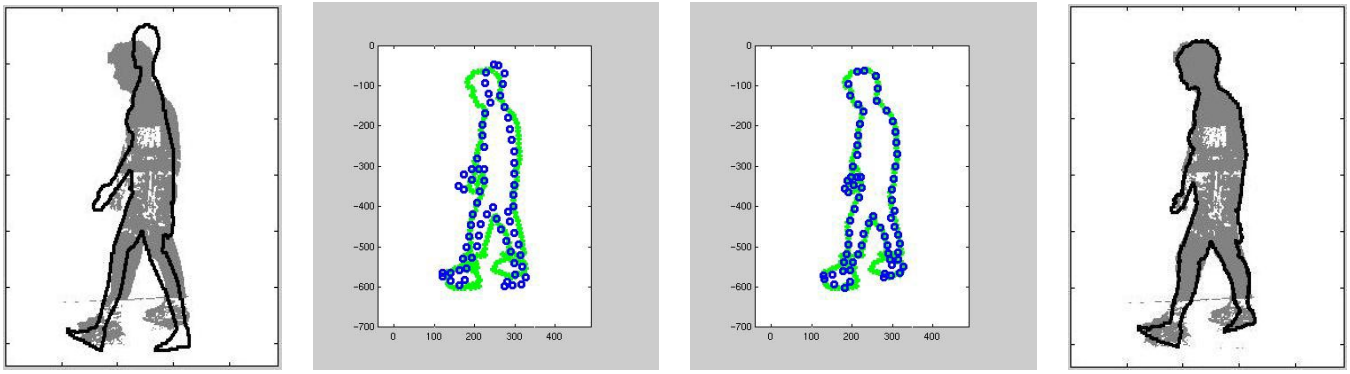

(a) (b) (c) (d)

Figure 10: (a) Moment-based alignment of model and data silhouettes. (b) Sampled points from model and data silhouette contours. (c) Results of non-rigid thin-plate spline alignment of the two sets of sample points. (d) Model silhouette warped by thin-plate spline transform, overlayed on data silhouette.

To illustrate what can potentially be done given this initial alignment between model and data silhouettes, we uniformly sample points from along each silhouette contour and use a program for non-rigid point matching to compute a thin-plate spline transformation between them [6][2]. Figure 10 shows, from left to right, the initial model contour alignment, the two sampled point sets, the resulting point sets after warping non-rigidly by a thin-plate spline, and the new warped model contour overlayed over the data silhouette. The agreement between contours is now much improved. The success of the non-rigid point matcher in this case is due in large part to the accuracy of the model silhouette topology, as determined by moment-based alignment of gait frieze patterns.

[2]The demo code for this technique was downloaded from `http://noodle.med.yale.edu/~chui/tps-rpm.html` and was found to work on our sampled point data "as is", with no reseting of parameters away from their default values.

More examples are shown in Figure 11. Model-based gait analysis using frieze patterns offers an efficient alternative to kinematic body part tracking for determining the location of individual body parts in each frame of a gait sequence.

6 Summary

We have presented a periodic pattern representation for analyzing gait sequences. Silhouette row and column projections are stacked over time to form frieze patterns that can be analyzed using the mathematical theory of symmetry groups. With the help of a walking humanoid avatar, we have studied the correlation between the seven frieze symmetry groups and gait viewing direction, and have had to develop practical techniques for classifying imperfect frieze patterns. Our future work will explore methods for more efficient and accurate viewpoint estimation from frieze patterns, and extend our mathematical methods for imperfect pattern analysis to patterns that are periodic along two dimensions. We have also presented a moment-based method for aligning frieze gait patterns both temporally and spatially. The method has applications in determining human identity from gait biometrics, and it provides an efficient alternative to frame-by-frame tracking approaches for locating and delineating body parts. We plan to apply this technique within a multi-camera motion capture system to perform 3D body reconstruction and motion analysis for human identification and sports entertainment.

