Robust Real-Time Periodic Motion Detection, Analysis, and Applications

Ross Cutler and Larry S. Davis, *Fellow*, *IEEE*

Abstract—We describe new techniques to detect and analyze periodic motion as seen from both a static and a moving camera. By tracking objects of interest, we compute an object's self-similarity as it evolves in time. For periodic motion, the self-similarity measure is also periodic and we apply Time-Frequency analysis to detect and characterize the periodic motion. The periodicity is also analyzed robustly using the 2D lattice structures inherent in similarity matrices. A real-time system has been implemented to track and classify objects using periodicity. Examples of object classification (people, running dogs, vehicles), person counting, and nonstationary periodicity are provided.

Index Terms—Periodic motion, motion segmention, object classification, person detection, motion symmetries, motion-based recognition.

1 Introduction

Object motions that repeat are common in both nature and the man-made environment in which we live. Perhaps the most prevalent periodic motions are the ambulatory motions made by humans and animals in their gaits (commonly referred to as "biological motion" [16]). Other examples include a person walking, a waving hand, a rotating wheel, ocean waves, and a flying bird. Knowing that an object's motion is periodic is a strong cue for object and action recognition [16], [11]. In addition, periodic motion can also aid in tracking objects. Furthermore, the periodic motion of people can be used to recognize individuals [20].

1.1 Motivation

Our work is motivated by the ability of animals and insects to utilize oscillatory motion for action and object recognition and navigation. There is behavioral evidence that pigeons are well-adapted to recognize the types of oscillatory movements that represent components of the motor behavior shown by many living organisms [9]. There is also evidence that certain insects use oscillatory motion for navigational purposes (hovering above flowers during feeding) [17]. Humans can recognize biological motion from viewing lights placed on the joints of moving people [16]. Humans can also recognize periodic movement of image sequences at very low resolutions, even when point correspondences are not possible. For example, Fig. 1 shows such a sequence. The effective resolution of this sequence is 9×15 pixels (it was created by resampling a 140×218 (8-bit, 30fps) image sequence to 9×15 and back to 140×218 using bicubic interpolation). In this sequence, note the similarity between frames 0 and 15. We will use image similarity to detect and analyze periodic motion.

1.2 Periodicity and Motion Symmetries

We define the motion of a point $\vec{X}(t)$, at time t, periodic if it repeats itself with a constant period p, i.e.,

$$\vec{X}(t+p) = \vec{X}(t) + \vec{T}(t), \tag{1}$$

where $\vec{T}(t)$ is a translation of the point. The period p is the smallest $p > 0$ that satisfies (1); the frequency of the motion is $1/p$. If p is not constant, then the motion is cyclic. In this work, we analyze locally (in time) periodic motion, which approximates many natural forms of cyclic motion.

Periodic motion can also be defined in terms of symmetry. Informally, spatial symmetry is self-similarity under a class of transformations, usually the group of Euclidean transformations in the plane (translations, rotations, and reflections) [36]. Periodic motion has a temporal (and sometimes spatial) symmetry. For example, Figs. 3a, 4a, 5a, and 6a show four simple dynamic systems (pendulums). For each system, the motion is such that $\vec{X}(t+p) = \vec{X}(t)$ for a point $\vec{X}(t)$ on the pendulum. However, each system exhibits qualitatively different types of periodic motion. Fig. 5a is a simple planar pendulum with a fixed rod under a gravitational field. The motion of this system gives it a temporal mirror symmetry along the shown vertical axis. The system in Fig. 4a is a similar pendulum, but with sufficient initial velocity such that it always travels in one angular direction. The motion of this system gives it a temporal mirror symmetry along the shown vertical axis. The system in Fig. 3a is a similar pendulum, but in zero gravity; note it has an infinite number of axes of symmetry that pass through the pivot of the pendulum. The system in Fig. 6a consists of a pair of uncoupled and 180° out of phase pendulums, a system which is often used to model the upper leg motion of humans [24]. This system has a temporal mirror symmetry along the shown vertical axis, as well as an approximate spatial mirror symmetry along the

• *The authors are with the University of Maryland, AV Williams Blvd. Computer Science Department, College Park, MD 20742. E-mail: {rgc, lsd}@cs.umd.edu.*

Manuscript received 21 Apr. 1999; revised 26 Feb. 2000; accepted 28 Mar. 2000.
Recommended for acceptance by R. Collins.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 109643.

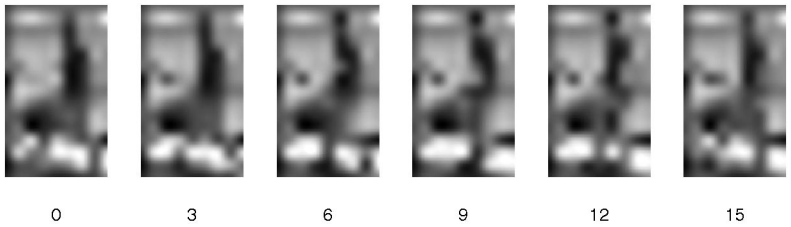

Fig. 1. Low resolution image sequences of a periodic motion (a person walking on a treadmill). The effective resolution is 9×15 pixels.

same vertical axis (it is approximate because the pendulums are not identical).

The above examples illustrate that while (1) can be used to detect periodicity, it is not sufficient to classify different types of periodic motion. For classification purposes, it is necessary to exploit the dynamics of the system of interest, which we do in Section 3.4.

1.3 Assumptions

In this work, we make the following assumptions: 1) the orientation and apparent size of the segmented objects do not change significantly during several periods (or do so periodically); 2) the frame rate is sufficiently fast for capturing the periodic motion (at least double the highest frequency in the periodic motion).

1.4 Contributions

The main contribution of this work is the introduction of novel techniques to robustly detect and analyze periodic motion. We have demonstrated these techniques with video of the quality typically found in both ground and airborne surveillance systems. Of particular interest is the utilization of the symmetries of motion exhibited in nature, which we use for object classification. We also provide several other novel applications of periodic motion, all related to automating a surveillance system.

1.5 Organization of the Paper

In Section 2, we review and critique the related work. The methodology is described in Section 3. Examples and applications of periodic motion, particularly for the automated surveillance domain, are given in Section 4. A real-time implementation of the methods is discussed in Section 5, followed by a summary of the paper in Section 6.

2 Related Work

There has been recent interest in segmenting and analyzing periodic or cyclic motion. Existing methods can be categorized as those requiring point correspondences [33], [35]; those analyzing periodicities of pixels [21], [30]; those analyzing features of periodic motion [27], [10], [14]; and those analyzing the periodicities of object similarities [6], [7], [33]. Related work has been done in analyzing the rigidity of moving objects [34], [25]. Below we review and critique each of these methods. Due to some similarities with the presented method, [33], [21], [30] are described in more detail than the other related work.

Seitz and Dyer [33] compute a temporal correlation plot for repeating motions using different image comparison functions, d_A and d_I. The affine comparison function d_A allows for view-invariant analysis of image motion, but requires point correspondences (which are achieved by tracking reflectors on the analyzed objects). The image comparison function d_I computes the sum of absolute differences between images. However, the objects are not tracked and, thus, must have nontranslational periodic motion in order for periodic motion to be detected. Cyclic motion is analyzed by computing the period-trace, which are curves that are fit to the surface d. Snakes are used to fit these curves, which assumes that d is well-behaved near zero so that near-matching configurations show up as local minima of d. The K-S test is utilized to classify periodic and nonperiodic motion. The samples used in the K-S test are the correlation matrix M and the hypothesized period-trace PT. The null hypothesis is that the motion is not periodic, i.e., the cumulative distribution function M and PT are not significantly different. The K-S test rejects the null hypothesis when periodic motion is present. However, it also rejects the null hypothesis if M is nonstationary. For example, when M has a trend, the cumulative distribution function of M and PT can be significantly different, resulting in classifying the motion as periodic (even if no periodic motion present). This can occur if the viewpoint of the object or lighting changes significantly during evaluation of M (see Fig. 19a). The basic weakness of this method is it uses a one-sided hypothesis test which assumes stationarity. A stronger test is needed to detect periodicity in nonstationary data, which we provide in Section 3.4.

Polana and Nelson [30] recognize periodic motions in an image sequence by first aligning the frames with respect to the centroid of an object so that the object remains stationary in time. Reference curves, which are lines parallel to the trajectory of the motion flow centroid, are extracted and the spectral power is estimated for the image signals along these curves. The periodicity measure of each reference curve is defined as the normalized difference between the sum of the spectral energy at the highest amplitude frequency and its multiples and the sum of the energy at the frequencies half way between.

Tsai et al. [35] analyze the periodic motion of a person walking parallel to the image plane. Both synthetic and real walking sequences are analyzed. For the real images, point correspondences were achieved by manually tracking the joints of the body. Periodicity was detected using Fourier analysis of the smoothed spatio-temporal curvature function of the trajectories created by specific points on the body as it performs periodic motion. A motion-based recognition application is described in which one complete cycle is stored as a model and a matching process is performed using one cycle of an input trajectory.

Allmen [1] used spatio-temporal flow curves of edge image sequences (with no background edges present) to analyze cyclic motion. Repeating patterns in the ST flow curves are detected using curvature scale-space. A potential problem with this technique is that the curvature of the ST flow curves is sensitive to noise. Such a technique would likely fail on very noisy sequences, such as that shown in Fig. 15.

Niyogi and Adelson [27] analyze human gait by first segmenting a person walking parallel to the image plane using background subtraction. A spatio-temporal surface is fit to the XYT pattern created by the walking person. This surface is approximately periodic and reflects the periodicity of the gait. Related work [26] used this surface (extracted differently) for gait recognition.

Liu and Picard [21] assume a static camera and use background subtraction to segment motion. Foreground objects are tracked and their path is fit to a line using a Hough transform (all examples have motion parallel to the image plane). The power spectrum of the temporal histories of each pixel is then analyzed using Fourier analysis and the harmonic energy caused by periodic motion is estimated. An implicit assumption in [21] is that the background is homogeneous (a sufficiently nonhomogeneous background will swamp the harmonic energy). Our work differs from [21] and [30] in that we analyze the periodicities of the image similarities of large areas of an object, not just individual pixels aligned with an object. Because of this difference (and the fact that we use a smooth image similarity metric), our Fourier analysis is much simpler since the signals we analyze do not have significant harmonics of the fundamental frequency. The harmonics in [21] and [30] are due to the large discontinuities in the signal of a single pixel; our self-similarity metric does not have such discontinuities.

Fujiyoshi and Lipton [10] segment moving objects from a static camera and extract the object boundaries. From the object boundary, a "star" skeleton is produced, which is then Fourier analyzed for periodic motion. This method requires accurate motion segmentation, which is not always possible (e.g., see Fig. 16). Also, objects must be segmented individually; no partial occlusions are allowed (as shown in Fig. 21a). In addition, since only the boundary of the object is analyzed for periodic change (and not the interior of the object), some periodic motions may not be detected (e.g., a textured rolling ball, or a person walking directly toward the camera).

