In *Motion-Based Recognition*, M. Shah and R. Jain, eds., Kluwer, Boston, 1997.

CYCLIC MOTION ANALYSIS USING THE PERIOD TRACE

STEVEN M. SEITZ AND CHARLES R. DYER
Department of Computer Sciences
University of Wisconsin
Madison, WI 53706

1. Introduction

Non-rigid motion analysis is complicated by a lack of general purpose rules or constraints governing how an object or scene evolves over time. In order to design practical algorithms, much of the work to date has focused on object-model-based techniques, such as interpretation of facial expressions, detection of human locomotion, cardiac image analysis, and gesture recognition. In contrast, repeating motions all share common temporal features that can be formally described and interpreted regardless of the particular object or scene that is moving. For instance, it is not necessary to recognize the runner or the motion in Fig. 1 in order to determine his stride frequency. Because many real-world motions repeat, e.g., a heart beating, an athlete running, and a wheel rotating, cyclic motion analysis techniques have broad applicability.

For motions that repeat regularly, detecting the period can yield important information about underlying object or scene properties. For instance, heart-rate relates to activity levels, the period of a hand on a clock to units of time, and a runner's stride frequency to velocity. For repeating motions whose cycle length can change as a function of time, which we call *cyclic* motions, detecting irregularities can illuminate physically important changes in the scene. By analyzing the variance of the period in several cycles of a runner's gait or a swimmer's stroke, anomalies can be located, providing feedback on specific areas which may need improvement [1]. Medical imagery can be analyzed using the same techniques, for instance to detect unevenness in a heart's beating motion.

Analyzing 3D cyclic motions from image sequences is challenging for several reasons, including

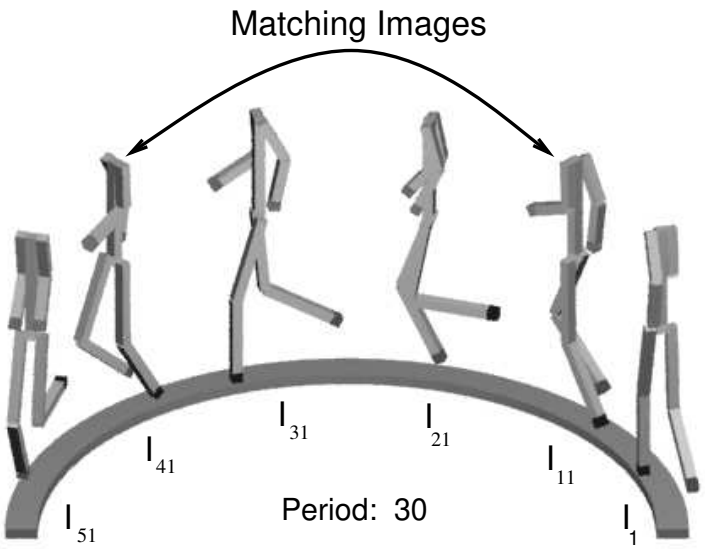

Figure 1. A runner's stride frequency can be determined by detecting *matching* images of the runner at the same point in his cycle. The time between repetitions is the period. Matching images, i.e., 2D views of the same 3D configuration, can be found by tracking the trajectories of five or more image features.

- Irregularities: The motion may not be strictly periodic due to fluctuations in speed from one cycle to the next. These fluctuations must be accounted for in order to robustly detect that a motion is cyclic and to interpret its properties.
- Complexity: The observed motion can be non-rigid and arbitrarily complex. Moreover, motion information must be integrated over long sequences of images.
- Camera motion: Changes in viewpoint affect the apparent motion and therefore complicate the analysis. In order to determine salient features of the underlying motion in the scene, viewpoint-dependent effects must be filtered out.

This chapter addresses the first two difficulties by introducing a compact, strictly temporal motion representation called the *period trace* for the detection and interpretation of cyclic motions. By accounting for changes in speed, uneven cyclic motions can be detected and irregularities located. Computational complexity is minimized by using simple and efficient image comparison operations in the recovery procedure. Using the period trace, an analog of the period is defined for cyclic motions, which we call the *instantaneous period*, giving cycle length as a function of time. Information about motion irregularities can be determined by considering derivatives of

the instantaneous period. The issue of camera motion is addressed by incorporation of a novel view-invariant match function for comparing images. The match function incorporates principles of affine invariance to allow for arbitrary changes in viewpoint.

The remainder of the chapter is structured as follows. Section 2 discusses related work on periodic motion analysis. Section 3 describes the *period trace* and its image-based detection using view-invariant techniques. Section 4 introduces several motion features that can be derived from the period trace. Section 5 presents a method for automatically recovering a motion's period trace using an optimization approach. Section 6 describes experiments on real and synthetic image sequences.

2. Related Work

Several researchers have investigated ways of measuring periodicity information from image sequences [2, 3, 4]. Allmen and Dyer [2] described an approach for detecting motion periods under orthographic projection. They used the curvature scale space of point trajectories to detect repeating patterns of curvature maxima and hence infer a period. Polana and Nelson [3] presented a method for detecting periodic motions using Fourier transforms of several point trajectories. In theory, the period of the motion could be detected as well by averaging the fundamental frequencies of the point trajectories, although the authors indicated that determining the period in this way was unreliable. Tsai and Shah [4] described a similar technique, using Fourier transforms of curvature values, where the period was determined from a single point trajectory.

None of these approaches are appropriate for analyzing repeating motions that lack a constant period; in particular, they cannot be used to locate motion irregularities or changes in speed. Furthermore, they rely on detecting periodicity in 2D and fail to account for changes in camera viewpoint and reference frame. For example, the runner's motion in Fig. 1 is periodic in a 3D reference frame that moves with the runner. However, during the course of a cycle, the image motion of points on the runner will not sweep out periodic paths, as a result of the runner's constantly changing attitude relative to the camera.

Motion information has proven useful in a number of related recognition problems. Johansson's pioneering work on moving light displays (MLD's) [5] demonstrated that human motions such as walking and running can be recognized solely based on the trajectories of a small number of bright spots attached to different parts of the body. Since then, motion-based recognition has become an active area of study within computer vision. These efforts include identification of pedestrians [6, 7], MLD-based motion recognition

[8], hand gesture recognition [9, 10], interpretation of facial expressions [11, 12], and temporal textures [13]. Details of many of these efforts are described in other chapters of this book, and in a recent survey article [14].

Cyclic motion analysis is unique among motion-based recognition approaches in that it does not require any type of object- or motion-specific model. Other motion-based recognition techniques [6, 7, 9, 10, 11, 12, 13] require *a priori* models of the underlying object and/or the motion in the scene, although these models can potentially be learned [9, 15]. In contrast, periodicity is a *universal* motion characteristic that can be detected and described without knowledge of the underlying object and without reference to a previous instance of the motion. Therefore, the techniques in this chapter can be used to interpret image sequences of completely unfamiliar objects and motions.

In [16] we introduced a match function that permits view-invariant comparison of images and removes the restriction of a stationary observer. In [17] we described a technique for analyzing irregular cyclic motions using the *period trace*. Here we present a unified framework for the analysis of periodic and cyclic motions from image sequences based on the period trace and view-invariant image matching.

3. Cyclic Motion and the Period Trace

Motion can be described both in spatial and temporal terms, as changes in position or variations in speed. Most existing representations use some combination of spatial and temporal information to describe how an object evolves over time. In many applications, however, we are interested in the temporal evolution of an object or a scene and *not* its instantaneous shape. Consequently, in this section we introduce an alternative, purely temporal approach that is independent of an object's spatial properties. Our temporal motion representation, called the *period trace*, describes how the period changes throughout the course of a motion and contains a wealth of useful motion information. Because no assumptions are made about the spatial structure or representation of the scene, the period trace can be used to analyze both rigidly and non-rigidly moving objects and scenes. To motivate the period trace, we first define what it means for a motion to be *periodic* and, more generally, *cyclic*.

