Finding Moving Shapes by Continuous-Model Evidence Gathering

Michael G. Grant, Mark S. Nixon and Paul H. Lewis
Department of Electronics and Computer Science
University of Southampton, Southampton SO17 1BJ, UK
(mgg97r|msn|phl)@ecs.soton.ac.uk

Abstract

Two recent approaches are combined in a new technique to find moving arbitrary shapes. We combine the Velocity Hough Transform, which extracts moving conic sections, with a continuous formulation for arbitrary shape extraction, which avoids discretisation errors associated with GHT methods. The new approach has been evaluated on synthetic and real imagery and is demonstrated to provide motion analysis that is resilient to noise and to be able to detect its target shapes, which are both moving and arbitrary. Further, it is shown to have performance advantages over contemporaneous single-image extraction techniques. Finally, it appears to offer improved immunity to noise and occlusion, consistent with evidence gathering techniques, as shown by results on real images.

1 Introduction

We describe a new technique to extract moving arbitrary shapes. This new technique has been created by fusing two evidence-gathering techniques, one for moving shapes and one for arbitrary shapes. Previous approaches to the problem of extracting moving shapes typically use a frame-by-frame approach where individual frames are analysed to locate the target shape. The data from each frame are then put together (e.g. using linear regression) to give a description of the motion of the shape. A large variety of methods has been developed previously to track shapes from frame to frame (including Kalman filters, neural networks, and a new velocity snake [1]). The Velocity Hough Transform (VHT) [2] is a recent evidence gathering approach that considers the whole image sequence as a single block of data and extracts the motion and structural parameters simultaneously. However, the VHT is limited by its inability to extract arbitrary shapes - it can only extract analytically parameterised shapes, such as lines, circles and ellipses. A modified version of the VHT isolates moving articulated objects (human subjects), by evidence gathering [3]. The model used combines two straight lines (representing legs) and a motion model that represents the swinging of the legs (by articulating the lines) and the motion of the hips. A model-based scenario was used in this extension, consistent with the original VHT, as opposed to an arbitrary one.

To find moving non-analytic shapes, there is a need to extend the VHT to handle general shapes such as in the Generalised Hough Transform (GHT) [4]. Rather than

BMVC 1999 doi:10.5244/C.13.55

using the GHT, we have chosen to use a new approach [5] that represents shapes using a continuous template described with Fourier descriptors which thus avoids the discretisation errors that are common with the GHT.

A future application of the new technique is the identification and recognition of walking subjects [6]. Gait is an attractive biometric because it can be measured non-invasively and remotely. There are presently two means of calculating gait biometrics - statistically and using model based techniques. The former method uses statistical techniques such as moments to measure changes in an image sequence of a walking person to derive a recognition metric. This method does not directly use the motion of walking to aid in the location or identification of the person but rather uses the motion content of the images as a whole [7, 8]. The latter approach uses models to identify regions on the body (e.g. the legs) and to make measurements directly upon these extracted regions (e.g. hip rotation cycles). One approach uses the HT for straight lines to extract the locations of a leg in a sequence [9] and, by analysing the angular motion, generates gait signatures for recognition, automatically.

The new technique will fit into either recognition category. For model-based recognition, we intend that the new technique will be used to extract the location and orientation of legs, and generate a gait signature as discussed later. For statistical-based recognition, the approach isolates the moving human body, as such giving a primer to extract features for statistical based techniques, and thus extending the stock of techniques for this purpose. The structure of the paper is as follows: Section 2.1 presents the VHT algorithm and Section 2.2 the new Fourier descriptor-based GHT. In Section 2.3 these are combined in a novel algorithm that simultaneously extracts motion and structural parameters for an arbitrary shape with linear motion. The new technique inherits the robustness of the original Hough Transform and is shown to be capable of withstanding considerable noise and occlusion in Section 3. Finally, in Sections 4 and 5, we suggest possible avenues of future research and present our conclusions.

2 Extracting moving objects

2.1 The VHT

A standard Hough transform based tracking algorithm will find instances of an object individually in each frame and use a technique such as linear regression to link each separate result into a description of the motion. If an image in the sequence is noisy or occluded and an incorrect location is extracted, this adversely influences the determination of the motion parameters. In contrast, the essential idea of the VHT [2] is to gather evidence from a sequence of images and to accumulate all the data from the sequence into a single parameter space, concurrently extracting optimal structural and motion parameters from the resulting data.