Selinger and Wixson [34] track objects and compute self-similarities of that object. A simple heuristic using the peaks of the 1D similarity measure is used to classify rigid and nonrigid moving objects, which in our tests fails to classify correctly for noisy images (e.g., the sequence in Fig. 15).

Heisele and Wohler [14] recognize pedestrians using color images from a moving camera. The images are segmented using a color/position feature space and the resulting clusters are tracked. A quadratic polynomial classifier extracts those clusters which represent the legs of pedestrians. The clusters are then classified by a time delay neural network, with spatio-temporal receptive fields. This method requires accurate object segmentation. A 3-CCD color camera was used to facilitate the color clustering and pedestrians are approximately 100 pixels in height. These image qualities and resolutions are typically not found in surveillance applications.

There has also been some work done in classifying periodic motion. Polana and Nelson [30] use the dominant frequency of the detected periodicity to determine the temporal scale of the motion. A temporally scaled XYT template, where XY is a feature based on optical flow, is used to match the given motion. The periodic motions include walking, running, swinging, jumping, skiing, jumping jacks, and a toy frog. This technique is view dependent and has not been demonstrated to generalize across different subjects and viewing conditions. Also, since optical flow is used, it will be highly susceptible to image noise.

Cohen et al. [5] classify oscillatory gestures of a moving light by modeling the gestures as simple one-dimensional ordinary differential equations. Six classes of gestures are considered (all circular and linear paths). This technique requires point correspondences and has not been shown to work on arbitrary oscillatory motions.

Area-based techniques, such as the present method, have several advantages over pixel-based techniques, such as [30], [21]. Specifically, area-based techniques allow the analysis of the dynamics of the entire object, which is not achievable by pixel-based techniques. This allows for classification of different types of periodic motion, such as those given in Section 4.1 and Section 4.4. In addition, area-based techniques allow detection and analysis of periodic motion that is not parallel to the image plane. All examples given in [30], [21] have motion parallel to the image plane, which ensures there is sufficient periodic pixel variation for the techniques to work. However, since area-based methods compute object similarities which span many pixels, the individual pixel variations do not have to be large. For example, our method can detect periodic motion from video sequences of people walking directly toward the camera. A related benefit is that area-based techniques allow the analysis of low S/N images, such as that shown in Fig. 16, since the S/N of the object similarity measure (such as (5)) is higher than that of a single pixel.

3 Method

The algorithm for periodicity detection and analysis consists of two parts. First, we segment the motion and track objects in the foreground. We then align each object along the temporal axis (using the object's tracking results) and compute the object's self-similarity as it evolves in time. For periodic motions, the self-similarity metric is periodic and we apply Time-Frequency analysis to detect and characterize the periodicity. The periodicity is also analyzed robustly using the 2D lattice structures inherent in similarity matrices.

3.1 Motion Segmentation and Tracking

Given an image sequence I_t from a moving camera, we segment regions of independent motion. The images I_t are first Gaussian filtered to reduce noise, resulting in I_t^*. The image I_t^* is then stabilized [12] with respect to image $I_{t-\tau}^*$, resulting in $V_{t,t-\tau}$. The images $V_{t,t-\tau}$ and I_t^* are differenced and thresholded to detect regions of motion, resulting in a binary motion image:

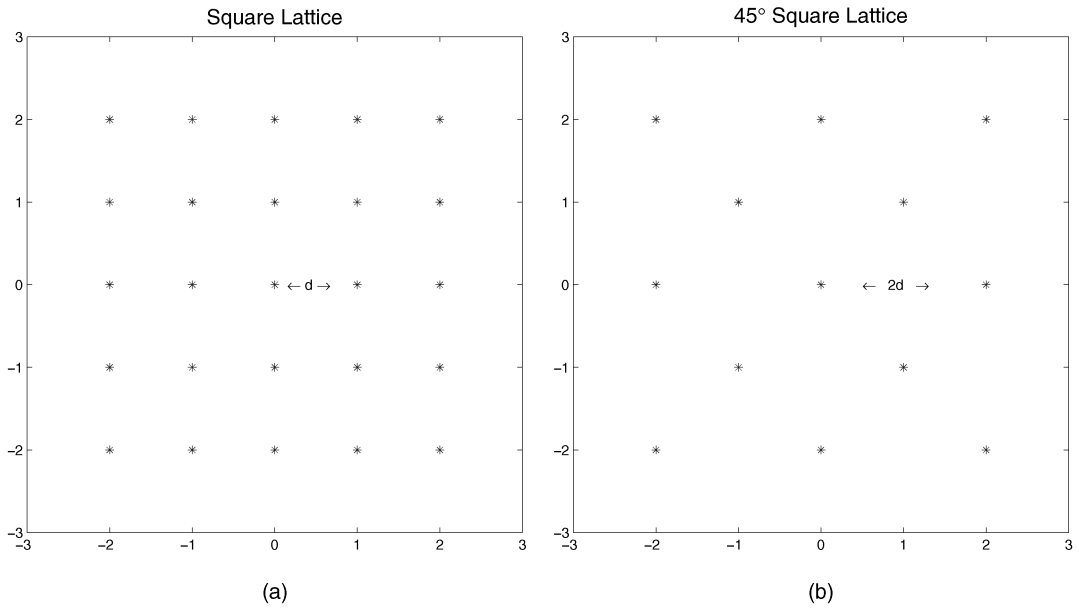

Fig. 2. Lattices used to match the peaks of the autocorrelation of S'. (a) Square lattice (b) 45° rotated square lattice.

$$\mathcal{M}_{t,-\tau} = \begin{cases} 1 & if \ \left| I_t^* - V_{t,t-\tau} \right| > T_M \\ 0 & otherwise, \end{cases} \quad (2)$$

where T_M is a threshold. In order to eliminate false motion at occlusion boundaries (and help filter spurious noise), the motion images $\mathcal{M}_{t,\tau}$ and $\mathcal{M}_{t,-\tau}$ are logically *and*ed together:

$$\mathcal{M}_t = \mathcal{M}_{t,-\tau} \wedge \mathcal{M}_{t,\tau}. \quad (3)$$

An example of $\mathcal{M}_t$ is shown in Fig. 21b. Note that, for large values of τ, motion parallax will cause false motion in $\mathcal{M}_t$. In our examples (for a moving camera), τ = 300 ms was used.

Note that, in many surveillance applications, images are acquired using a camera with automatic gain, shutter, and exposure. In these cases, normalizing the image mean before comparing images I_{t_1} and I_{t_2} will help minimize false motion due to a change in the gain, shutter, or exposure.

A morphological open operation is performed on $\mathcal{M}_t$ (yielding $\mathcal{M}_t^*$), which reduces motion due to image noise. The connected components for $\mathcal{M}_t^*$ are computed and small components are eliminated (further reducing image noise). The connected components which are spatially similar (in distance) are then merged and the merged connected components are added to a list of objects O_t to be tracked. An object has the following attributes: area, centroid, bounding box, velocity, ID number, and age (in frames). Objects in O_t and O_{t+k}, $k > 0$, are corresponded using spatial and temporal coherency.

It should be noted that the tracker is not required to be very accurate, as the self-similarity metric we use is robust and can handle tracking errors of several pixels (as measured in our examples).

Also note that, when the background of a tracked object is sufficiently homogeneous, and the tracked object does not change size significantly during several periods, then accurate object segmentation is not necessary. In these cases, we can allow O_t to include both the foreground and background. Examples of such backgrounds include grassy

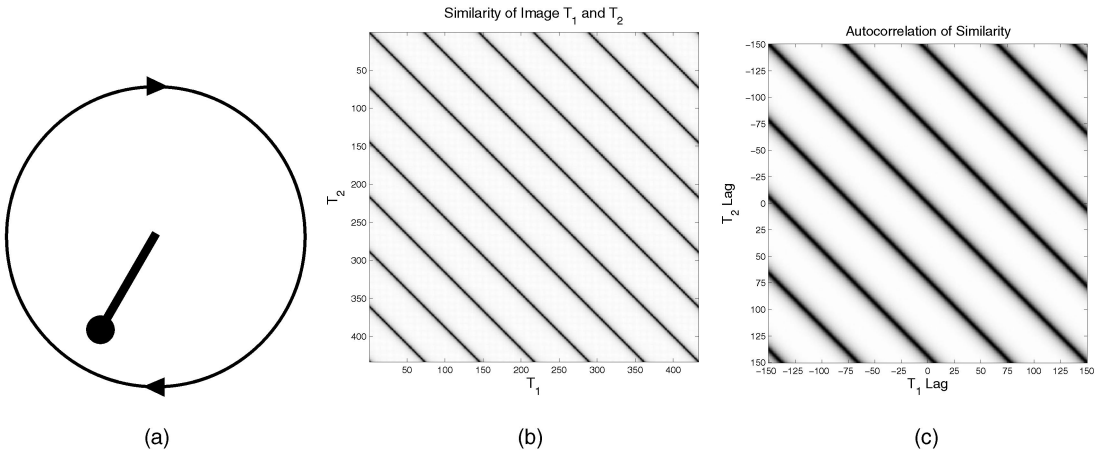

Fig. 3. (a) Pendulum in zero gravity with a constant angular velocity. The arrows denote the direction of motion. (b) Similarity plot for pendulum. Darker pixels are more similar. (c) Autocorrelation of similarity plot.

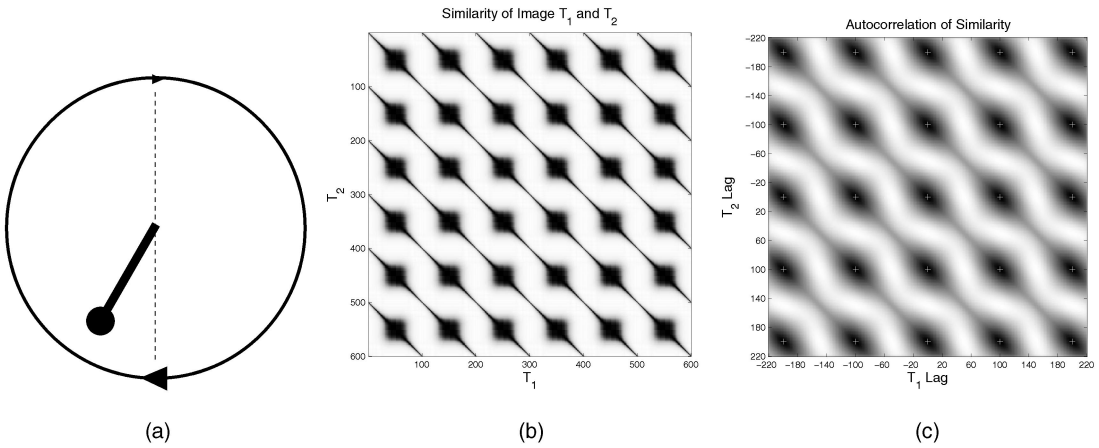

Fig. 4. (a) Pendulum in gravity with single angular direction. The arrows denote the direction and magnitude of motion; the pendulum travels faster at the bottom of its trajectory than at the top. (b) Similarity plot for pendulum. (c) Autocorrelation of similarity plot. The peaks are denoted by "+" symbols.