3.1. PERIODIC AND CYCLIC MOTIONS

We define a *motion* $M(t)$ to be a function whose value at time t is the instantaneous configuration of a continuously-moving object or scene[1]. We

[1] Although time is implicit, t could represent another quantity.

call a motion P *periodic* if it repeats with period p, i.e.,

$$P(t+p) = P(t) \tag{1}$$

for some constant $p > 0$ and all times t in a given time domain. The smallest such constant p is the *period* and the set $P_{t_0} = \{P(t) \mid t_0 \leq t < t_0 + p\}$ is called the *cycle* beginning at time t_0.

The notion of periodicity given in Eq. (1) has a very restrictive temporal constraint, namely, that the motion is perfectly regular from one cycle to the next. This constraint is relaxed with the introduction of a *period-warping* function, ϕ, as follows:

Definition 1 *A motion C is called* **cyclic** *if*

$$C(\phi(t)) = C(t)$$

for all times t in a given time domain and some increasing continuous function ϕ satisfying $\phi(t) > t$. A function ϕ satisfying these properties is called C-warping.

Intuitively, $\phi(t)$ corresponds to the start of the next cycle after the cycle beginning at time t. The increasing condition on ϕ ensures that a cyclic motion is order preserving, i.e., $t_1 < t_2$ implies $\phi(t_1) < \phi(t_2)$. Notice that all periodic motions are cyclic and that a cyclic motion is periodic when Definition 1 is satisfied for $\phi(t) = t + p$, with p the period. Although not all cyclic motions are periodic, any cyclic motion can be *warped* into a periodic motion by an appropriate temporal transformation (see Section 4.4).

3.2. PROJECTIONS OF 3D CYCLIC MOTIONS

Since we are concerned with image-based analysis of 3D cyclic motions, the effects of camera projection must be considered. For instance, consider the problem of measuring the stride frequency of an athlete filmed while running along a path. Both the relative movement of the runner's limbs and his change in attitude relative to the camera contribute to the motion in the image sequence, but only the former is relevant to the stride frequency. We would like to be able to determine the stride frequency in a way that is not affected by changes in viewpoint. This is accomplished using a view-invariant match function that equates a set of images if and only if they represent views of an object in the same 3D configuration.

3.2.1. *Projected Match Criterion*

Formally, we want to determine if a set of images *match*, i.e., correspond to views of the same 3D configuration of an object or a scene. Our formulation assumes an *affine* camera model [18, 19] which is a generalization of

orthographic, weak perspective, and paraperspective projections. An image sequence $\mathbf{I}_1, \ldots, \mathbf{I}_m$ is represented as a sequence of $2 \times n$ matrices where each column is the instantaneous position of one of n image features. It is assumed that corresponding columns of $\mathbf{I}_t$ and $\mathbf{I}_s$ represent projections of the same scene point. Hence, the correspondence between features in different images is assumed to be known. An object or scene, $\mathbf{S}$, is represented as a $3 \times n$ matrix with columns corresponding to the 3D coordinates of features in an object-centered affine reference frame. Without loss of generality, assume that each row of $\mathbf{I}_t$ has zero mean; choose the origin of the image coordinates to be the centroid of the feature points. Under an affine projection model a set of images, $\Gamma = \{\mathbf{I}_{t_1}, \ldots, \mathbf{I}_{t_k}\}$, match if and only if

$$\mathbf{I}_{t_i} = \mathbf{\Pi}_i \mathbf{S} \tag{2}$$

for some fixed $3 \times n$ matrix $\mathbf{S}$, $i = 1, \ldots, k$, and 2×3 matrices $\mathbf{\Pi}_1, \ldots, \mathbf{\Pi}_k$.

Define the measurement matrix $\mathbf{M}_\Gamma$ of a set of images to be the concatenation of the image measurements:

$$\mathbf{M}_\Gamma = \begin{bmatrix} \mathbf{I}_{t_1} \\ \vdots \\ \mathbf{I}_{t_k} \end{bmatrix}$$

If the images Γ match, then, by Eq. (2), we can express $\mathbf{M}_\Gamma$ as

$$\mathbf{M}_\Gamma = \begin{bmatrix} \mathbf{\Pi}_1 \\ \vdots \\ \mathbf{\Pi}_k \end{bmatrix} \mathbf{S} \tag{3}$$

Therefore $\mathbf{M}_\Gamma$ is the product of two matrices, each having rank at most 3. It follows that $\mathbf{M}_\Gamma$ has rank 3 or less. This is the Rank Theorem, due to Tomasi and Kanade [20], and generalized for an affine camera. Conversely, any measurement matrix of rank 3 or less can be decomposed as in Eq. (3) using singular value decomposition [20]. Therefore, we have the following:

Generalized Rank Theorem: *A set of images* Γ *match if and only if* $\mathbf{M}_\Gamma$ *is of rank at most 3.*

Under orthographic projection (the case considered by Tomasi and Kanade), the rank condition alone is not sufficient to determine that a set of images match. However, it is both necessary and sufficient under the more general affine camera model. Hence, the Generalized Rank Theorem can be used to verify if a set of images could have been produced by a 3D cyclic motion. Accordingly, let $\mathbf{I}_1, \ldots, \mathbf{I}_m$ be a sequence of images, let ϕ be an integer-valued function, and denote $\Gamma_t^\phi = \{\mathbf{I}_{\phi^i(t)} \mid 1 \leq \phi^i(t) \leq m\}$. The following

result is a consequence of the Generalized Rank Theorem and Eq. (1):

Projected Match Criterion: *An image sequence,* $\mathbf{I}_1, \ldots, \mathbf{I}_m$*, is the affine projection of a 3D cyclic motion if and only if there exists an increasing function satisfying* $\phi(t) > t$ *such that* $rank(\mathbf{M}_{\Gamma_t^\phi}) \leq 3$ *for* $t = 1, \ldots, m$.

Such a function ϕ is called **I**-*warping*. The Projected Match Criterion suggests a provably-correct way of detecting cyclic motions: check all possible **I**-warping functions. Under an assumption of periodicity, this brute-force strategy is feasible [16] but exhaustive search is too costly in general. A more efficient method is described in Section 5.

3.2.2. *View-Invariant Image Matching*

Because rank measurements are highly sensitive to numerical errors, a more robust measure of match quality is needed. We can characterize the residual error of a sequence of images by the amount by which we have to perturb features in order to make the images match. Define the match error of a set of images as follows:

$$dist_{\mathcal{A}}(\Gamma) = min\ \{\|\mathbf{E}\|_{rms} \mid rank(\mathbf{M}_\Gamma + \mathbf{E}) \leq 3\}$$

$\|\mathbf{E}\|_{rms}$ is the root-mean-squared norm of the matrix $\mathbf{E}$ defined by $\|\mathbf{E}\|_{rms} = \sqrt{\frac{1}{2kn}\sum_{i,j}\mathbf{E}_{ij}^2}$. Alternatively, $dist_{\mathcal{A}}$ can be expressed in terms of the singular values of $\mathbf{M}_\Gamma$:

$$dist_{\mathcal{A}}(\Gamma) = \sqrt{\frac{1}{2kn}\sum_{i=4}^{n}\sigma_i^2} \tag{4}$$

where $\sigma_1 \ldots \sigma_n$ are the singular values of Γ. To derive Eq. (4), observe that singular value decomposition gives

$$\mathbf{M}_\Gamma = \mathbf{U}\mathbf{\Sigma}\mathbf{V}$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices and the singular values of $\mathbf{M}_\Gamma$ appear along the diagonal of $\mathbf{\Sigma}$ (a $2k \times n$ diagonal matrix). The above equation can be rewritten:

$$\mathbf{M}_\Gamma = \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V} + \mathbf{U}\mathbf{\Sigma}'\mathbf{V} \tag{5}$$

where $\tilde{\mathbf{\Sigma}}$ and $\mathbf{\Sigma}'$ are structured as

$$\tilde{\mathbf{\Sigma}} = \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \mathbf{0} \\ & & \sigma_3 & \\ & \mathbf{0} & & \mathbf{0} \end{bmatrix} \qquad \mathbf{\Sigma}' = \begin{bmatrix} \mathbf{0} & & \mathbf{0} & \\ & \sigma_4 & & \\ \mathbf{0} & & \ddots & \\ & & & \sigma_n \end{bmatrix}$$

Eq. (5) can be expressed as

$$\mathbf{M}_\Gamma = \tilde{\mathbf{M}}_\Gamma + \mathbf{M}'_\Gamma$$

where $\tilde{\mathbf{M}}_\Gamma = \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}$ and $\mathbf{M}'_\Gamma = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}$.