Moving objects in a sequence have temporal correlation that is often unused by many tracking algorithms. Using this, the VHT is more robust than standard tracking implementations, particularly when the object is occluded or the frames are noisy. By extending the evidence gathering across the entire sequence, deficiencies in some frames are offset by redundant information in others (e.g. structural information in the target shape that is frequently repeated in each frame). A circle (centre $c_{x,y}$, radius r) moving with linear motion is described by co-ordinates a_x and a_y as:

$$\begin{aligned} a_x &= c_x + r \cdot \cos(\theta) + v_x \cdot t \\ a_y &= c_y + r \cdot \sin(\theta) + v_y \cdot t \end{aligned} \qquad 1$$

where v_x and v_y are the velocities of the circle along the x and y axes, and t is the time reference of the frame relative to time $t = 0$ (the initial frame). The time reference of each frame is used to calculate how far back to push the vote co-ordinates along the path of motion. For example, if a circle moves forward at 1 pixel per frame, at $t=3$ the centre of the circle of votes in the accumulator should be moved 3 pixels back to ensure the votes pass through the initial circle centre co-ordinates. These equations calculate the vote co-ordinates in an accumulator for a feature point in a given frame of a sequence. After processing, the peak in the accumulator represents the best estimates of the circle's parameters and its centre co-ordinates at time $t = 0$.

2.2 Template representation with Fourier descriptors

In the GHT, the template shape is represented by an R-table, which is a discrete lookup table storing an angle and radius from a predefined reference point to each point on the border of the shape [4]. The problems with this representation are well described in the literature [10], but essentially derive from the fact that it is a discrete representation sampled at a particular scale. When the template is scaled or rotated, there are problems with aliasing and rounding errors. Figure 2 shows the effects of scaling and rotating the discrete set of points comprising the shape in Figure 1. Clearly, the new sets of points can have missing data and show the effects of discretisation.

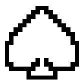

Figure 1: Original image

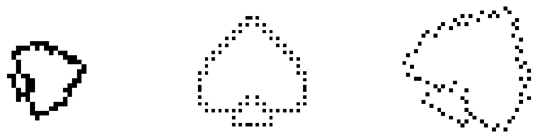

Figure 2: Rotated and scaled images

This difficulty can be avoided by using Fourier descriptors (FDs), which give a continuous representation that can be sampled at any resolution without the aliasing problems of the R-table. The actual FDs used are elliptic Fourier descriptors [11], which were chosen for their completeness, simple geometric interpretation and the fact that they can be simply produced from a chain code of a contour.

A curve defined by two sets of orthogonal co-ordinates, $c_x(s)$ and $c_y(s)$, parameterised by $s \in [0,2\pi)$ has elliptic Fourier descriptors as follows:

$$
\begin{aligned}
a_{xk} &= \frac{1}{\pi}\int_{-\pi}^{\pi} c_x(s)\cos(ks)ds \quad & a_{yk} &= \frac{1}{\pi}\int_{-\pi}^{\pi} c_y(s)\cos(ks)ds \\
b_{xk} &= \frac{1}{\pi}\int_{-\pi}^{\pi} c_x(s)\sin(ks)ds \quad & b_{yk} &= \frac{1}{\pi}\int_{-\pi}^{\pi} c_y(s)\sin(ks)ds
\end{aligned} \tag{2}
$$

These four parameters describe an ellipse, with k being the harmonic number. The range of k defines how many ellipses are used to represent a model shape and, thus, how accurate a shape representation the FDs allow. In the general case, using more harmonics gives a better representation of a shape. However, FDs are mostly used to describe a model taken from a discrete image, which must have a limited sampling rate defined by the pixel granularity. Before FDs can be used to draw the shape in the accumulator, they must be converted from the frequency domain to vectors in the spatial domain. The FDs can be converted to vectors (along the x- and y-axes) from the origin to a point on the curve by the following expressions:

$$v_x\left(s,\overline{FD}_x\right)=\sum_{k=1}^{n} a_{xk}\cos(ks)+b_{xk}\sin(ks)$$
$$v_y\left(s,\overline{FD}_y\right)=\sum_{k=1}^{n} a_{yk}\cos(ks)+b_{yk}\sin(ks) \qquad 3$$

where n is the number of harmonics used in the FDs:

$$\overline{FD}_x=\left\{a_{x_1},b_{x_1},a_{x_2},b_{x_2}\ldots,a_{x_n},b_{x_n}\right\}$$
$$\overline{FD}_y=\left\{a_{y_1},b_{y_1},a_{y_2},b_{y_2}\ldots,a_{y_n},b_{y_n}\right\} \qquad 4$$