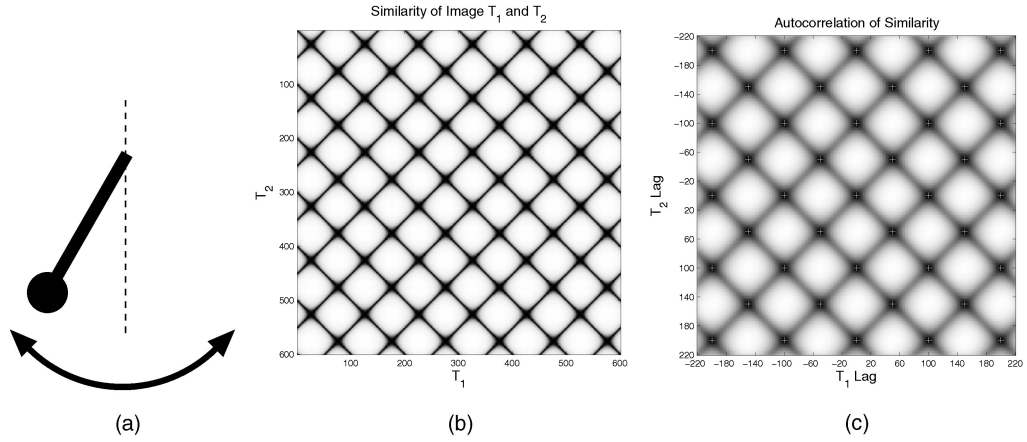

Fig. 5. (a) Pendulum in gravity with an oscillating angular direction. The arrows denote the direction of motion. (b) Similarity plot for pendulum. (c) Autocorrelation of similarity plot. The peaks are denoted by "+" symbols.

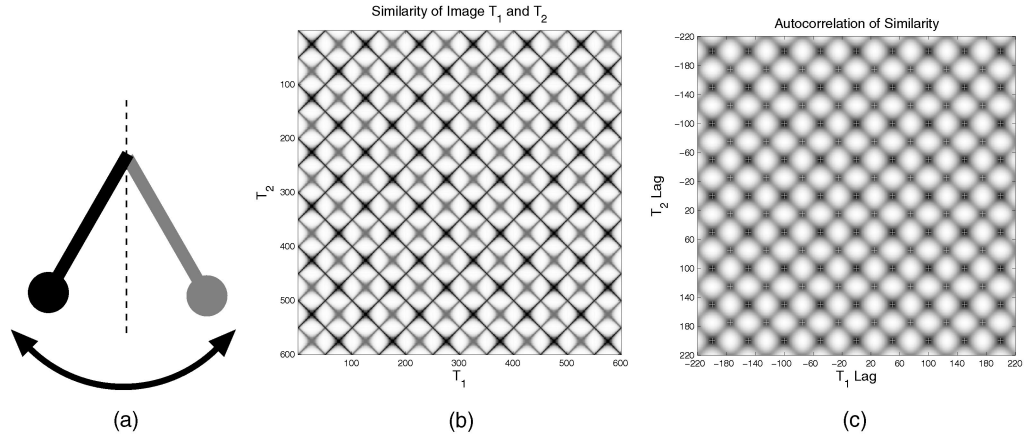

Fig. 6. (a) Two pendulum out of phase 180° in gravity. The arrows denote the direction of motion. (b) Similarity plot for pendulums. (c) Autocorrelation of similarity plot. The peaks are denoted by "+" symbols.

fields, dirt roads, and parking lots. An example of such a sequence is given in Fig. 15.

3.2 Periodicity Detection and Analysis

The output of the motion segmentation and tracking algorithm is a set of foreground objects, each of which has a centroid and size. To detect periodicity for each object, we first align the segmented object (for each frame) using the object's centroid and resize the objects (using a Mitchell filter [32]) so that they all have the same dimensions. The scaling is required to account for an apparent size change due to change in distance from the object to the camera. Because the object segmentation can be noisy, the object dimensions are estimated using the median of N frames (where N is the number of frames we analyze the object over). The object O_t's self-similarity is then computed at times t_1 and t_2. While

Fig. 7. Person walking on a treadmill.

many image similarity metrics can be defined (e.g., normalized cross-correlation, Hausdorff distance [15], color indexing [2]), perhaps the simplest is absolute correlation:

$$St_1, t_2 = \sum_{(x,y)\in B_{t_1}} |O_{t_1}(x,y) - O_{t_2}(x,y)|, \quad (4)$$

where B_{t_1} is the bounding box of object O_{t_1}. In order to account for tracking errors, the minimal S is found by translating over a small search radius r:

$$S't_1, t_2 = \min_{|dx,dy|<r} \sum_{(x,y)\in B_{t_1}} |O_{t_1}(x+dx, y+dy) - O_{t_2}(x,y)|. \quad (5)$$

For periodic motions, S' will also be periodic. For example, Fig. 8a shows a plot of S' for all combinations of t_1 and t_2 for a walking sequence (the similarity values have been linearly scaled to the gray-scale intensity range [0, 255]; dark regions show more similarity). Note that a similarity plot should be symmetric along the main diagonal; however, if substantial image scaling is required, this will not be the case. In addition, there will always be a dark line on the main diagonal (since an object is similar to itself at any given

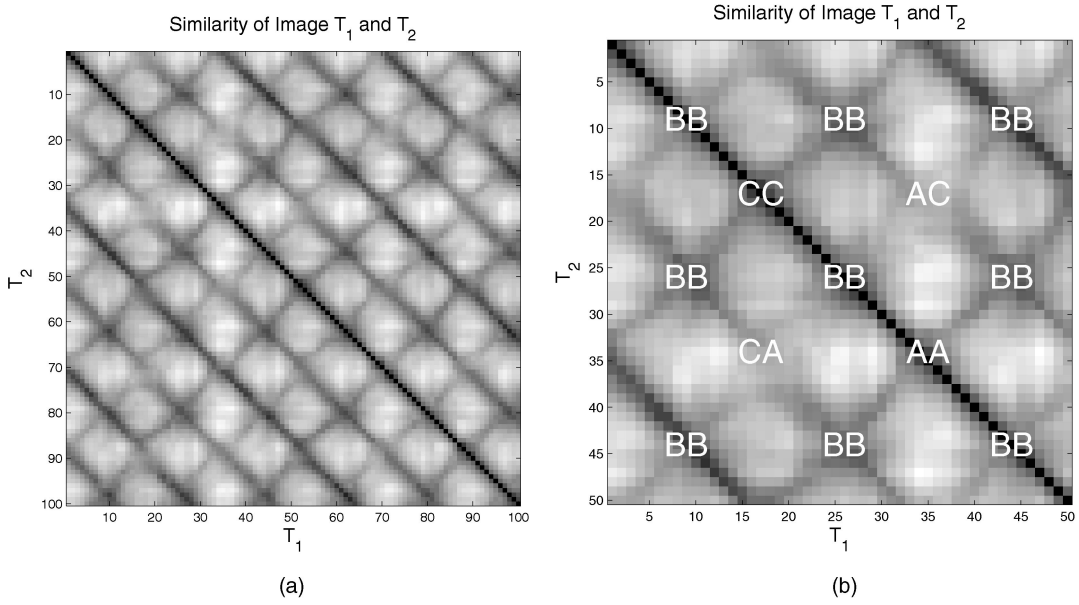

Fig. 8. (a) Similarity plot for the person walking in Fig. 7. (b) Lattice structure for the upper left quadrant of (a). At the intersections of the diagonal and cross diagonal lines, the images are similar to Fig. 7a, 7b, and 7c. This can be used to determine the phase of the walking person.

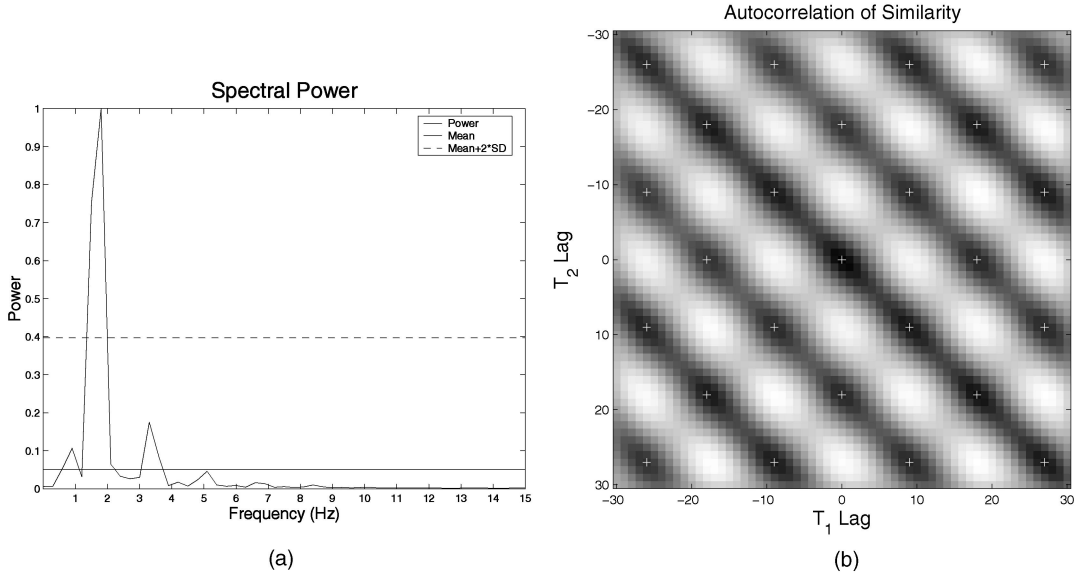

Fig. 9. (a) Power spectrum of similarity of a walking person. (b) Autocorrelation of the similarity of the walking person in Fig. 7 (smoothed with a 5×5 $\sigma = 1$ Gaussian filter). The peaks (shown with white "+" symbols) are used to fit the 45° rotated square lattice in Fig. 2b.

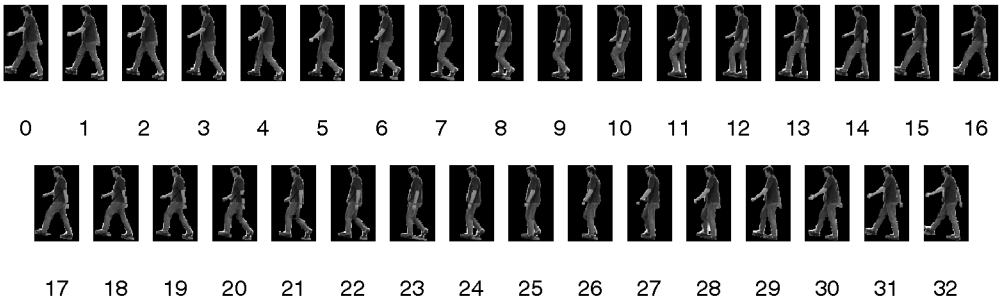

Fig. 10. Cycle of a person walking ($p = 32$). Note the similarity of frame t and $p/2 - t$, and the similarity of frame t and $p/2 + t$.