$\tilde{\mathbf{M}}_\Gamma$ is the optimal (in an RMS sense) rank-3 approximation of $\mathbf{M}_\Gamma$ [21]. Hence, $\mathbf{M}'_\Gamma$ is the minimal perturbation of $\mathbf{M}_\Gamma$ that produces a match. Therefore, $dist_{\mathcal{A}}(\Gamma) = \|\mathbf{M}'_\Gamma\|_{rms}$. The latter term is just $\|\mathbf{\Sigma}'\|_{rms}$ since $\mathbf{U}$ and $\mathbf{V}$ are orthogonal, and the result follows. □

The measure $dist_{\mathcal{A}}$ gives the average amount (in pixels) necessary to additively perturb the coordinates of each image feature in order to produce a set of matching images. In the case where there are only two images we abbreviate $dist_{\mathcal{A}}(\{\mathbf{I}_s, \mathbf{I}_t\})$ as $d_{\mathcal{A}}(\mathbf{I}_s, \mathbf{I}_t)$. $dist_{\mathcal{A}}$ has the following properties:

- $dist_{\mathcal{A}}(\Gamma) = 0$ if and only if the images Γ match exactly.
- $dist_{\mathcal{A}}$ is well-behaved with respect to noise since it is defined in terms of feature measurement perturbations. See, for example, [21].
- $dist_{\mathcal{A}}$ is defined in image coordinates and can be directly related to measurement errors.
- $dist_{\mathcal{A}}$ is always zero when less than five features are considered.
- $dist_{\mathcal{A}}$ may be non-zero when the features are co-planar. Therefore co-planarity of the feature set is not a limitation, in contrast to related techniques [18, 19, 20].
- For n features, $d_{\mathcal{A}}(\mathbf{X}, \mathbf{Y}) = \frac{\sigma_4}{2\sqrt{n}}$ and the evaluation cost is $O(n)$ arithmetic operations. For m images, the evaluation cost is the smaller of $O(nm^2)$ and $O(mn^2)$.

3.2.3. *Occlusions and Lost Features*

Matching several images at once is appropriate when a large number of features are continuously visible throughout an image sequence. For real world applications however, this is seldom the case due to factors such as noise and occlusion. Therefore, in practice we compute only pairwise correlations between images that have at least 5 features in common. In terms of the Projected Match Criterion, this amounts to labeling a motion as cyclic if $d(\mathbf{I}_t, \mathbf{I}_{\phi(t)}) = 0$ for all image pairs $\mathbf{I}_t$ and $\mathbf{I}_{\phi(t)}$ having at least 5

features in common. Note that this condition is weaker than the Projected Match Criterion; it is only a necessary condition for a motion to be cyclic.

3.3. THE PERIOD TRACE

A C-warping function ϕ contains information concerning the temporal variation of a cyclic motion C. Unfortunately, Definition 1 does not determine a unique C-warping ϕ for a given cyclic motion C. For instance, if C is periodic with period p, $\phi(t) = t + kp$ for any positive integer k satisfies Definition 1. For definiteness we introduce the notion of *instantaneous period*:

Definition 2 *Let C be a cyclic motion. Let*

$$\boldsymbol{\phi}_1 = \text{pointwise-infimum } \{\phi \ C\text{-warping}\}$$

Define $\tau_1(t) = \phi_1(t) - t$. τ_1 is called the **instantaneous period** *of C.*

The fact that ϕ_1 is C-warping follows from the continuity of C: Let $\mathcal{A} = \{\phi \ C\text{-warping}\}$ and let $\{\psi_i\}_{i=1}^{\infty}$ be a subset of $\mathcal{A}$ converging to ϕ_1 at t. Then $C(\psi_i(t)) = C(t)$ for each i, so $\lim_{i\to\infty} C(\psi_i(t)) = C(t)$. But $\lim_{i\to\infty} C(\psi_i(t)) = C(\phi_1(t))$ by continuity of C. Hence $C(\phi_1(t)) = C(t)$ for all values of t, i.e., ϕ_1 is C-warping.

Intuitively, $\tau_1(t)$ is the length of the cycle beginning at time t. For instance, if C is periodic, τ_1 is the period. From τ_1, several useful quantities are computable, including τ_n, the instantaneous combined length of the next n cycles. Accordingly, let

$$\phi_n(t) \ = \ (\overbrace{\phi \circ \phi \circ \ldots \circ \phi}^{n})(t)$$

for integers $n > 0$. Then the *nth instantaneous period*, τ_n, is defined as

$$\tau_n(t) = \phi_n(t) - t \tag{6}$$

In addition to being continuous, τ_n has the following property:

$$\frac{\tau_n(t+h) - \tau_n(t)}{h} > -1 \tag{7}$$

which follows from Eq. (6) and the increasing condition on ϕ_1. If τ_n is differentiable, Eq. (7) is equivalent to the condition $\tau_n' > -1$.

The inverse functions exist and are defined as

$$\begin{aligned} \phi_{-n} &= (\phi_n)^{-1} \\ \tau_{-n}(t) &= \phi_{-n}(t) - t \end{aligned}$$

for $n > 0$. $\tau_{-n}(t)$ has the intuitive interpretation as the combined length of the previous n cycles ending at time t. Its value is always negative.

We refer to the set of functions $\{\tau_n \mid n \neq 0\}$ as the *period trace* of a cyclic motion C. For instance, the period trace of a periodic motion is a set of constant functions: $\tau_n = np$, where p is the period, as shown in Fig. 2. The period trace is a comprehensive map of the cycles in a repeating motion and identifies all corresponding configurations in different cycles. These properties make the period trace a useful tool for comparing different cycles and tracking changes in a repeating motion over time.

4. Trends and Irregularities

The period trace of a cyclic motion describes how the motion varies from one cycle to the next. Consequently, many important attributes of a motion can be computed directly from the period trace, without reference to the spatial structure of the underlying scene. In this section we present several features that quantify temporal trends and irregularities in cyclic motions.

4.1. LOCAL FEATURES

Because the period trace reflects the cumulative motion of n cycles, it is not obvious that local temporal information can be derived from it. In this section we show that local temporal irregularities can indeed be computed from derivatives of the instantaneous periods. Any interval $[t, t+h]$ in which τ_n is not constant, for some n, is said to be an *irregular interval*. The quantity $\tau_n(t+h) - \tau_n(t)$ gives the cumulative change in the nth period, and $\frac{\tau_n(t+h)-\tau_n(t)}{h}$ is the mean rate of change of τ_n in the interval $[t, t+h]$. Letting $h \to 0$, the mean rate of change converges to the instantaneous rate of change, $\tau_n'(t)$. Henceforth, we denote τ_n' as the one-sided derivative defined by $\tau_n'(t) = \lim_{h \to 0} \frac{\tau_n(t+h)-\tau_n(t)}{h}$.

Values of t for which $\tau_n'(t) \neq 0$ are *irregular points*, i.e., points where the motion is faster or slower at $\phi_n(t)$ than at t. Moreover, τ_n' provides both the sign and magnitude of the irregularity. For instance, $\tau_1'(t) = 1$ indicates that the period at time t is increasing at a rate of 1 unit, i.e., motion is faster at t than at $\phi_1(t)$ by 1 unit. Because the motion is faster at the beginning of the cycle than at the end, it must slow down during the course of the cycle, thereby causing a net increase in the period.

Generally, not all points t where $\tau_n'(t) \neq 0$, for some $n \neq 0$, correspond to times where the motion is irregular relative to the norm. A nonzero value of $\tau_n'(t)$ only indicates that motion is irregular at time t *relative* to time $\phi_n(t)$. A point t where the motion is globally irregular satisfies $\tau_n'(t) \neq 0$, for all $n \neq 0$ (see Fig. 2(f)).