Note that the DC terms have been omitted (by not summing for $k=0$), which translates the curve so that its centre is at the origin. These co-ordinates can be rotated and scaled by the following expressions:

$$R_x(s,l,\rho)=lv_x\left(s,\overline{FD}_x\right)\cos(\rho)-lv_y\left(s,\overline{FD}_y\right)\sin(\rho)$$
$$R_y(s,l,\rho)=lv_x\left(s,\overline{FD}_x\right)\sin(\rho)+lv_y\left(s,\overline{FD}_y\right)\cos(\rho) \qquad 5$$

where l is the scale factor and ρ is the angle of rotation (measured anticlockwise). Following from the formal theory of the HT [5], we require a kernel that defines the shape of votes to be laid down in the accumulator for each feature point (e.g. an edge pixel). This is a combination of curves, each with its origin on the reference point to be voted for (typically at the centre of the template shape), at a number of orientations and scales (for similarity transform invariance). This combination can be obtained from:

$$\overline{\varpi}(s,l,\rho)=R_x(s,l,\rho)U_x+R_y(s,l,\rho)U_y \qquad 6$$

where U_x and U_y are two orthogonal vectors defining the x- and y- axis respectively. This curve is inserted into the accumulator by offsetting it from the co-ordinates of each feature point in the image I, which is defined by:

$$I=\left\{\overline{\lambda}(t)\mid t\in D_I\right\} \qquad 7$$

Here, $\overline{\lambda}(t)$ is a parametric function that defines the points in the image, where a suffix on the domain indicates its extent (i.e. here, D_I is the domain of the image). The following expression defines how votes are placed in the accumulator for an image I:

$$A_t=\left\{\overline{\lambda}_t-\overline{\varpi}(s,l,\rho)\mid s\in D_s\right\}, t\in D_t \qquad 8$$

These equations describe the concept of the HT but do not formalise the actual technique used, namely the accumulation phase. Parameter space can be mapped into an accumulator by using a matching function, which determines whether a point, $\overline{c}$, in parameter space should be incremented for a point, $\overline{d}$, in the set A_t. The equation below defines the simplest accumulation strategy, that of incrementing an accumulator cell by unity for each match. Changing the matching function M can accommodate more complex strategies.

$$M\left(\overline{c},\overline{d}\right)=\begin{cases}1 & \text{if } \overline{c}=\overline{d}\\ 0 & \text{if } \overline{c}\neq\overline{d}\end{cases} \qquad 9$$

Next, this function is used on A_t for a range of parameter values. The expression for this defines the HT for arbitrary shapes using FDs. The continuous form is:

$$S_F\left(\overline{b},l,\rho\right)= M\left(\overline{b},\ \overline{\lambda}(t)-\overline{\varpi}(s,l,\rho)\right)ds\,dt \qquad 10$$

where $\bar{b}$ is the translation vector (i.e. the location of the reference point). To implement this on a computer it needs to be in discrete form. This is achieved by converting integrals to summations and limiting the range of the parameters being examined, as follows:

$$S_{DF}\left(\bar{b}, l, \rho\right) = \sum_{t \in D_t}\sum_{s \in D_s} M\left(\bar{b},\ \bar{\lambda}(t) - \bar{\varpi}(s, l, \rho)\right) \qquad 11$$

This expression defines the HT for arbitrary shapes using FDs. Essentially, for a given feature point, a locus of points is plotted through the four-dimensional (x centre location, y centre location, orientation and scale) accumulator space. This locus is formed from scaled and rotated representations of the template in two-dimensional planes along the x- and y-axes of the accumulator.

Instead of recovering individual vote co-ordinates from an R-table, as in the GHT, they are calculated from the FDs. Since the FDs are continuous and can be sampled at any resolution, this avoids the problems of discretisation. Clearly, if the original template is smaller than the reconstructed one, the FDs will provide no additional detail as they only provide a continuous representation of the original template.