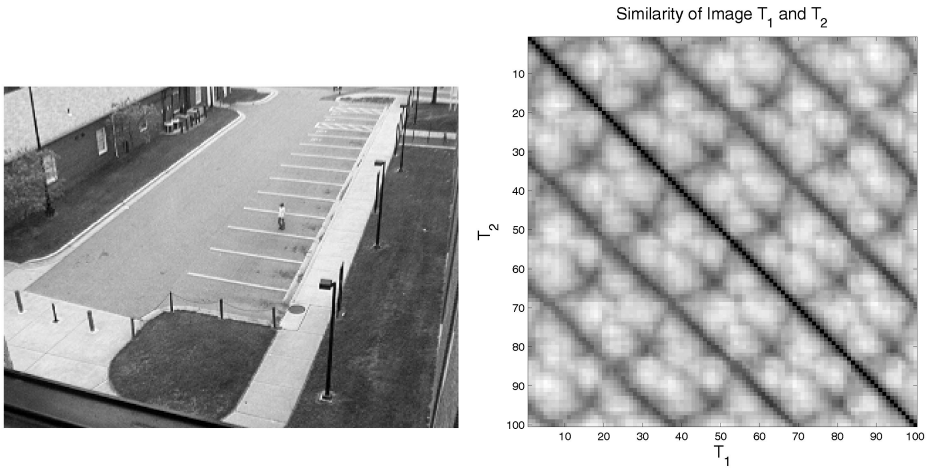

(a) (b)

Fig. 11. (a) First image of a 100-image walking sequence (the subject is walking approximately $25°$ offset from the camera's image plane). (b) Walking sequence similarity plot, which shows the similarity of the object (person) at times t_1 and t_2. Dark regions show greater degrees of similarity.

time) and periodic motions will have dark lines (or curves if the period is not constant) parallel to the diagonal.

To determine if an object exhibits periodicity, we estimate the 1D power spectrum of $S'(t_1, t_2)$ for a fixed t_1 and all values of t_2 (i.e., the columns of S'). In estimating the spectral power, the columns of S' are linearly detrended and a Hanning filter is applied. A more accurate spectrum is estimated by averaging the spectra of multiple t_1s [31] to get a final power estimate $P(f_i)$, where f_i is the frequency. Periodic motion will show up as peaks in this spectrum at the motion's fundamental frequencies. A peak at frequency f_i is significant if

$$P(f_i) > \mu_P + K\sigma_P, \tag{6}$$

where K is a threshold value (typically 3), μ_P is the mean of P, and σ_P is the standard deviation of P. Note that multiple peaks can be significant, as we will see in the examples.

In the above test, we assume that the period is locally constant. The locality is made precise using Time-Frequency analysis given in Section 3.3. We also assume that there are only linear amplitude modulations to the columns of S' (so that linear detrending is sufficient to make the data stationary) and that any additive noise to S' is Gaussian. Both of these assumptions are relaxed in the method given in Section 3.4.

3.2.1 Fisher's Test

If we assume that columns of S' are stationary and contaminated with white noise, and that any periodicity present consists of a single fundamental frequency, then we can apply the well-known Fisher's test [29], [3]. Fisher's test will reject the null hypothesis (that S' is only white noise) if $P(f_i)$ is substantially larger than the average value. Assuming N is even, let $q = \lfloor (N-1)/2 \rfloor$ and

$$\mathcal{E}_q = q \frac{\max_{1 \le i \le q} P(f_i)}{\sum_{i=1}^{q} P(f_i)}. \tag{7}$$

To apply the test, we compute the realized value x of $\mathcal{E}_q$ from S' and then compute the probability:

$$P(\mathcal{E}_q \ge x) = 1 - \sum_{j=0}^{q} (-1)^j \binom{q}{j} (1 - jx/q)_+^{q-1}, \tag{8}$$

where $z_+ = \max(z, 0)$. If this probability is less than α, then we reject the null hypothesis at level α (in practice, we use $\alpha = 0.05$). This test is optimal if there exists a single periodic component at a Fourier frequency f_i in white noise stationary data [29]. To test for periodicities containing multiple frequencies, Seigel's test [29] can be applied.

In practice, Fisher's test, like the K-S test used by [33], works well if the periodic data is stationary with white noise. However, in most of our nonperiodic test data (e.g., Fig. 19a), which is not stationary, both Fisher's and the K-S test yield false periodicities with high confidence.

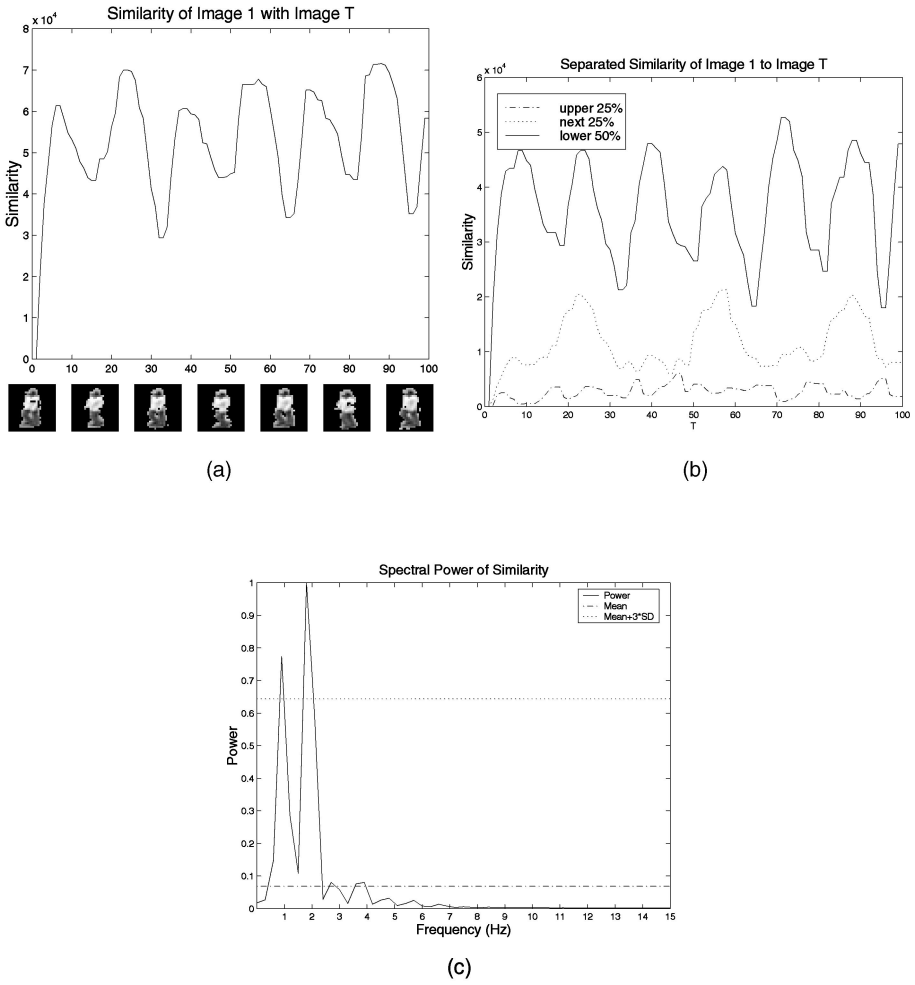

Fig. 12. (a) Column 1 of Fig. 11b, with the corresponding segmented object for the local minima. (b) Image similarity for upper 25 percent, next 25 percent, and lower 50 percent of body. (c) Average power spectrum of all columns of Fig. 11b.

3.2.2 Recurrence Matrices

It is interesting to note that S' is a recurrence matrix [8], [4], without using time-delayed embedded dimensions. Recurrence matrices are a qualitative tool used to perform time series analysis of nonlinear dynamical systems (both periodic and nonperiodic). Recurrence matrices make no assumptions on the stationarity of the data and do not require many data points to be used (a few cycles of periodic data is sufficient). The input for a recurrence matrix is a multidimensional temporally sampled signal. In our use, the input signal is the tracked object image sequence O_t and the distance measure is image similarity. Given a recurrence matrix, the initial trajectory $\vec{X}(t)$ of a point on an object can be recovered up to an isometry [23]. Therefore, the recurrence plot encodes the spatiotemporal dynamics of the moving object. The similarity plot encodes a projection of the spatiotemporal dynamics of the moving object.

3.3 Time-Frequency Analysis

For stationary periodicity (i.e., periodicity with statistics that don't change with time), the above analysis is sufficient. However, for nonstationary periodicity, Fourier analysis is not appropriate. Instead, we use Time-Frequency analysis and the Short-Time Fourier Transform (STFT) [28]:

$$F_x(t, v; h) = \int_{-\infty}^{\infty} x(u) h^*(u - t) e^{i2\pi v u} du, \tag{9}$$

where $h^*(u - t)$ is a short-time analysis window and $x(u)$ is the signal to analyze (S' in our case). The short-time analysis window effectively suppresses the signal $x(u)$ outside a neighborhood around the analysis time point $u = t$. Therefore, the STFT is a "local" spectrum of the signal $x(u)$ around t.

We use a Hanning windowing function as the short-time analysis window. The window length should be chosen to

Fig. 13. Cycle of a running dog ($p = 12$). Note the lack of similarity for any two frames t_1 and t_2, $0 < t_1 < t_2 < p$.

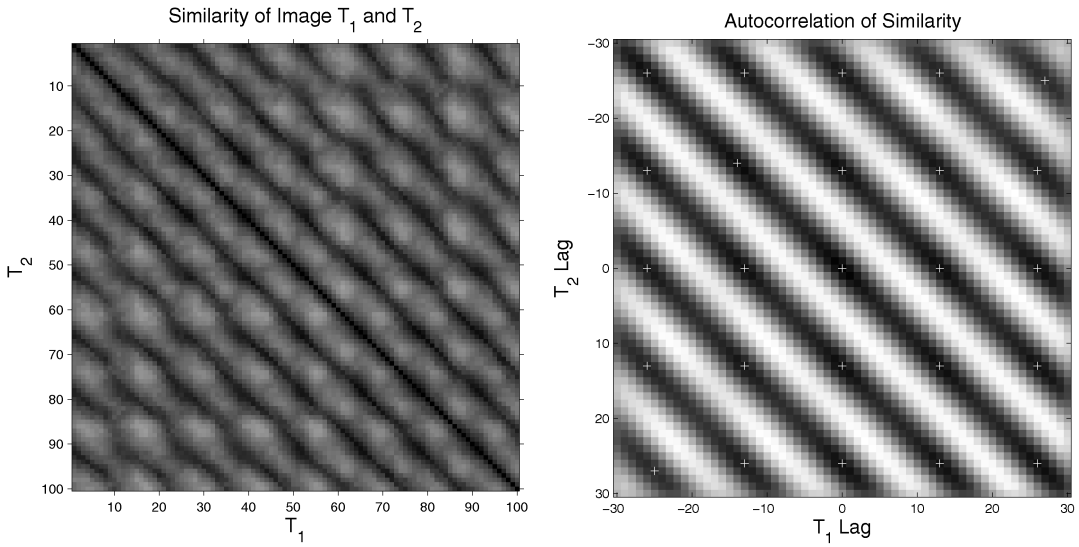

Fig. 14. (a) Similarity plot of the running dog in Fig. 13. Note that there are no dark lines perpendicular to the main diagonal, as shown in Fig. 8a. (b) Autocorrelation of the similarity plot (smoothed with a 5×5 $\sigma = 1$ Gaussian filter). The peaks (shown with white "+" symbols) are used to fit the square lattice in Fig. 2a.

be long enough to achieve a good power spectrum estimate, but short enough to capture a local change in the periodicity. In practice, a window length equal to several periods works well for typical human motions. An example of nonstationary periodicity is given in Section 4.7.