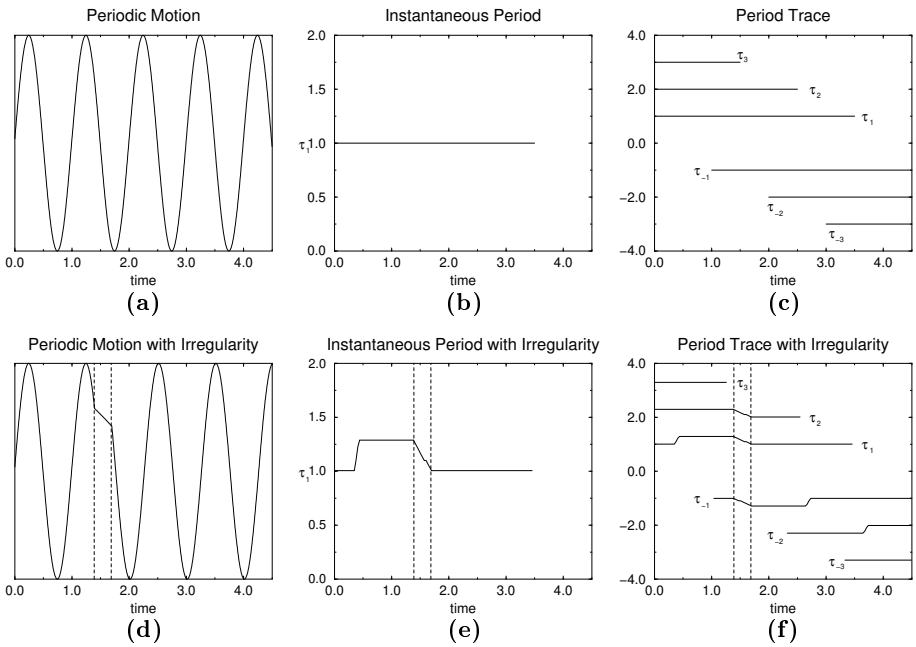

Figure 2. Effects of an irregularity on the instantaneous period. A periodic sinusoid signal (a) has a constant period (b) and a period trace consisting of multiples of the period (c). A small irregularity is introduced (d) by temporarily slowing by a factor of 10. The irregularity shows up in two places in the instantaneous period (e), where a ramp occurs once when the irregularity enters the current cycle and once where it leaves. The actual location of the irregularity is detected using the period trace (f) where the ramps "line-up" only at the true irregularity (i.e., in a small interval around time 1.5).

The total first-order irregularity of a cyclic motion C at a point in time t can be defined as

$$irreg_C(t) = mean\ \{|\tau_n'(t)| \mid n \neq 0\} \tag{8}$$

Higher-order irregularities may also be relevant in certain situations. For instance, consider a jogger whose stride frequency is steadily decreasing. Due to the changes in speed, τ_n' will be nonzero and second-order irregularities may be more interesting, i.e., where $\tau_n''(t) \neq 0$. Unless otherwise qualified, the term *irregularity* henceforth refers to a first-order irregularity.

4.2. GLOBAL FEATURES

The period trace can also be used to compute various global features of a cyclic motion. These include

– period

- cyclic acceleration
- shortest and longest cycle
- regular and irregular points, cycles, and intervals

The simplest of these is the period. If $\tau_1'(t)$ is uniformly 0, then τ_1 is constant and corresponds to the period. Similarly, if $\tau_1' = c$, where $c > 0$, then the period is uniformly increasing at a rate of c. In the case of a decelerating jogger, τ_1' gives the change in speed in terms of the rate of increase in stride period. The mean value of τ_1' gives the mean rate of increase of the period, which we call the average *cyclic acceleration*.

In addition to trends such as the period and cyclic acceleration, the period trace can be used to compute globally significant points, cycles, or intervals in a motion. For instance, a longest (shortest) cycle is given by any interval $[t, \phi_1(t))$ for which $\tau_1(t)$ is maximal (minimal).

Highly irregular points can be found by maximizing Eq. (8). Similarly, the irregularity of a cycle is measured by integrating Eq. (8) over the entire cycle:

$$irreg(C_t) = \frac{1}{\tau_1(t)} \int_t^{\phi_1(t)} irreg_C(s)\ ds$$

Another useful feature is the median[2] cumulative irregularity of an interval:

$$irreg_C(t_1, t_2) = median\ \{|\tau_n(t_2) - \tau_n(t_1)| \ \mid\ n \neq 0\}$$

4.3. MOTION SIGNATURES

Somes types of cyclic motion produce telltale temporal signatures. Examples include periodic motions and motions with uniform cyclic acceleration, both of which have instantaneous periods that are linear. The form of a motion's period trace may also indicate something about the distribution of irregularities. Cyclic motions that contain isolated irregularities have very distinct signatures (see Figs. 2 and 7); the instantaneous period of a motion of this sort is piecewise-constant except at isolated intervals whose locations and extents correspond to the irregularities.

The locations of irregularities in the period trace provide other important qualitative clues about the behavior of a moving object. Any cyclic motion containing at least one irregularity has a period trace with ramps in various places, e.g., see Fig. 2(f). There are two categories of ramps; those that line-up, recurring at the same place in each cycle, and isolated ramps that don't recur. Generally, irregular motions will have both types of ramps,

[2]We have found the *median* to be more robust than the *mean* in this context.

but certain motions have only ramps that line-up. These are motions where irregularities occur at roughly the same point in each cycle. For instance, a heart may contract at a different rate in each cycle, resulting in a series of ramps that line-up at each contraction.

4.4. MEASURING RELATIVE SPEED AND REMOVING IRREGULARITIES

The period trace provides detailed information about cycles and changes in speed but does not explicitly give instantaneous speed as a function of time. Consider the motion of a jogger whose stride frequency varies through time. Suppose we wish to describe these variations not as changes in stride frequency, but in terms of differences in speed relative to the norm. For instance, we might say that the jogger is currently running twice as fast as normal (e.g., 15 mph versus 7.5 mph). This can be achieved by recovering a *time-warping* function that maps a perfectly regular, e.g. periodic, motion into the uneven motion of the runner.

Any cyclic motion can be warped into a periodic motion by *flattening* its period trace. This flattening transformation and its inverse are the time-warping functions we're after. The first step is to find a function that maps a cyclic motion into a periodic motion. In other words, given a cyclic motion, $C(t)$, we wish to find a reparameterization $s = \sigma(t)$ such that $\tilde{\tau}_1(s)$ is constant, where

$$\tilde{\tau}_1(s) = (\sigma \circ \phi_1 \circ \sigma^{-1})(s) - s$$

The result is that $C(s)$ is periodic, satisfying Eq. (1).

For any cyclic motion $C(t)$, there exist innumerable reparameterizations, $s = \sigma(t)$, such that $C(s)$ is periodic. A natural choice is one that synchronizes the motion with respect to a particular cycle. The result is that the reference cycle remains intact and all other cycles are aligned to match the reference cycle. This approach is attractive because a cycle may be chosen that satisfies a particular criterion and the whole motion will be rectified to conform to the criterion. For instance, irregularities may be removed by choosing a *regular* reference cycle, using the techniques of Section 4.2. Similarly, choosing a cycle that has a particular anomaly will result in a motion with a corresponding anomaly in every cycle.