2.3 Moving arbitrary shape extraction

To produce an algorithm capable of extracting arbitrary shapes undergoing motion, the Fourier descriptor version of the GHT is extended in the same way as the HT for circles was extended into the VHT. Instead of drawing a circle in the accumulator (as in the VHT), the Fourier descriptors are used to trace a locus of votes in the form of the template shape, adjusted for the motion of the object relative to the time reference of each frame. Once the voting process is complete, peaks in the accumulator indicate the location (at time $t = 0$) of an instance of the template shape moving linearly.

The new technique combines the VHT approach with the HT for arbitrary shapes using FDs. This complementary blending of techniques maximises the robustness of the composite algorithm. The FDs minimise discretisation error in the plotting of the template shape and the VHT exploits the temporal correlation across a sequence, increasing the amount of valid evidence and reducing the effect of noise and occlusion. In common with the VHT, this new technique does not need initialisation nor does the problem of correspondence need to be addressed directly.

Since the technique extracts arbitrary shapes, the theory is closer to the HT for arbitrary shapes. The first difference is that image sequences are processed instead of single frames. Consequently, the image I is replaced with an image sequence IS, defined below:

$$IS = \left\{\bar{\lambda}(t, f) \mid t \in D_I, f \in D_f\right\} \qquad 12$$

where t is a parameter to define which part of the image is required and f is the frame time (where $f = 0$ is the start frame of analysis). The template representation is identical to the original (static) version until it starts to describe the kernel of the HT. In the new technique, this is different because velocity parameters are also being accumulated. Hence, the kernel is:

$$\bar{\varpi}\left(s, f, l, \rho, v_x, v_y\right) = R_x(s, l, \rho)U_x + R_y(s, l, \rho)U_y + fv_xU_x + fv_yU_y \qquad 13$$

v_x and v_y are respectively the x centre and y centre velocity parameters. The accumulator vote-pattern expression is then:

$$A_t = \left\{\bar{\lambda}_{t,f} - \bar{\varpi}\left(s, f, l, \rho, v_x, v_y\right) \mid s \in D_s\right\}, t \in D_t, f \in D_f \qquad 14$$

and the expression for the new technique in continuous form is:

$$S_F\left(\bar{b}, l, \rho, v_x, v_y\right) = M\left(\bar{b}, \bar{\lambda}(t, f) - \bar{\omega}(s, f, l, \rho, v_x, v_y)\right) ds\ dt\ df \qquad 15$$

Finally, in discrete form the expression is:

$$S_{DF}\left(\bar{b}, l, \rho, v_x, v_y\right) = \sum_{f \in D_f} \sum_{t \in D_t} \sum_{s \in D_s} M\left(\bar{b}, \bar{\lambda}(t, f) - \bar{\omega}(s, f, l, \rho, v_x, v_y)\right) \qquad 16$$

This gives an accumulation strategy for finding moving arbitrary shapes. The voting algorithm is a combination of the VHT and the Fourier descriptor version of the GHT. For each edge pixel in a frame, a locus of votes is calculated from the Fourier description of the template shape and entered into the unified accumulator. The co-ordinates of loci are adjusted to allow for the predicted motion of the shape, dependent on the frame time, as in the VHT.

The other notable point about the algorithm is that the complete range of votes is made in the accumulator (i.e. the entire shape is traced in the accumulator, as in the Merlin-Farber method [12]). In the GHT, the only template points allowed to vote are those with the same gradient direction as the edge pixel being processed. With good gradient direction data, this restriction removes a lot of unnecessary votes (and hence noise) from the accumulator. The voting algorithm could be changed to perform the same reduction of votes as the GHT by incorporating a function that calculates the gradient direction at a point on the Fourier-described curve. This value would then be compared against edge pixel gradient direction to restrict the votes cast into the accumulator.

The new technique constructs a 6-dimensional accumulator space where the parameters are the centre co-ordinates, scale, orientation and the velocity along both directions of freedom. The maximum peak is indexed by values for the co-ordinates of the moving shape that best matched the image data. The accumulator requires a large amount of memory, even for small images. This is unavoidable if the algorithm is to be used in its most basic form. However, most of the HT-related speed-up and memory-reduction modifications are applicable. Some examples are parameter space decomposition [13], randomised sampling algorithms [14] or the use of genetic algorithms [3].