3.4 Robust Periodicity Analysis

In Section 3.2, we used a hypothesis test on the 1D power spectrum of S' to determine if S' contained any periodic motion. The null hypothesis is that there is only white noise in the spectrum, which is rejected by (6) if significant periodic motion is present. However, the null hypothesis can also be rejected if S' contains significant non-Gaussian noise, or if the period is locally nonconstant, or if S' is amplitude modulated nonlinearly. We seek a technique that minimizes the number of false periodicities while maximizing the number of true periodicities. Toward this end, we devise a test that performs well when the assumptions stated in Section 3.2 are satisfied, but does not yield false periodicities when these assumptions are violated.

An alternative technique to Fourier analysis of the 1D columns of S is to analyze the 2D power spectrum of S'. However, as noted in [19], the autocorrelation of S' for regular textures has more prominent peaks than those in the 2D Fourier spectrum. Let A be the normalized autocorrelation of S':

$$A(d_x, d_y) = \frac{\sum_{(x,y)\in R}\left(S'(x,y)-\bar{S'}_R\right)\left(S'(x+d_x,y+d_y)-\bar{S'}_{R_L}\right)}{\left(\sum_{(x,y)\in R}\left(S'(x,y)-\bar{S'}_R\right)^2 \sum_{(x,y)\in R}\left(S'(x+d_x,y+d_y)-\bar{S'}_{R_L}\right)^2\right)^{0.5}}, \quad (10)$$

where $\bar{S'}_R$ is the mean of S' over the region R, $\bar{S'}_{R_L}$ is the mean of S' over the region R shifted by the lag (dx, dy), and the regions R and R_L cover S' and the lagged S'. If S' is periodic, then A will have peaks regularly spaced in a planar lattice M_d, where d is the distance between the lattice points. In our examples, we will consider two lattices, a square lattice $M_{S,d}$ (Fig. 2a), and a 45° rotated square lattice $M_{R,d}$ (Fig. 2b). The peaks $\mathcal{P}$ in A are matched to M_d using the match error measure e:

Fig. 15. Person running across a parking lot, viewed from a moving camera at an altitude of 1,500 feet.

Fig. 16. Zoomed images of the person in Fig. 15, which correspond to the poses in Fig. 7. The person is 12×7 pixels in size.

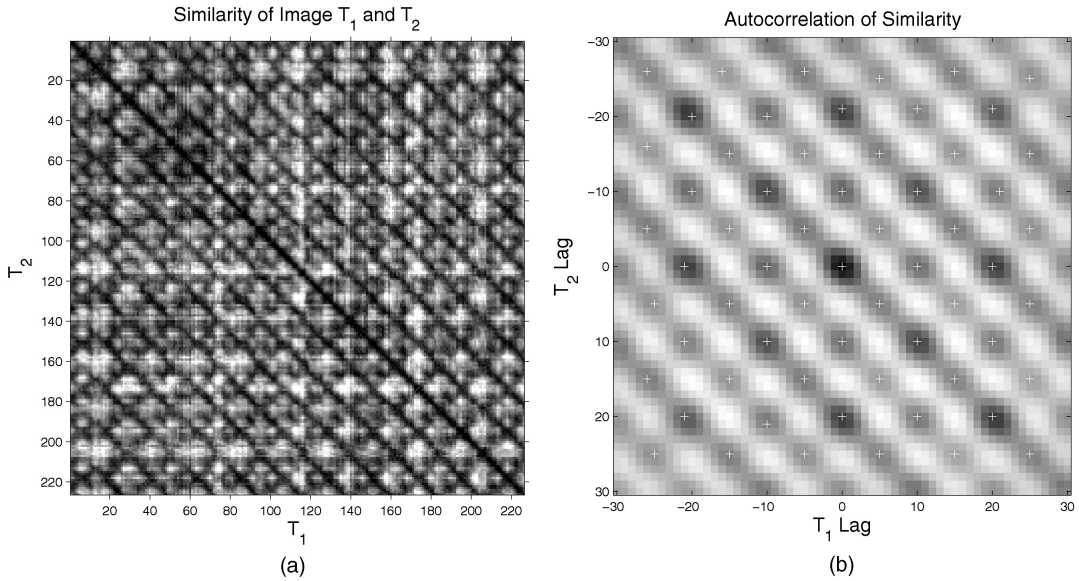

Fig. 17. (a) Similarity plot of the running person in Fig. 15. (b) Autocorrelation of upper quadrant of S'. The peaks are used to fit the 45° rotated square lattice in Fig. 2b.

$$B_i = \left\{ \mathcal{P}_i \,\middle|\, M_{d,i} - \mathcal{P}_i \le \min_{j \neq i} |M_{d,i} - \mathcal{P}_j| \wedge |M_{d,i} - \mathcal{P}_i| \le T_D \wedge A(\mathcal{P}_i) \ge T_A \right\} \quad (11)$$

$$e(M_d) = \sum_i |M_{d,i} - B_i|, \quad (12)$$

where B_i is the closest peak to the lattice point $M_{d,i}$, T_D ($T_D < d/2$) is the maximum distance $\mathcal{P}_i$ can deviate from $M_{d,i}$, and T_A is the minimum autocorrelation value that the matched peak may have. M_d matches $\mathcal{P}$ if all the following are satisfied:

$$\min_{d_1 \le d \le d_2} e(M_d) < T_e, \quad (13)$$

$$|B| \ge T_M, \quad (14)$$

where T_e is a match threshold; $[d_1, d_2]$ is the range of d; T_M is the minimum number of points in M_d to match. In practice, we let $T_D = 1$, $T_e = 2|M_d|$, $T_e = 2|M_d|$, and $T_M = 0.9|M_d|, T_A = 0.25$. The range $[d_1, d_2]$ determines the possible range of the expected period, with the requirement $0 < d_1 < d_2 < L$, where L is the maximum lag used in computing A. The number of points in M_R and M_S can be based on the period of the expected periodicity and frame-rate of the camera. The period $p = 2d\tau$, where τ is the sampling interval (e.g., $\tau = 33$ ms for NTSC video).

Peaks in A are determined by first smoothing A with a Gaussian filter G, yielding A^*. $A^*(i,j)$ is a peak if $A^*(i,j)$ is a strict maximum in a local neighborhood with radius N. In our examples, G is a 5×5 filter with $\sigma = 1$ and $N = 5$. Lin et al. [19] provide an automatic method for determining the optimal size of G.

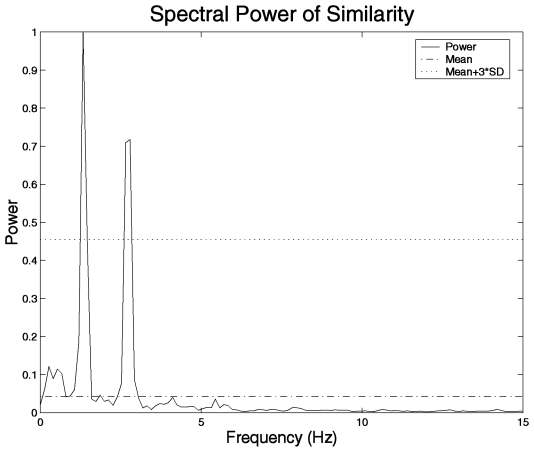

Fig. 18. Spectral power of the running person in Fig. 15.

4 Examples and Applications

4.1 Synthetic Data

In this section, we demonstrate the methods on synthetic data examples. We generated images of a periodic planar pendulum, with different initial conditions, parameters, and configurations. Note that the equation of motion for a simple planar pendulum is

$$\frac{d^2\theta}{dt^2} + \frac{g}{L}\sin\theta = 0, \quad (15)$$

where g is the gravitational acceleration, L is the length of the rigid rod, and θ is the angle between the pendulum rod and vertical axis [22]. In the first example (see Fig. 3a), we set $g = 0$ so that the pendulum has a circular motion with a constant angular velocity. The diagonal lines in the similarity plot (Fig. 3b) are formed due to the self-similarity of the pendulum at every complete cycle. The autocorrelation (Fig. 3c) has no peaks.

In the next example, we use the same configuration, but set $g > 0$ and the initial angular velocity to be sufficient so that the pendulum still has a single angular direction. However, in this configuration, the angular velocity is not

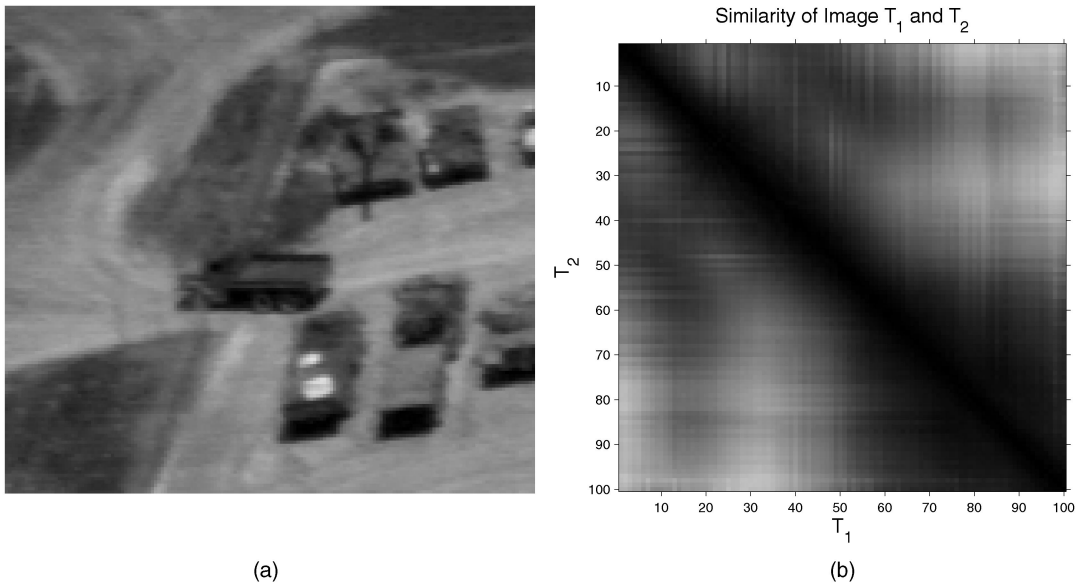

Fig. 19. (a) Vehicle driving across a parking lot. (b) Similarity plot of the vehicle.