Let $C(t)$ be a cyclic motion and choose a cycle, C_{t_0}, starting at time t_0. Define time warping function σ and its inverse as

$$\begin{array}{lcll} \sigma(t) & = & \phi_{-n}(t) + n\tau_1(t_0) & \text{for } t \in [\phi_n(t_0), \phi_{n+1}(t_0)) \\ \sigma^{-1}(s) & = & \phi_n(s - n\tau_1(t_0)) & \text{for } \sigma^{-1}(s) \in [\phi_n(t_0), \phi_{n+1}(t_0)) \end{array}$$

Claim: *$C(s)$ is periodic, where $s = \sigma(t)$.*

Proof: Fix a value of s. Suppose, without loss of generality, that $\sigma^{-1}(s) \in [\phi_n(t_0), \phi_{n+1}(t_0))$. Then

$$\begin{aligned}
\tilde{\tau}_1(s) &= (\sigma \circ \phi_1 \circ \phi_n)(s - n\tau_1(t_0)) - s \\
&= (\phi_{-n-1} \circ \phi_1 \circ \phi_n)(s - n\tau_1(t_0)) + (n+1)\tau_1(t_0) - s \\
&= s - n\tau_1(t_0) + (n+1)\tau_1(t_0) - s \\
&= \tau_1(t_0)
\end{aligned}$$

Therefore $\tilde{\tau}_1(s)$ is constant, so $C(s)$ is periodic with period $\tau_1(t_0)$. □

An image sequence is rectified using σ^{-1} as follows: given an initial image sequence $I_1, I_2, \ldots, I_m$, the rectified sequence is $I_{\sigma^{-1}(1)}, I_{\sigma^{-1}(2)}, \ldots, I_{\sigma^{-1}(k)}$, where $k = \lfloor \sigma(m) \rfloor$. The effect is that any irregularities are removed and there is no variation in speed from one cycle to the next.

Whereas σ^{-1} can be used for making a sequence periodic, σ is useful for describing how a cyclic motion deviates from periodic. The derivative of σ gives the factor by which motion is locally faster with respect to the reference cycle. Contrast this with τ_1', which gives the change in cycle length as a function of time; $\frac{d\sigma}{dt}$ gives relative speed directly, but requires a reference cycle. For instance, $\frac{d\sigma}{dt}(t_1) = 2$ indicates that the motion at time t_1 is twice as fast as normal. In Section 6 these techniques are used to measure relative speed and to demonstrate the rectification process for a real image sequence.

5. Recovering the Period Trace

In order to detect cyclic motion features such as those introduced in Section 4, we must be able to recover functions $\tau_n(t)$ for each integer $n \neq 0$ from a cyclic motion C. Because the complete period trace can be computed from τ_1, it is sufficient to determine τ_1. For periodic motions, τ_1 is constant and can be determined using unconstrained scalar optimization methods. More generally, the recovery of τ_1 is posed as a constrained functional minimization problem.

5.1. THE PERIODIC CASE

If a cyclic motion is known to have a certain form, this information may be used to simplify the recovery of its period trace. Particularly, if an *a priori* motion model is available, e.g., linear or piecewise constant, the model can ameliorate the task of determining τ_1 by reducing the search space and making the recovery procedure less sensitive to noise. Existing techniques for recovering periodicity information [2, 3, 4, 16]. implicitly use this strategy, exploiting the assumption of a constant period. The case where such a model is not available is treated in the next section.

The instantaneous periods of a periodic motion are constant and correspond to the multiples of the period. Therefore, the period trace of a periodic motion can be found by optimizing some function over the set of candidate periods. Towards this end, we introduce a tool for assessing the significance of a candidate period or period trace.

5.1.1. *A Significance Measure*

In the presence of noise and quantization effects it is unlikely that corresponding images of a 3D cyclic motion will match precisely, as dictated by the Projected Match Criterion. Still, any candidate period $\tilde{p}$ induces a correspondence between images spaced $\tilde{p}$ frames apart. In order to discriminate between cyclic and non-cyclic motions, it is necessary to evaluate the quality of this correspondence, based on how well the images match. For this purpose we use a statistical measure based on the Kolmogorov-Smirnov (K-S) test [22].

The quality of a computed period trace can be assessed by considering the problem within the context of sampling theory. The operation of comparing two images $\mathbf{I}_s$ and $\mathbf{I}_t$ can be thought of as taking a random point sample from *match space*: $\mathcal{M} = \{d_{\mathcal{A}}(\mathbf{I}_s, \mathbf{I}_t) \mid 1 \leq s < t \leq m\}$. A period trace $\pi = \{\tau_n \mid n \neq 0\}$ determines a specific sample $\mathcal{S}_\pi \subset \mathcal{M}$ given by

$$\mathcal{S}_\pi = \{d_{\mathcal{A}}(\mathbf{I}_t, \mathbf{I}_{t+\tau_n(t)}) \mid n > 0\}$$

For example, a period of p corresponds to the sample:

$$\mathcal{S}_p = \{d_{\mathcal{A}}(\mathbf{I}_t, \mathbf{I}_{t+np}) \mid n > 0\}$$

$\mathcal{S}_\pi$ contains the match scores of precisely those images that correspond, under the hypothesis that scene motion is cyclic with period trace π. If this hypothesis is correct, the distribution of $\mathcal{S}_\pi$ should differ significantly from that of $\mathcal{M}$, since matching images minimize $d_{\mathcal{A}}$. The K-S test gives the probability, or *P-value*, of a random sample of $\mathcal{M}$ matching the cumulative histogram of $\mathcal{S}_\pi$, i.e., the probability that the motion is not cyclic. For details on how to calculate the K-S statistic, consult [22].

The P-value of a candidate period trace π is evaluated by applying the K-S test to $\mathcal{S}_\pi$ and $\mathcal{M}$. Since the objective is to find a period trace that minimizes $d_{\mathcal{A}}$ it is only necessary to assess P-values of period traces for which

$$mean(\mathcal{S}_\pi) < mean(\mathcal{M}) \tag{9}$$

Candidate period traces for which this inequality is violated are normalized by assigning P-values of 1. The P-value of a motion is defined to be the minimum P-value of all possible period traces.

In practice, a motion's exact P-value is never computed; the P-value of the period trace obtained by the optimization procedure provides an upper bound that is generally sufficient to differentiate cyclic from non-cyclic motions. We have found this test to be quite robust; real cyclic motions yield P-values that are exceedingly small, on the order of 10^{-20} or smaller, whereas non-cyclic motions have P-values at or near 1. In fact, every one of the non-cyclic motions we evaluated had a P-value of 1. The reason is that image sequences are decidedly non-random; adjacent images (i.e., $\mathbf{I}_t$ and $\mathbf{I}_{t+1}$) tend to be similar, resulting in a natural bias towards small periods. $\mathcal{M}$ contains match scores for *all* adjacent frames and it is therefore unlikely that Eq. (9) will be satisfied for any period trace π chosen at random. As a result, most period traces will have P-values of 1. Note that this bias helps to differentiate cyclic from non-cyclic motions since the latter tend to have P-values of 1. For cyclic images sequences there is an implicit assumption that the instantaneous period is uniformly greater than 1. In other words, the frame rate (images per second) should be greater than the frequency (cycles per second), or else the period trace cannot be detected.

5.1.2. *Detecting Periods Using the K-S Test*

Using the K-S test it is straightforward both to calculate the probability that a motion is periodic and to compute the most likely period. Given an image sequence $\mathbf{I}_1, \ldots, \mathbf{I}_m$, evaluate the K-S test for every candidate period $\tilde{p} \leq (m-1)/2$. The value of $\tilde{p}$ with smallest P-value constitutes the most likely period and the corresponding P-value gives an upper bound on the probability that the motion is not cyclic.

A potential complication arises in discriminating between the true period and its multiples. If p is the true period, note that $S_{kp} \subset S_p$ for any positive integer multiple k of p. Hence, it is possible that $mean(S_{kp}) \leq mean(S_p)$, for some value of k. However, the K-S test assigns greater significance to larger samples, adding a built-in bias in favor of smaller periods. Because S_{kp} represents a smaller sample than S_p, the latter is automatically assigned more significance, resulting in a smaller P-value. Consequently, the P-value of the true period will tend to be more significant than those of its multiples due to the natural bias of the K-S test. This phenomenon can be seen in Fig. 5, where the true period has a much smaller P-value than its multiple.