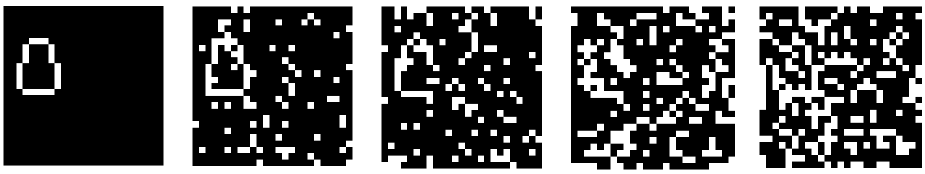

Figure 3: Original image and 40%, 60%, 80% and 100% added Gaussian noise

3 Results

3.1 Gaussian noise testing

The new algorithm was run on a five-frame sequence based on the small (20x20) image (in Figure 3) moving linearly along the x-axis at a velocity of one pixel per frame. A small image was chosen to reduce memory requirements and make computation of large-scale tests practical. Noise was added at random to each frame of the sequence at eleven noise levels from 0% random coverage to 100% random coverage of the frame. The noise distribution was zero-mean Gaussian with a standard deviation of three.

Examples of the effects of the increasing noise levels can be seen in Figure 3 where the grey-level images produced were thresholded. Note that at the maximum noise level the shape is completely obliterated.

The results for the normal GHT tracking algorithm were generated using a standard GHT on each frame of the sequence and using linear regression on the results to calculate the velocity terms. The testing conditions were as described above.

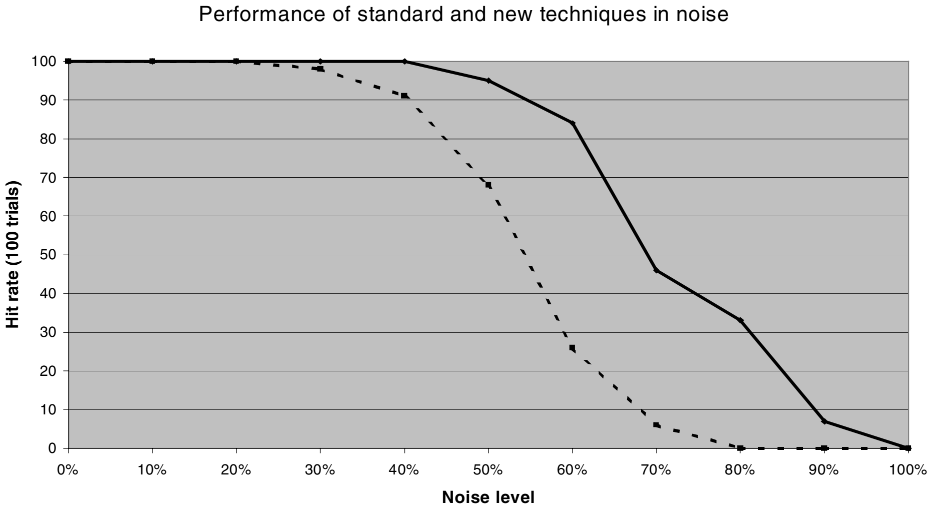

Figure 4: Noise performance (dashed = GHT-based, solid = new technique)

The graph above shows the new technique is significantly more accurate than the GHT-based technique. By accumulating the additional temporally correlated evidence the new technique is able to handle noise levels that are approximately 20 percent greater than the standard technique.

The irregularity in the results for the new algorithm at 80% noise is probably an anomaly - the result shown for that amount of noise seems to be a little too accurate for the noise level. A result of about 25 votes would be more in line with the general curve and only represents an overshoot of 7 “hits”, which is not significant enough to conflict with the difference being chance.

3.2 Occlusion testing

Figure 5: Frames 1, 3 and 5 of occluded sequence (bar is 5 pixels wide)

A simple test of the effects of occlusion was carried out on the five-frame sequence described above. No speckle noise was added to the sequence since this would be an unnecessary factor. Instead, a number of vertical lines of pixels starting from column 2 were blanked out and both algorithms run on the resulting image. Figure 5 shows some frames from the occluded sequence (the occlusion bar is shown in grey to make it visible - it is normally black).