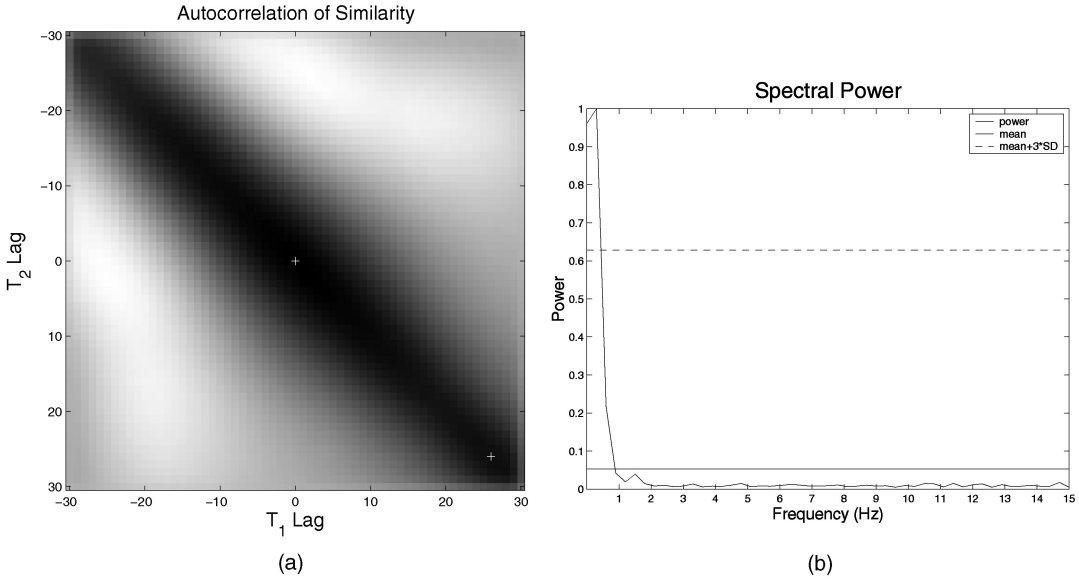

Fig. 20. (a) Spectral power of the vehicle in Fig. 19a. (b) Autocorrelation of S' of the vehicle in Fig. 19a (smoothed with a 5×5 $\sigma = 1$Gaussian filter). The peaks are denoted by "+" symbols.

constant, which is reflected in the qualitatively different similarity plot (Fig. 4b) and autocorrelation (Fig. 4c). Note that the peaks in A match the lattice structure in Fig. 2a.

By decreasing the initial angular velocity, the pendulum will oscillate with a changing angular direction, as shown in Fig. 5a. The similarity plot for this system is shown in Fig. 5b and the autocorrelation in Fig. 5c. Note that the peaks in A match the lattice structure in Fig. 2b.

Finally, for the system of two pendulums 180° out of phase, shown in Fig. 6a, the similarity plot is shown in Fig. 6b and the autocorrelation is shown in Fig. 6c. Note that the peaks in A match the lattice structure in Fig. 2b. Also note the lower measures of similarity for the diagonal lines $S'(t, t + (k + 1/2)p)$ and the cross-diagonal lines $S(t, (k + 1/2)p - t)$, and the corresponding effect on A.

4.2 The Symmetry of a Walking Person

In this example, we first analyze periodic motion with no (little) translational motion, a person walking on a treadmill (Fig. 7). This sequence was captured using a static JVC KY-F55B color camera at 640×480 at 30fps, deinterlaced, and scaled to 160×120. Since the camera is static and there is no

TABLE 1
Confusion Matrix for Person, Dog, and Other Classification

	Other	Person	Dog
Other	30	0	0
Person	0	55	0
Dog	0	0	4

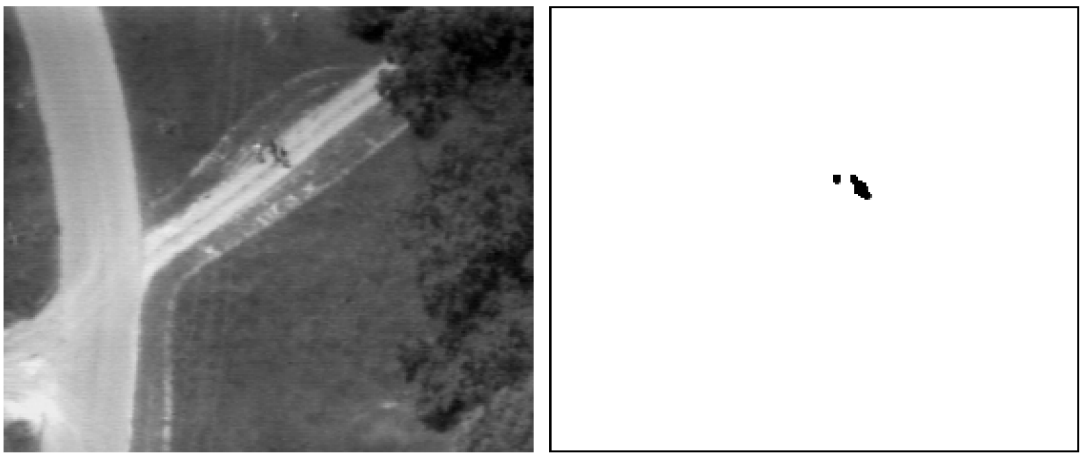

(a) (b)

Fig. 21. (a) Three people running, viewed from a moving camera at an altitude of 1,500 feet. (b) Segmented motion.

translational motion, background subtraction was used to segment the motion [6].

The similarity plot S' for this sequence is shown in Fig. 8a. The dark lines correspond to two images in the sequence that are similar. The darkest line is the main diagonal since $S'(t,t)=0$. The dark lines parallel to the main diagonal are formed since $S'(t, kp/2+t) \simeq 0$, where p is the period, and k is an integer. The dark lines perpendicular to the main diagonal are formed since $S'(t, kp/2-t) \simeq 0$ and is due to the symmetry of human walking (see Fig. 10).

It is interesting to note that, at the intersections of these lines, these images are similar to either Fig. 7a, 7b, or 7c, (see Fig. 8b). That is, S' encodes the phase of the person walking, not just the period. This fact is exploited in the example in Section 4.5.

The autocorrelation A of S' is shown in Fig. 9b. The peaks in A form a 45° rotated square lattice (Fig. 2b), which is used for object classification (Section 4.4). Note that the magnitude of the peaks in A (Fig. 9b) have a pattern similar to the A in Fig. 6c.

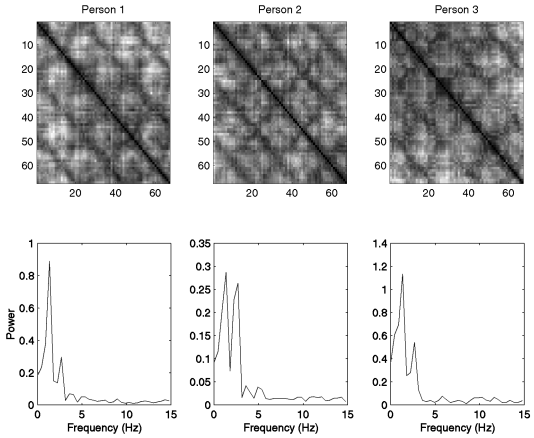

Fig. 22. Similarity plots and spectral power for three people in Fig. 21a. Note that the frequency resolution is not as high as in Fig. 18 since fewer frames are used to estimate the power.

We next analyze the motion of a person who is walking at an approximately 25° offset to the camera's image plane from a static camera. (Fig. 11a). The segmented person is approximately 20 pixels in height, and is shown in Fig. 12a. The similarity plot (Fig. 11a) shows dark diagonal lines at a period of approximately 1 second (32 frames), which correspond to the period of the person's walking. The lighter diagonal lines shown with a period of approximately 0.5 seconds (16 frames) are explained by first noting that the person's right arm swing is not fully visible (due to the 25° offset to the image plane). Therefore, it takes two steps for the body to be maximally self-similar, while the legs become very self-similar at every step. The effect of this is that the similarity measure S' is the composition of two periodic signals, with periods differing by a factor of two. This is shown in Fig. 12b, where the aligned object image is partitioned into three segments (the upper 25 percent, next 25 percent, and lower 50 percent of the body) and S' is computed for each segment. The upper 25 percent, which includes the head and shoulders, shows no periodic motion; the next 25 percent, which includes the one visible arm, has a period double that of the lower 50 percent (which includes the legs). Fig. 12c shows the average power spectrum for all the columns in S'.

4.3 The Symmetry of a Running Dog

In this example, we look at the periodicity of a running dog from a static camera. Fig. 13 shows a complete cycle of a dog (a Black Labrador). Unlike the symmetry of a walking or running person, a running dog has a lack of similarity for $S'(t, kp-t)$. This results in the similarity plot (Fig. 14a) having dark lines parallel to the main diagonal, formed by $S'(t, kp+t)$, but no lines perpendicular to the main diagonal (as with a walking/running person). The similarity plot has peaks (Fig. 14a) that correspond to poses of the dog at frame 0 in Fig. 13. The autocorrelation A of S' is shown in Fig. 14b; the peaks in A form a square lattice (Fig. 2a), which is used in Section 4.4 for object classification.