5.2. THE CYCLIC CASE

When nothing is known *a priori* about a cyclic motion, the recovery of its period trace can be posed as a constrained functional optimization problem. Because the complete period trace can be computed from τ_1, it is sufficient

to determine $\tau_1(t)$ for each value of t. Given an image sequence $\mathbf{I}_1, \ldots, \mathbf{I}_m$, we seek a τ_1 that minimizes the energy function E, comprised of two terms:

$$\begin{aligned} E_1(\tau_1, t) &= mean\ \{d_A(\mathbf{I}_t, \mathbf{I}_{t+\tau_1(t)}) \mid 1 \leq t,\ t + \tau_1(t) \leq m\} \\ E_2(\tau_1, t) &= |\tau_1'(t)| \\ E(\tau_1) &= \int_{t_1}^{t_2} E_1(\tau_1, t) + \alpha E_2(\tau_1, t)\ dt \end{aligned} \tag{10}$$

The first term seeks to maximize correlations between corresponding images, whereas the second term is added as a continuity constraint on τ_1. α is a scalar weighting factor that balances E_1 and E_2. τ_1 is also subject to the following hard constraints, as imposed in Section 3.3:

$$\tau_1(t) > 0 \tag{11}$$

$$\tau_1(t+1) - \tau_1(t) \geq -1 \tag{12}$$

Note that the strict inequality in Eq. (7) is relaxed above since it cannot be enforced on the discrete domain $t \in [1, m]$.

We used a multiscale *snake* algorithm adapted from [23, 24] for recovering the period trace by iteratively minimizing the energy function E. The snake is initialized to a rough estimate of τ_1 and is then incrementally refined so as to reach a state of locally minimal energy. At each iteration a value of t is selected and $\tau_1(t)$ is adjusted by 0, 1, or -1 unit so as to globally decrease $E(\tau_1)$ while checking that Eqs. (11) and (12) are satisfied. A pass consists of iterating over all successive values of $t = 1, 2, \ldots, m$. The algorithm performs repeated passes until E converges. Convergence is guaranteed since energy decreases monotonically in each iteration. To avoid problems with local minima, the optimization approach is repeated at three scales of increasing resolution using the output of each stage as the initialization for the next.

Initialization is performed by setting τ_1 to the constant function that minimizes the K-S test, as described in Section 5.1. In combination with the multiscale approach, we found this simple strategy to provide adequate initialization for the optimization procedure to converge correctly for the sequences that we tested. However, this strategy is designed for near-periodic sequences; highly irregular cyclic sequences may require a more sophisticated initialization procedure, perhaps incorporating user-interaction, as in [23].

5.3. THE PERIOD TRACE RECOVERY ALGORITHM

The following sequence of operations is used to compute the period trace and its P-value from a sequence of images $\mathbf{I}_1, \ldots, \mathbf{I}_m$:

1. Compute P-values for each candidate period $\bar{p} = 2, \ldots, \frac{m-1}{2}$. The candidate p with smallest P-value is the most likely period. The associated P-value gives the probability that the motion is not periodic.
2. Initialize the snake to $\tau_1 = p$ at the coarsest scale and apply energy minimization passes until the snake converges. Upon convergence, the new snake configuration is used to initialize energy minimization at the next finer scale. The procedure repeats until convergence is reached at the finest scale.
3. Obtain the P-value of the computed period trace using the K-S test, giving the probability that the motion is not cyclic.

As noted in Section 5.1.1, this algorithm cannot detect periods of one frame or less. Such motions will generally be found to be *non-cyclic* with a P-value of 1.

6. Experiments

One synthetic and two real image sequences are used in this section to show the performance of the optimization methods described in Section 5 and the detection of temporal motion features. Selected images from the sequences are shown in Figs. 3 and 4. Additional experiments with cardiac X-ray images are presented in [17, 25].

6.1. SYNTHETIC JOGGER

In order to study the effects of noise and perspective distortion on the algorithm, we created a simulation of a human jogging around a short track using a volumetric model and real motion data (see Fig. 4). To model the human torso, we used nine parallelepipeds connected by revolute joints. The model was roughly 250 units tall and the track had a diameter of 200 units. The camera was fixed at an angle of 30 degrees from horizontal. The periodic motion consisted of a sequence of joint angles extrapolated from real motion data of a human running, provided by N. Goddard [8]. The projected positions of vertices of the parallelepipeds were used as input to the algorithm. No attempt was made to distinguish visible from occluded vertices.

We found the algorithm of Section 5.1 to be highly robust to both feature noise and perspective effects (see Fig. 5). At the highest noise level of 10 pixels (the projected jogger was roughly 25 units wide and 250 tall), the true period was correctly detected with a P-value of 10^{-41}. At the highest level of perspective distortion (shown in Fig. 4), the detected period had a P-value of 10^{-47}. In all 123 trials the period with lowest P-value matched the

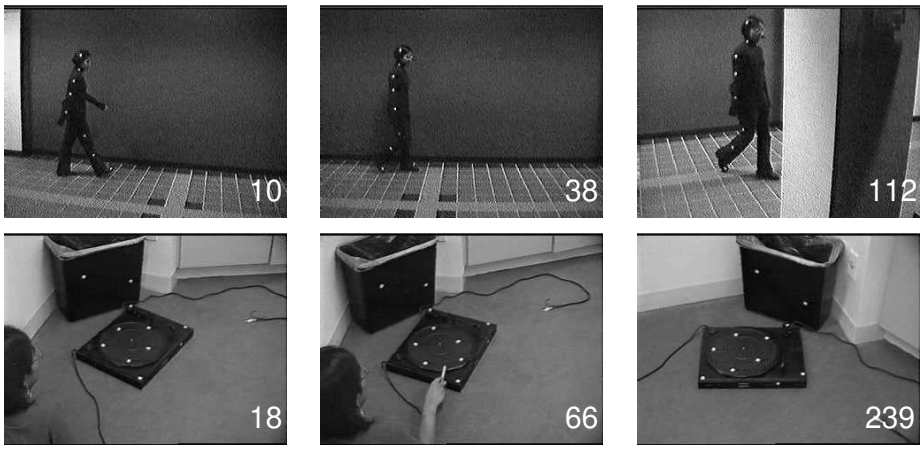

Figure 3. Selected images from image sequences of cyclic motions. Top: Human subject walking. Bottom: Rotating turntable. The second and third images match in each set shown above. These and other correspondences were computed automatically from the period trace.

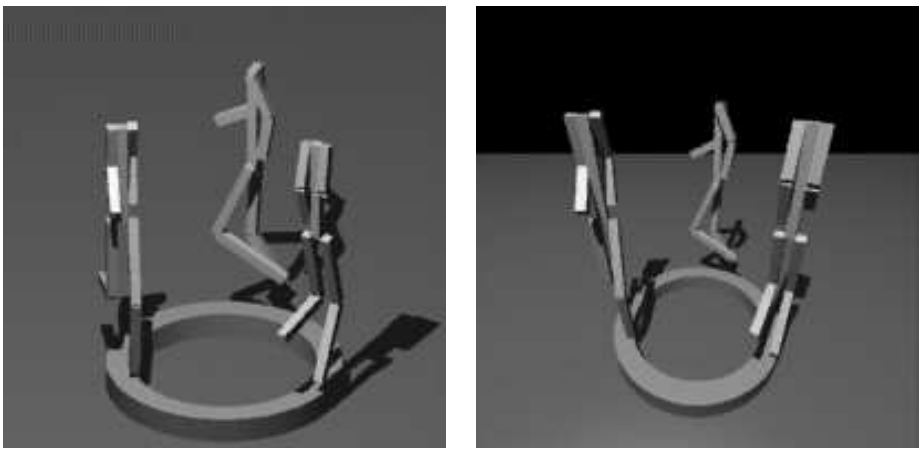

Figure 4. Projections of a simulated jogger moving around a track. Left: The jogger at three different points along the path under orthographic projection. Right: Same scene under perspective projection with the camera 300 units from the track center. Note the projective distortions in the perspective image that are not modeled by orthographic projection.