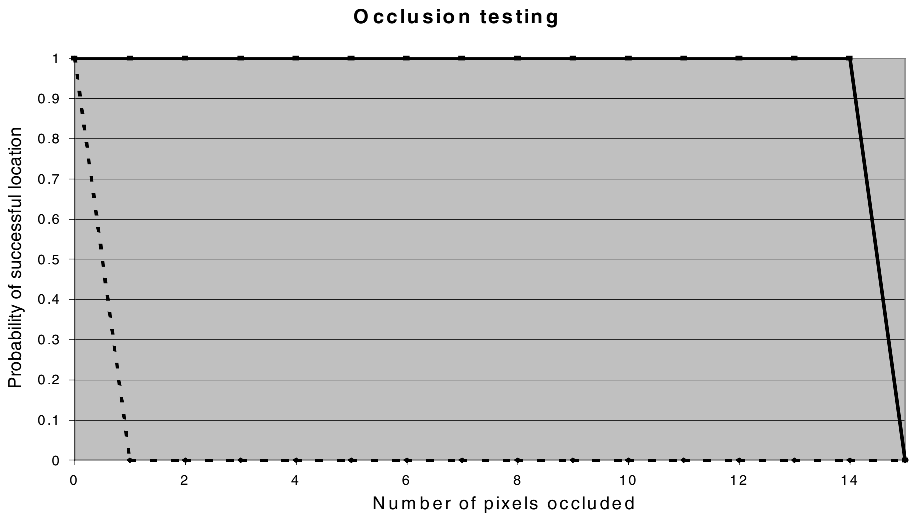

Figure 6: Occlusion Tests (dashed = GHT-based, solid = new technique)

The results reveal that the new technique keeps track of the shape until the blanking is 15 pixels wide – which obscures the shape for the duration of the entire sequence. The GHT based algorithm failed as soon as any blanking was introduced. This failure reveals more about the algorithm's implementation than about its resilience to occlusion. The current implementation uses the estimated location of the template shape in every frame as an input to the linear regression stage. Therefore, when a frame is corrupted and gives an incorrect result, the output of the linear regression stage is affected causing a "global" estimation error. A more intelligent implementation might include a heuristic that ignores frames giving evidence inconsistent with the majority of frames.

3.3 Finding People

For purposes of illustration, the technique is now applied to locate a moving human body in a sequence of images. The current implementation of the new technique locates rigid shapes moving with linear velocity. Clearly, its formulation is general so plastic deformation could be included, as it was for pulsating arteries in the original VHT formulation [2]. In the case of a human walking, the torso is approximately a constant shape and, if the camera is far enough away, the bobbing motion of gait is minimal. Consequently, it is possible to detect people using the technique in its current form by searching for the torso. However, no meaningful gait data can be gathered from just the location of the torso so this method of locating a human is only useful as a primer for another technique. Nonetheless, using the new technique to locate a human demonstrates that it is equally applicable to real world images.

Self-occlusion of the body due to the motion of the arms and legs is a problem that affects the performance of many person-tracking algorithms. By the nature of evidence gathering, the new algorithm copes with occlusions that do not reduce the number of correct votes (from the remainder of the true contour) below the level of noise in the image. Consequently, there is no need for special precautions in the new technique.

Figure 7: Template shape and frames 1, 3, 5, 7 and 10 of walker sequence

Figure 7 shows a reconstructed template of a walker's torso, which was originally created by manually tracing the torso in the first frame of the sequence. To accurately reconstruct the template requires a large number of Fourier descriptors because of the discontinuities at the end of the open curve. Also shown in Figure 7 are several frames of the walker sequence with the template superimposed. During the first part of the sequence, the walker's location is accurately extracted - the initial location is exact and the extracted speed (thirteen pixels per second) is correct. Shortly after frame seven the walker rises up on his leg (vertically, there is a rise of fifteen pixels), which will cause the votes to "miss" in the accumulator, since this movement has not been accommodated in the evidence gathering strategy. The subject also slows down slightly thus violating the assumption of linear velocity and consequently the template "overtakes" the walker - frame 10 shows it some pixels ahead.

4 Further Work

As it stands, the new technique is suitable as a primer for other gait recognition techniques. To enable the new method to be used as a primary gait recognition technique, there need to be improvements in the motion model. Introducing articulation, so that two legs can be tracked, will enable a gait signature to be calculated from hip rotation cycles. This would improve on the work in [9] since it would be possible to model the legs directly rather than with straight lines.