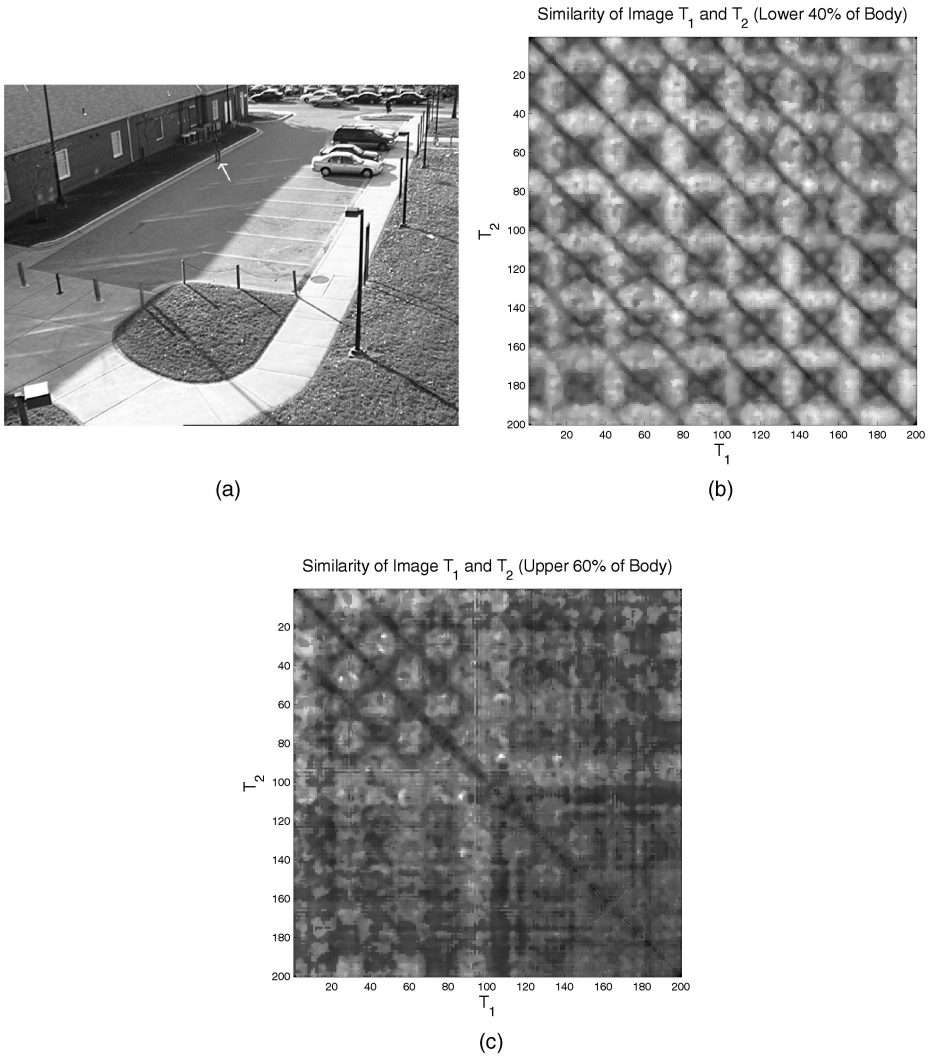

Fig. 23. (a) Frame 100 from a low contrast 200 frame sequence; the subject (marked with a white arrow) puts his hands in his pockets halfway through the sequence. (b) Similarity plot of the lower 40 percent of the body. (c) Similarity plot of the upper 60 percent of the body. The periodicity ceases after the middle of the sequence.

4.4 Object Classification Using Periodicity

A common task in an automated surveillance system is to classify moving objects. In this example, we classify three types of moving objects: people, dogs, and other. We use the lattice fitting method described in Section 3.4 for classification, which is motivated by texture classification methods. Specifically, the square lattice M_S (Fig. 2a) is used to classify running dogs and the 45° square lattice M_R (Fig. 2b) is used to classify walking or running people. Note that $M_{R,d}$ is a subset of $M_{S,d}$, so if both $M_{S,d}$ and $M_{R,d}$ match, M_R is declared the winner. If neither lattice provides a good match to A, then the moving object is classified as other.

The video database used to test the classification consists of video from both airborne surveillance (people and vehicles) and ground surveillance (people, vehicles, and dogs). The database consists of 30 vehicle sequences (25 from airborne video); 55 person sequences (50 from airborne video); and 4 dog sequences (all from ground video).

For the airborne video and dog sequences, the background was not segmented from the foreground object. For these sequences, the background was sufficiently homogeneous (e.g., dirt roads, parking lots, grassy fields) for this method to work. For the other sequences (taken with a static camera), the background was segmented as described in [6].

The airborne video in Fig. 15 was recorded from a Sony XC-999 camera (640×240 at 30fps) at an altitude of about 1,500 feet. There is significant motion blur due to a slow shutter speed and fast camera motion. Additional noise is induced by the analog capture of the video from duplicated SVHS tape. Fig. 15 shows a person running across a parking lot. The person is approximately 12×7 pixels in size (Fig. 16). The similarity plot in Fig. 17a shows a clearly periodic motion, which corresponds to the person running. Fig. 18 shows that the person is running with a frequency of 1.3Hz; the second peak at 2.6Hz is due to the symmetry of the human motion described in Section 4.2. The autocorrelation of S' is shown in Fig. 17b. Fig. 19b shows the similarity plot for the vehicle in Fig. 19a, which has no periodicity. The spectral power for the vehicle (Fig. 20b) is flat. The autocorrelation of S' has only 2 peaks (Fig. 20a).

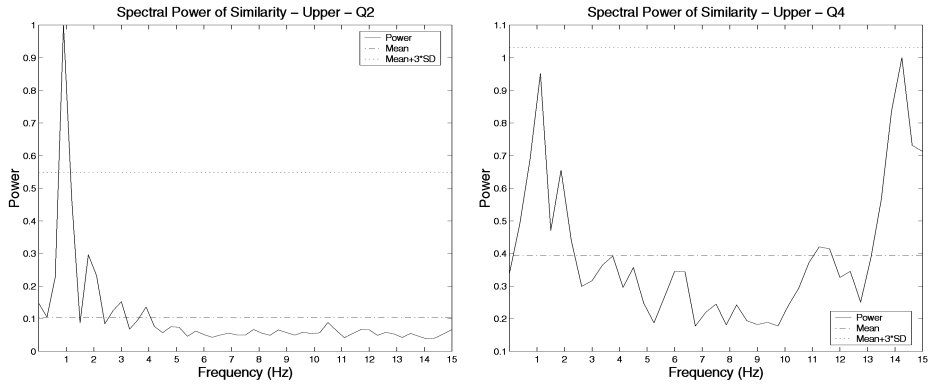

(a) (b)

Fig. 24. (a) Power spectra of the upper left quadrant of Fig. 23c. (b) Power spectra of the lower right quadrant of Fig. 23c.

The results of the classifications are shown in Table 1. The thresholds used for the lattice matching are those given in Section 3.4. Each sequence is 100 images (30 fps); a lag time of 30 images (1 second) is used to compute A.

4.5 Counting People

Another common task in an automated surveillance system is to count the number of people entering and leaving an area. This task is difficult since, when people are close to each other, it is not always simple to distinguish the individuals. For example, Fig. 21a is a frame from an airborne video sequence that shows three people running along a road and the result of the motion segmentation (Fig. 21b). Simple motion blob counting will give an inaccurate estimate of the number of people. However, if we know the approximate location of the airplane (via GPS) and have an approximate site model (a ground plane), we can estimate what the expected image size of an "average" person should be. This size is used to window a region with motion for periodic detection. In this example, three nonoverlapping windows were found to have periodic motion, each corresponding to a person. The similarity plots and spectral powers are shown in Fig. 22.

The similarity plots in Fig. 22 can also be used to extract the phase angle of the running person. The phase angle is encoded in the position of the cross diagonals of S'. In this example, the phase angles are all significantly different from one another, giving further evidence that we have not over counted the number of people.

4.6 Simple Event Detection

In this example, we show how periodicity can be used as input for event detection. Fig. 23a shows a person walking through a low contrast area (in a shadow) toward the camera; halfway through the 200 image sequence the person stops swinging his arms and puts them into his pockets. This action is shown on the similarity plots for the upper and lower portions of the body. Specifically, in Fig. 23c, a periodic pattern for the upper part of the body is visible for the images [1, 100], but not for [101, 200]. This is further shown by the significant peak in the power spectrum for the images [1,100] (Fig. 24a) and the lack of significant peaks in the power spectrum for the images [101, 200] (Fig. 24b). Thus, while the image of the person is only 37 pixels high in this sequence and we are not tracking his body parts, we can deduce that he stopped swinging his arms at about frame 100. An automated surveillance system can use this technique to help decide if someone is carrying an object. In [13], we combine periodicity and shape analysis to detect if someone is carrying an object.

4.7 Nonstationary Periodicity

In this example, a person is walking, and roughly half way through the sequence, starts to run (see Fig. 25a). The similarity plot (Fig. 25b) clearly shows this transition. Using a short-time analysis windowing Hanning function of length 3,300 ms (100 frames), the power is estimated in the walking and running stages (Fig. 26).

4.8 Estimating Human Stride Using Periodicity

In [26] and [20], human gait was used for person recognition. In this example, we do not analyze the gait (which is how people walk or run), but rather estimate the stride length of a walking or running person. The stride itself can be useful for person recognition, particularly during tracking. For example, stride length can help object (person) correspondence after occlusions. Stride length can also be used for input to a surveillance system to detect auto theft in a parking area (e.g., a person of different size and stride length drove off with a car than the person who drove in with the car).

Assume the area of surveillance has a site model, and the camera is calibrated. The estimated stride length is $L = v_g p$, where v_g is the ground velocity of the person, and p is the period. For best results, v_g and p should be filtered to reduce the inherent noise in the tracking and period estimation. For example, in Fig. 25a, the estimated stride of the person is 22 inches when walking and 42 inches when running, which is within 2 inches of the person's actual stride.

5 Real-Time System

A real-time system has been implemented to track and classify objects using periodicity. The system uses a dual Processor 550MHz Pentium III Xeon-based PC and runs at 15Hz with 640×240 gray-scale images captured from an airborne video camera. The system uses the real-time stabilization results from [12].

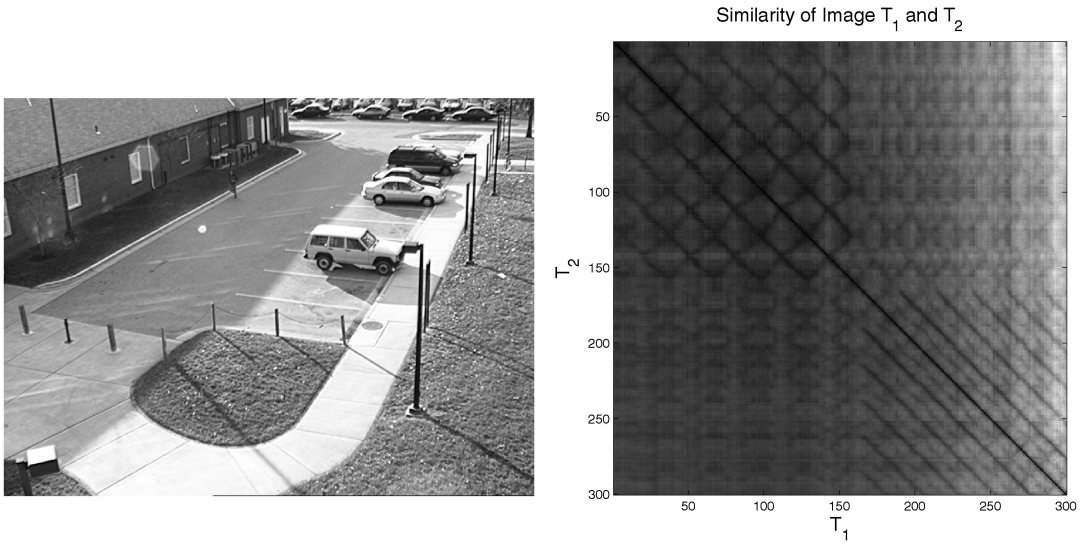

Fig. 25. (a) Person walking, then running. (b) Similarity plot of walking/running sequence.