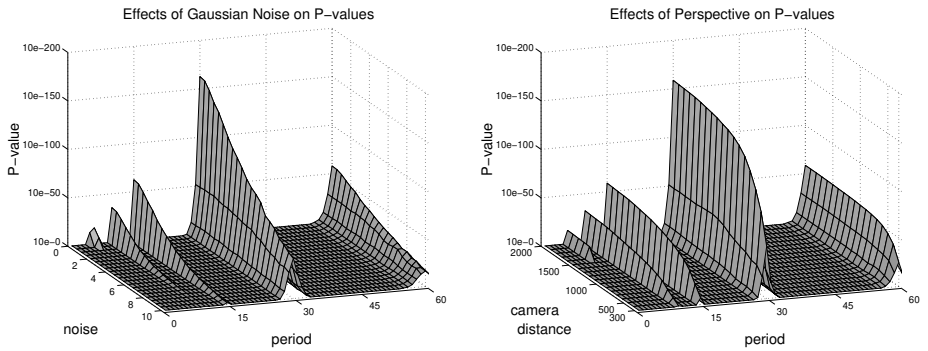

Figure 5. Effects of noise and perspective on period detection. Left: Plot of recovered period P-values under Gaussian noise with standard deviation varying from 0 to 10 pixels. Results were averaged over five trials. Right: Recovered P-values under perspective as camera distance from track center varied from 300 to 2000 units. In all cases the detected period of 30 frames exactly matched the true period.

true period of 30 frames. The most evident effect of noise and perspective was a gradual degradation of P-values, as seen in the Fig. 5.

6.2. HUMAN WALKER

A human subject was filmed walking in an arc subtending about 70 degrees. The camera was manually rotated to keep the subject in view, and the focal length was slowly changed during filming. To aid in feature detection and tracking, reflective markers were placed in areas which were visible for the duration of the sequence (i.e., right arm, right leg, mid torso, and head). We used the method described in Section 5.1 to determine the period with smallest P-value. Fig. 3 shows that the algorithm successfully detected matching images; images 38 and 112 are views of the subject at roughly the same point in her walk cycle. The algorithm detected a period of 37 with a P-value of 10^{-18}. The P-value of the detected period clearly beats out all other candidates (see Fig. 6).

Because features on only one side of the torso were tracked, the feature set was nearly planar. Co-planarity of the feature set is known to be problematic for many shape-from-motion algorithms [18, 19, 20]. Note, however, that it does not cause a problem for our algorithm because an explicit 3D scene representation is never computed. Note also that the entire motion is nearly affine; the features in each image can be roughly approximated by a horizontal shear and/or reflection of the features in the first image. In fact, we have found this property to be true of other human locomotory motions such as running, skipping and jumping. Since the match function d_A equates all images that are related by a 3D affine transformation, purely

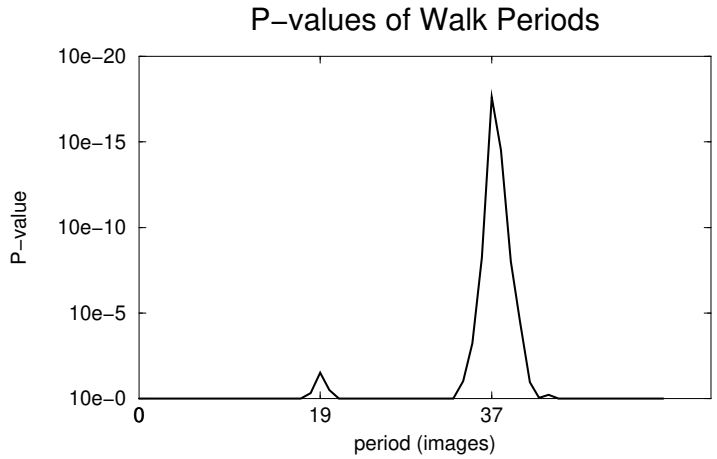

Figure 6. Period P-values of a walking human. The period with the highest level of significance is 37, with a P-value of 10^{-18}.

affine motions appear stationary, and period detection fails. The periods of nearly-affine motions, however, can be reliably detected with our approach, but the resulting P-values tend to be larger.

6.3. ROTATING TURNTABLE

A rotating turntable was filmed using a moving, hand-held video camera. Reflective markers were placed both on the turntable and elsewhere in the static scene. Note that although the motion of the turntable is rigid, the scene as a whole moves non-rigidly. Custom software was used to track the markers as the camera moved around the scene and the turntable simultaneously rotated. Twice the turntable was briefly touched to temporarily slow the rotation and produce an irregularity.

Fig. 7 shows the recovered period trace. The figure illustrates that the optimization procedure successfully located low-energy "valleys" in the match space, corresponding to the dark contours in the temporal correlation plot. The two regions with the highest irregularity values correspond to the two brief intervals in which the turntable was touched. The width of each interval indicates how long the turntable was touched and the irregularity value determines the extent to which the rotation was slowed. Notice that there are ramps in several locations in the period trace. A ramp's presence at point (t,s) indicates that the turntable's speed at time t is different than at time s, since the period is locally changing. Therefore, a motion irregularity occurs at time t *or* at time s. In order to determine where

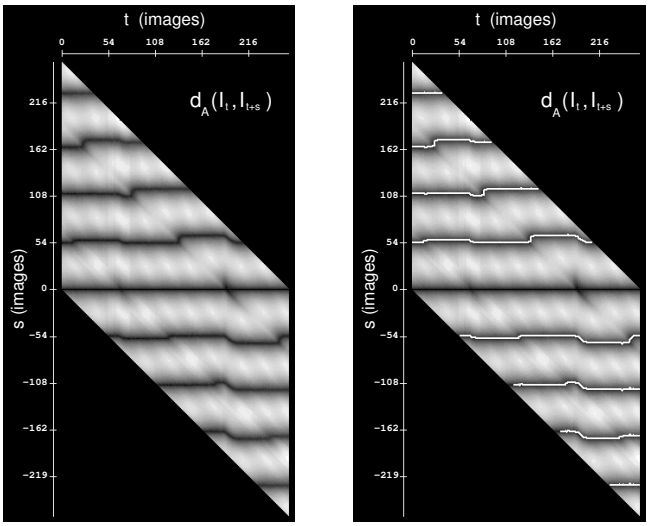

(a) **(b)**

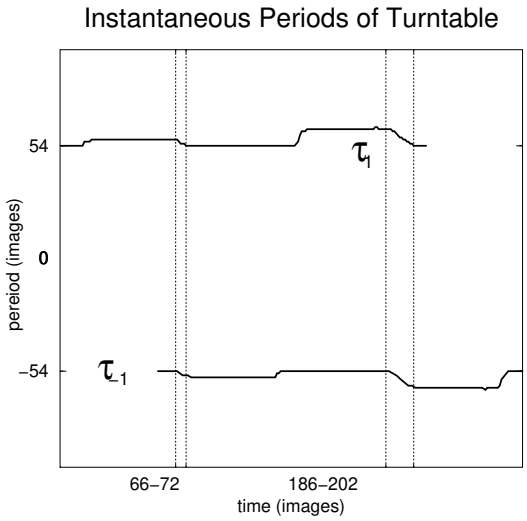

(c)

Figure 7. Period trace of the turntable motion sequence. (a) Temporal correlation plot of $d_{\mathcal{A}}(I_t, I_{t+s})$ with values shown as intensity (light – low correlation, dark – high correlation). (b) The recovered period trace (white) is superimposed on the correlation plot. (c) A graph of τ_1 and τ_{-1} with the two most irregular intervals marked. Notice that discontinuities appear in various places but "line up" only at actual motion irregularities. The turntable frequency is $33\frac{1}{3}$ revolutions per minute, which corresponds to 54 frames per revolution at NTSC video rate.