Another possible direction for development of the algorithm is towards improving the generality of the tracking of arbitrary objects. The eventual aim is to be able to extract arbitrary shapes undergoing arbitrary motion. The simplistic approach is to increase the complexity of the HT kernel to represent an increasingly complex curved motion path. However, an accurate polynomial description of an arbitrary path will require a large number of terms, which massively increases the dimensionality of the problem. This has parallels with the early development of the HT for arbitrary shapes. The solution presented by Merlin and Farber was to use templates, losing the flexibility of extracting all possible arbitrary shapes for the ability to efficiently find a specified arbitrary shape. Similarly, we feel there is a case for the use of motion “templates” to avoid the curse of increasing dimensionality, although the issue is less clear cut than in the shape specification case.

5 Conclusions

We have presented a new technique that robustly estimates structural and motion parameters for specified moving arbitrary shapes in an image sequence. The technique has excellent tolerance to noise and occlusion in comparison to the standard GHT based tracking technique. Discretisation errors are also minimised in the accumulator by the use of Fourier descriptors to represent template shapes in a continuous form, which eliminate common problems to do with rotation and scaling of templates. Further work will aim to capitalise on the advantages of this new technique, especially in gait recognition and biometric extraction.

References

[1] N. Peterfreund, "The Velocity Snake: Deformable Contour for Tracking in Spatio-Velocity Space", *CVIU* **73**(3), pp 346-356, 1999

[2] J. M. Nash, J. N. Carter and M. S. Nixon, Dynamic Feature Extraction via the Velocity Hough Transform, *Pattern Recog. Letters*, **18**(10), pp. 1035-1047, 1997

[3] J. M. Nash, J. N. Carter and M. S. Nixon, Extracting Moving Articulated Objects by Evidence Gathering, *Proc. BMVC 98*, **2**, pp. 609-618, Sept. 1998

[4] D. H. Ballard, "Generalising the Hough Transform to Find Arbitrary Shapes", *CVGIP* **13**, pp 111-122, 1981

[5] A. S. Aguado, M. S. Nixon and M. E. Montiel, Parameterising Arbitrary Shapes via Fourier Descriptors for Evidence-Gathering Detection, *CVIU*, **69**(2), pp 202-221, 1998

[6] M. S. Nixon, J. N. Carter, D. Cunado, P. S. Huang, S. V. Stevenage, Automatic Gait Recognition, *In*: A. K. Jain, R. Bolle and S. Pankanti Eds.: *Biometrics: Personal Identification in Networked Society*, pp. 231-250, 1999

[7] H. Murase and R. Sakai, Moving Object Recognition in Eigenspace Representation: Gait Analysis and Lip Reading, *Pattern Recog. Letters*, **17**, pp. 155-162, 1996

[8] P. S. Huang, C. J. Harris and M. S. Nixon, Comparing Different Template Features for Recognizing People by their Gait, *Proc. BMVC 98*, **2**, pp. 639-648, Sept. 1998

[9] D. Cunado, J. M. Nash, M. S. Nixon and J. N. Carter, Gait Extraction and Description by Evidence-Gathering, *Proc. 2^{nd} Int. Conf. on Audio- and Video-Based Biometric Person Authentication AVBPA99,* Washington D.C., pp 43-48, March 1999

[10] W. E. L. Grimson and D. P. Huttenglocher, "On Sensitivity of the Hough Transform for Object Recognition", *IEEE Trans. on PAMI*, **12**, pp 255-275, 1990

[11] F. P. Kuhl and C. R. Giardina, "Elliptic Fourier Feature of a Closed Contour", *Computer Graphics and Image Processing*, **18**, pp 236-258, 1982

[12] P. M. Merlin and D. J. Farber, "A Parallel Mechanism for Detecting Curves in Pictures", *IEEE Trans. on Computers*, **24**, pp 96-98 1975

[13] A. S. Aguado, M. E. Montiel and M. S. Nixon, "Improving Parameter Space Decomposition for the Generalised Hough Transform", *Proc.ICIP 96* **III**, pp 627-630, 1996

[14] L. Xu, E. Oja and P. Kultanen, "A New Curve Detection Method: Randomised Hough Transform (RHT)", *Pattern Recog. Letters* **11**, pp 331-338, 1990