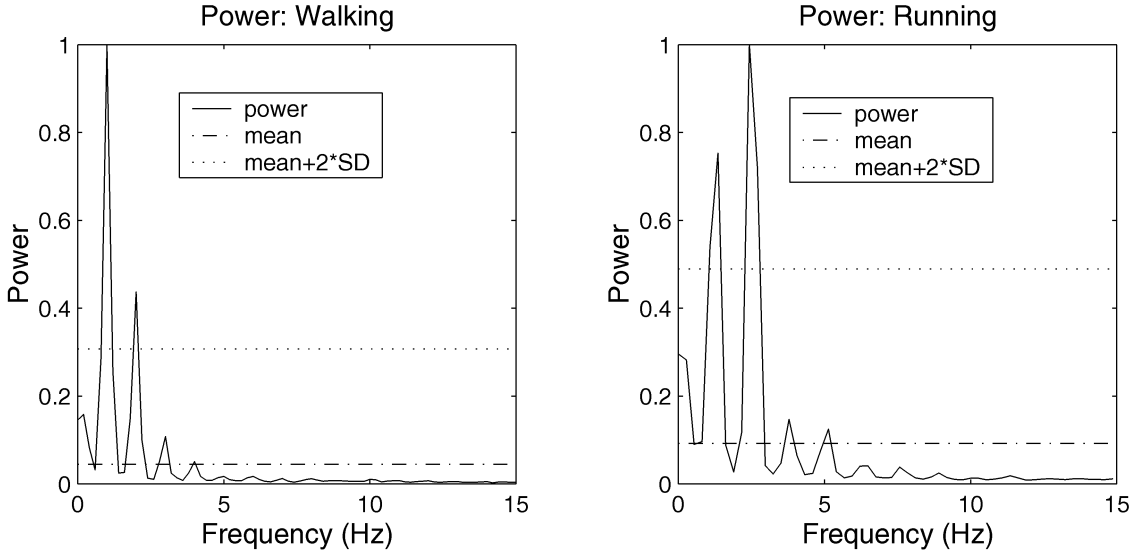

Fig. 26. Spectral power for walking/running sequence in Fig. 25a.

We will briefly discuss how the method can be efficiently implemented to run on a real-time system. In computing S', for each new frame, only a single column that corresponds to the new frame needs to be recomputed; the remaining entries can be reused (shifted) for the updated S'. Therefore, for each new frame, only $O(N)$ $S'(i,j)$ computations need to be done, where N is the number of rows and columns in S'.

For computing A, the 2D FFT can be utilized to greatly decrease the computational cost [18].

Finally, SIMD instructions, such as those available on the Pentium III, can be utilized for computing $S'(i,j)$, as well as A (either directly or using the FFT).

6 Conclusions

We have described new techniques to detect and analyze periodic motion as seen from both a static and moving camera. By tracking objects of interest, we compute an object's self-similarity as it evolves in time. For periodic motion, the self-similarity measure is also periodic and we apply Time-Frequency analysis to detect and characterize the periodic motion. The periodicity is also analyzed robustly using the 2D lattice structures inherent in similarity matrices.

Future work includes using alternative independent motion algorithms for moving camera video, which could make the analysis more robust for nonhomogeneous backgrounds for the case of a moving camera. Further use of the symmetries of motion for use in classification of additional types of periodic motion is also being investigated.

Acknowledgments

The airborne video was provided by the DARPA Airborne Video Surveillance project. This paper was written under the support of contract DAAL-01-97-K-0102 (ARPA Order E653), DAAB07-98-C-J019, and AASERT Grant DAAH-04-96-1-0221.

REFERENCES

[1] M. Allmen, "Image Sequence Description Using Spatiotemporal Flow Curves: Toward Motion-Based Recognition," PhD thesis, Univ. of Wisconsin, Madison, 1991.

[2] D.H. Ballard and M.J. Swain, "Color Indexing," *Int'l J. Computer Vision,* vol. 7, no. 1, 11-32, 1991.

[3] P. Brockwell and R. Davis, *Time Series: Theory and Methods.* Spinger-Verlag, 1987.

[4] M. Casdagli, "Recurrence Plots Revisited," *Physica D,* vol. 108, pp. 12-44, 1997.

[5] C.J. Cohen, L. Conway, and D. Koditschek, "Dynamic System Representation, Generation, and Recognition of Basic Oscillatory Motion Gestures," *Proc. IEEE Int'l Conf. Automatic Face and Gesture Recognition,* 1996.

[6] R. Cutler and L. Davis, "View-Based Detection and Analysis of Periodic Motion," *Proc. Int'l Conf. Pattern Recognition,* p. SA14, Aug. 1998.

[7] R. Cutler and L. Davis, "Real-Time Periodic Motion Detection, Analysis, and Applications," *Proc. Computer Vision and Pattern Recognition,* pp. 326-332, June 1999.

[8] P. Eckmann, S.O. Kamphorst, and D. Ruelle, "Recurrence Plots of Dynamical Systems," *J. Europhysics Letters,* vol. 4, pp. 973-977, 1987.

[9] J. Emmerton, "The Pigeon's Discrimination of Movement Patterns (Lissajous Figures) and Contour-Dependent Rotational Invariance," *Perception,* vol. 15, no. 5, pp. 573-588, 1986.

[10] H. Fujiyoshi and A. Lipton, "Real-Time Human Motion Analysis by Image Skeletonization," *Proc. IEEE Workshop Applications of Computer Vision,* p. session 1A, Oct. 1998.

[11] N. Goddard, "The Interpretation of Visual Motion: Recognizing Moving Light Displays," *Proc. IEEE Workshop Motion,* pp. 212-220, Mar. 1989.

[12] M. Hansen, P. Anandan, K. Dana, G. van der Wal, and P. Burt, "Real-Time Scene Stabilization and Mosaic Construction," *Proc. DARPA Image Understanding Workshop,* pp. 457-465, Nov. 1994.

[13] I. Haritaoglu, R. Cutler, D. Harwood, and L. Davis, "Backpack: Detection of People Carrying Objects Using Silhouettes," *Proc. Int'l Conf. Computer Vision,* pp. 102-107, 1999.

[14] B. Heisele and C. Wohler, "Motion-Based Recognition of Pedestrians," *Proc. Int'l Conf. Pattern Recognition,* Aug. 1998.

[15] D. Huttenlocher, G.A. Klanderman, and W. Rucklidge, "Comparing Images Using the Hausdorff Distance," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 15, no. 9, pp. 805-863, Sept. 1993.

[16] G. Johansson, "Visual Motion Perception," *Scientific Am.,* vol. 232, pp. 75-88, June 1976.

[17] R. Kern, "Visual Position Stabilization in the Hummingbird Hawk Moth, Macroglossum Stellatarum L. I. Behavioural Analysis," *J. Computer Physiology A,* vol. 182, no. 2, pp. 225-237, 1998.

[18] J.P. Lewis, "Fast Normalized Cross-Correlation," *Vision Interface,* 1995.

[19] H.-C. Lin, L.-L. Wang, and S.-N. Yang, "Extracting Periodicity of a Regular Texture Based on Autocorrelation Functions," *Pattern Recognition Letters,* vol. 18, pp. 433-443, 1997.

[20] J. Little and J. Boyd, "Recognizing People by Their Gate: The Shape of Motion," *Videre,* vol. 1, no. 2, 1998.

[21] F. Liu and R. Picard, "Finding Periodicity in Space and Time," *Proc. Int'l Conf. Computer Vision,* pp. 376-383, Jan. 1998.

[22] J. Marion, *Classical Dynamics of Particles and Systems.* Academic Press, 1970.

[23] G. McGuire, N.B. Azar, and M. Shelhamer, "Recurrence Matrices and the Preservation of Dynamical Properties," *Physics Letters A,* vol. 237, pp. 43-47, 1997.

[24] T.A. McMahon, *Muscles, Reflexes, and Locomotion.* Princeton Univ. Press, 1984.

[25] D. McReynolds and D. Lowe, "Rigidity Checking of 3D Point Correspondences Under Perspective Projection," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 18, no. 12, pp. 1,174-1,185, 1996.

[26] S. Niyogi and E. Adelson, "Analyzing and Recognizing Walking Figures in XYT," *Proc. Int'l Conf. Computer Vision and Pattern Recognition,* pp. 469-474, Dec. 1994.

[27] S. Niyogi and E. Adelson, "Analyzing Gait with Spatiotemporal Surfaces," *Proc. IEEE Workshop Motion of Non-Rigid and Articulated Objects,* pp. 64-69, 1994.

[28] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing.* Prentice Hall, 1989.

[29] D.B. Percival, "Harmonic Analysis," *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques.* Cambridge Univ. Press, 1993.

[30] R. Polana and R. Nelson, "Detection and Recognition of Periodic, Non-Rigid Motion," *Int'l J. Computer Vision,* vol. 23, no. 3, pp. 261-282, June/July 1997.

[31] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C.* Cambridge Univ. Press, 1988.

[32] D. Schumacher, "General Filtered Image Rescaling," *Graphics Gems III,* D. Kirk, ed. Harcourt Brace Jovanovich, 1992.

[33] S.M. Seitz and C.R. Dyer, "View-Invariant Analysis of Cyclic Motion," *Int'l J. Computer Vision,* vol. 25, no. 3, pp. 1-23, 1997.

[34] A. Selinger and L. Wixson, "Classifying Moving Objects as Rigid or Non-Rigid without Correspondences," *Proc. DARPA Image Understanding Workshop,* pp. 341-347, Nov. 1998.

[35] P. Tsai, M. Shah, K. Keiter, and T. Kasparis, "Cyclic Motion Detection for Motion Based Recognition," *Pattern Recognition,* vol. 27, no. 12, pp. 1,591-1,603, 1994.

[36] H. Weyl, *Symmetry.* Princeton Univ. Press, 1952.

Ross Cutler received BS degrees in mathematics, computer science, and physics in 1992. The MS degree in computer science in 1996 and the PhD degree in computer science in 2000 from the University of Maryland, College Park. His research interests include motion-based recognition, motion segmentation, HCI, video indexing, multimedia databases, gesture recognition, augmented reality, and real-time systems. He is currently a faculty research assistant in the University of Maryland Institute of Advanced Computer Studies. He has previously been employed at Johns Hopkins University (research staff in neurology/neurosurgery), and has consulted for Microsoft Research, Emory University, the University of Pennsylvania, and Scientech Inc.

Larry S. Davis received the BA degree from Colgate University in 1970, the MS, and PhD degrees in computer science from the University of Maryland in 1974 and 1976, respectively. From 1977-1981, he was an assistant professor in the Department of Computer Science at the University of Texas, Austin. He returned to the University of Maryland as an associate professor in 1981. From 1985-1994, he was the director of the University of Maryland Institute for Advanced Computer Studies. He is currently a professor in the Institute and the Computer Science Department, as well as acting chair of the Computer Science Department. He was named a fellow of the IEEE in 1997. Dr. Davis is known for his research in computer vision and high performance computing. He has published more than 75 papers in journals and has supervised more than 12 PhD students. He is an associate editor of the *International Journal of Computer Vision* and an area editor for *Computer Models for Image Processor: Image Understanding.* He has served as program or general chair for most of the field's major conferences and workshops, including the Fifth International Conference on Computer Vision, the field's leading international conference.