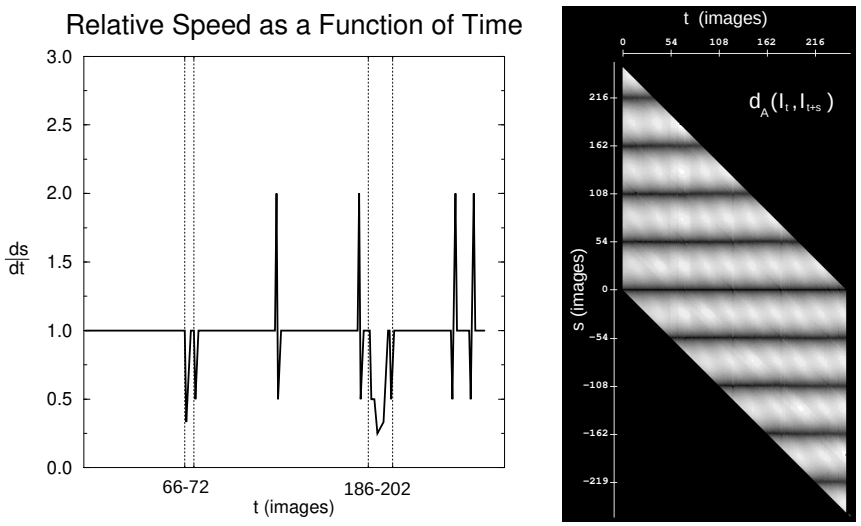

Figure 8. Image sequence rectification using the period trace. The turntable sequence was rectified, producing a periodic sequence with irregularities automatically removed. Left: Speed relative to the first cycle. Touching the turntable slowed the rotational speed, producing momentary reductions in $\frac{ds}{dt}$. Other fluctuations are present but are not significant since the local variations in speed cancel. Right: Temporal correlation plot of the rectified sequence. Observe that the period trace (dark lines) consists of constant functions, indicative of a periodic motion. The period with smallest P-value is 54, exactly matching the ground truth of $33\frac{1}{3}$ rpm at NTSC video rate.

specifically the irregularity occurs, simply look for ramps that "line-up", i.e., that occur at the same interval in $\tau_{-n}, \ldots, \tau_n$. The only ramps that line-up are the two intervals marked in Fig. 7, both of which correspond to intervals where the turntable was touched. Image 66 in Fig. 3 is contained in a detected irregular interval and shows the turntable being touched.

The rectification procedure of Section 4.4 was applied to the image sequence, producing a periodic image sequence (see Fig. 8) comprised of a subset of the images in the original turntable sequence. Notice that the period trace has been *flattened* and no longer contains signs of irregularities. The first complete cycle was automatically detected and used as the reference cycle. In the rectified sequence, the turntable rotates at a constant rate of 54 frames per revolution, exactly matching the ground truth rate of $33\frac{1}{3}$ rpm at NTSC video rate. The turntable sequence was found to be cyclic with a P-value of 10^{-229}. The rectified sequence was determined periodic with a P-value smaller than 10^{-324}, the smallest number representable with 8-byte double floating point precision.

The left graph in Fig. 8 shows instantaneous speed relative to the reference cycle. The slowing of the turntable is reflected by a reduction of

speed in the two intervals in which the turntable was touched. Notice that there are four other momentary fluctuations in speed. These fluctuations, however, are not significant because the local changes in speed cancel out. For example, the rotational speed is twice as fast as normal in frame 126, and half as fast as normal in frame 127. These reciprocal variations exactly cancel each other out, producing no net change in cycle length. This characteristic pattern is common in graphs of this type and is indicative of a small error in the recovered period trace rather than a motion irregularity in the scene. In contrast, motion irregularities generally leave variations that do not cancel, resulting in a net increase or decrease in cycle length during the time in which the irregularity occurs (e.g., frames 66-72 and 186-202 in Fig. 8).

7. Conclusion

This chapter presented a new tool called the *period trace* for describing temporal properties of cyclic motions. By delimiting cycles and locating inter-cycle image correspondences, the period trace provides a framework for tracking changes in an object's motion over time. From this representation, several types of cyclic-motion features were derived, relating to motion trends and irregularities.

A second contribution was an affine-invariant image match function for comparing images under variable viewpoint. Using this function, we derived necessary and sufficient conditions for determining whether an uncalibrated image sequence is the projection of a 3D cyclic motion in the scene. Whereas previous work on affine-invariance was limited to static scenes or rigid-body motion, our results indicate that affine-invariant techniques have application to objects and scenes whose motion is arbitrarily non-rigid.

An energy-based optimization technique was described for recovering the period trace from image sequences obtained with a moving uncalibrated camera. Results on both real and synthetic image sequences indicate that the technique detects the period and other features of cyclic motions with high reliability and is robust with respect to measurement errors and perspective effects.

Acknowledgements

The support of the National Science Foundation under Grant Nos. IRI-9220782 and CDA-9222948 is gratefully acknowledged.

References

1. T. S. Perry, "Biomechanically engineered athletes," *IEEE Spectrum*, pp. 43–44,

April 1990.
2. M. Allmen and C. R. Dyer, "Cyclic motion detection using spatiotemporal surfaces and curves," in *Proc. 10th Int. Conf. on Pattern Recognition*, pp. 365–370, 1990.
3. R. Polana and R. Nelson, "Detecting activities," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 2–7, 1993.
4. P. Tsai, M. Shah, K. Keiter, and T. Kasparis, "Cyclic motion detection for motion based recognition," *Pattern Recognition*, vol. 27, no. 12, pp. 1591–1603, 1994.
5. G. Johansson, "Visual perception of biological motion and a model for its analysis," *Perception and Psychophysics*, vol. 14, no. 2, pp. 201–211, 1973.
6. D. Hogg, "Model-based vision: A program to see a walking person," *Image and Vision Computing*, vol. 1, no. 1, pp. 5–20, 1983.
7. K. Rohr, "Incremental recognition of pedestrians from image sequences," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 8–13, 1993.
8. N. H. Goddard, *The Perception of Articulated Motion: Recognizing Moving Light Displays*. PhD thesis, University of Rochester, Rochester, NY, 1992.
9. T. Darrell and A. Pentland, "Space-time gestures," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 335–340, 1993.
10. J. W. Davis and M. Shah, "Visual gesture recognition," *IEE Proc. Vision, Image and Signal Processing*, vol. 141, no. 2, pp. 101–106, 1994.
11. Y. Yacoob and L. Davis, "Computing spatio-temporal representations of human faces," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 70–75, 1994.
12. I. A. Essa and A. Pentland, "A vision system for observing and extracting facial action parameters," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 76–83, 1994.
13. R. Polana and R. Nelson, "Recognition of motion from temporal texture," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 129–134, 1992.
14. C. Cedras and M. Shah, "Motion-based recognition: A survey," *Image and Vision Computing*, vol. 13, no. 2, pp. 129–155, 1995.
15. A. M. Baumberg and D. C. Hogg, "An efficient method for contour tracking using active shape models," in *Proc. Workshop on Motion of Non-Rigid and Articulated Objects*, pp. 194–199, 1994.
16. S. M. Seitz and C. R. Dyer, "Affine invariant detection of periodic motion," in *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 970–975, 1994.
17. S. M. Seitz and C. R. Dyer, "Detecting irregularities in cyclic motion," in *Proc. Workshop on Motion of Non-Rigid and Articulated Objects*, pp. 178–185, 1994.
18. J. J. Koenderink and A. J. van Doorn, "Affine structure from motion," *J. Opt. Soc. Am. A*, vol. 8, pp. 377–385, 1991.
19. L. S. Shapiro, A. Zisserman, and M. Brady, "3D motion recovery via affine epipolar geometry," *Int. J. of Computer Vision*, vol. 16, pp. 147–182, 1995.
20. C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: A factorization method," *Int. J. of Computer Vision*, vol. 9, no. 2, pp. 137–154, 1992.
21. G. W. Stewart, *Introduction to Matrix Computations*. New York, NY: Academic Press, 1973.
22. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*. Cambridge, MA: Cambridge University Press, 1988.
23. M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. of Computer Vision*, vol. 1, no. 4, pp. 321–331, 1988.
24. D. J. Williams and M. Shah, "A fast algorithm for active contours and curvature information," *CVGIP: Image Understanding*, vol. 1, no. 55, pp. 14–26, 1992.
25. S. M. Seitz and C. R. Dyer, "View-invariant analysis of cyclic motion," *Int. J. of Computer Vision*, 1997. To appear.