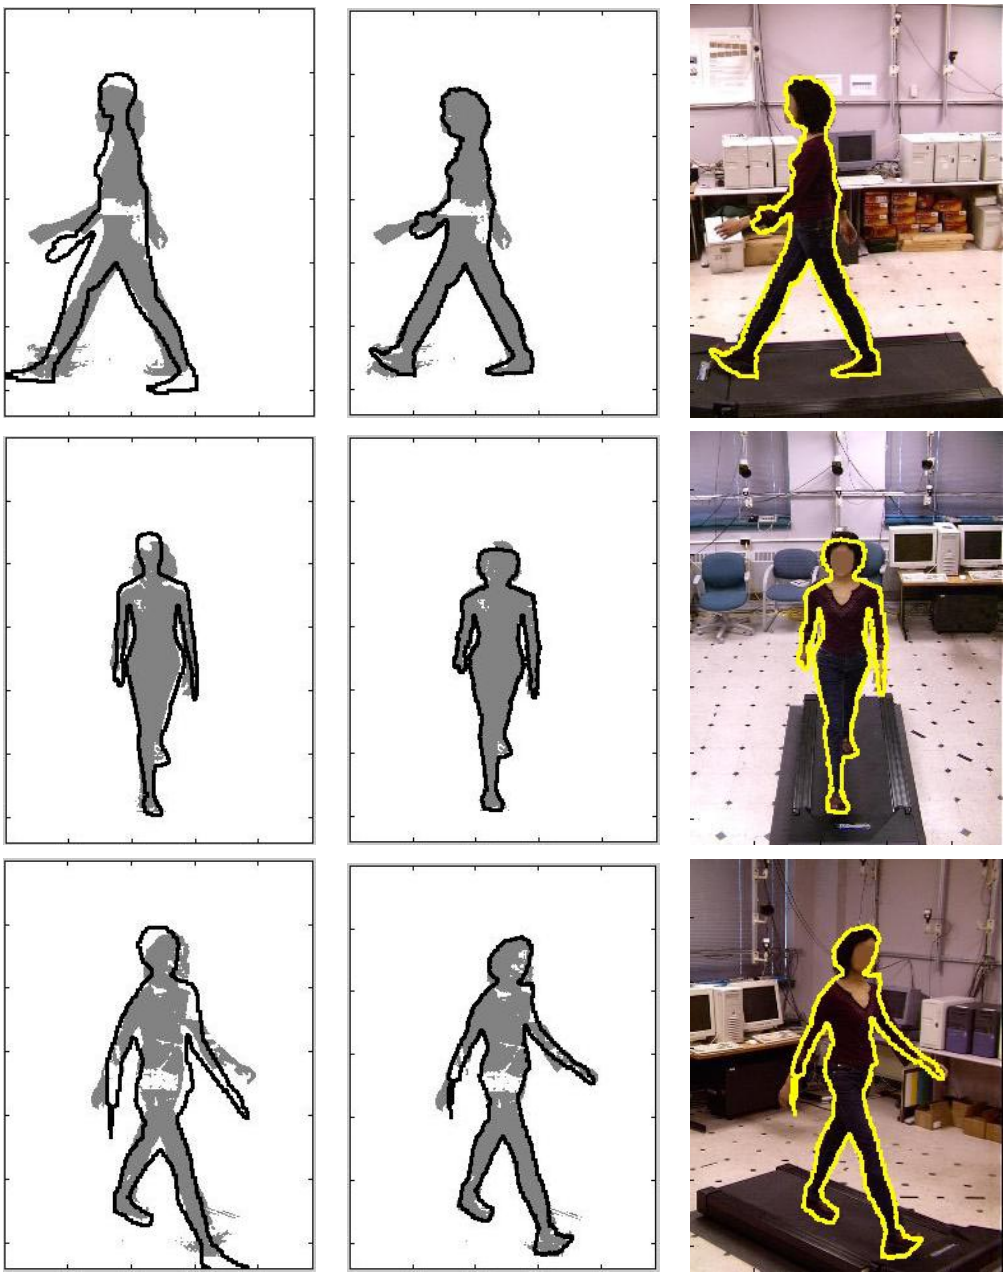

(a) (b) (c)

Figure 11: Three views (one per row) taken simultaneously by synchronized cameras. From left to right in each row, (a) Moment-based alignment of model and data silhouettes. (b) Model silhouette warped by computed thin-plate spline transform. (c) Warped model silhouette overlayed on original image. We plan to use results like these for further model-based body part analysis, ultimately leading to a sequence of 3D body reconstructions.

References

[1] B.K.P.Horn. *Robot Vision*. MIT Press, 1986.

[2] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 8–15, 1998.

[3] T.J. Cham and J.M. Rehg. A multiple hypothesis approach to figure tracking. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages II:239–245, 1999.

[4] R. Cutler and L.. Davis. Robust real-time periodic motion detection, analysis, and applications. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, To Appear.

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

[6] H.Chui and A.Rangarajan. A new algorithm for non-rigid point matching. *IEEE Computer Vision and Pattern Recognition*, pages 44–51, 2000.

[7] D. Hogg. Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.

[8] K. Kanatani. *Statistical Optimization for Geometric Computation : Theory and Practice (Machine Intelligence and Pattern Recognition, Vol 18)*. North-Holland, 1996.

[9] K. Kanatani. Comments on "Symmetry as a Continuous Feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):246–247, 1997.

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

[11] F. Liu and R. W Picard. Finding periodicity in space and time. In *IEEE International Conference on Computer Vision (ICCV)*, 1998.

[12] Y. Liu and R. T. Collins. A Computational Model for Repeated Pattern Perception using Frieze and Wallpaper Groups. In *Computer Vision and Pattern Recognition Conference*, pages 537–544, Los Alamitos, CA, June 2000. IEEE Computer Society Press. (http://www.ri.cmu.edu/pubs/pub_3302.html).

[13] Y. Liu and R.T. Collins. Frieze and wallpaper symmetry groups classification under affine and perspective distortion. Technical Report CMU-RI-TR-98-37, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 1998.

[14] S.A. Niyogi and E.H. Adelson. Analyzing and recognizing walking figures in xyt. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 469–474, 1994.

[15] S.M. Seitz and C.R. Dyer. View-invariant analysis of cyclic motion. *International Journal of Computer Vision*, 25:1–23, 1997.

[16] H. Zabrodsky, S. Peleg, and D. Avnir. Symmetry as a continuous feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(12):1154–1165, December 1995.