Phase Models in Gait Analysis

Jeffrey E. Boyd
Department of Computer Science
University of Calgary
Calgary AB T2N 1N4
boyd@cpsc.ucalgary.ca

James J. Little
Department of Computer Science
University of British Columbia
Vancouver BC V6T 1Z4
little@cs.ubc.ca

Abstract

We examine both model-free and model-based gait analysis using a framework built upon three requirements in gait perception: frequency entrainment, phase locking, and physical plausibility. A key element of the framework is the use of frequency and phase information present in the oscillations derived from a gait. This leads to phase-based models of gait. We present a set of tools for phase models and describe their use in both model-free and model-based gait analysis. The common framework allows a comparison of methods and suggests a path to connect model-free representations with kinematic models.

1. Introduction

Computer vision research in human motion analysis can be divided into two broad categories: model-based and model-free. The model kinematic, describing the interconnected parts that make up the human body and their kinetic relationship to each other. Model-free methods are sometimes referred to as *appearance-based*, and may also be referred to as exemplar methods.

In the model-free category Little and Boyd [16] describe a *shape-of-motion* system that analyzes periodicities in the shapes of people as they walk. The analysis yields features that are suitable for recognizing individual gaits. Baumberg and Hogg [2, 3] describe the shape of the silhouette of walkers. Bobick and Davis [5, 6] accumulate motion information into a motion energy image and a motion history image to recognize human activities. Boyd [7], Liu and Picard [17], and Polana and Nelson [21] look at oscillations in image intensities due to gait in order to analyze gaits. Cutler and Davis [10] describe a system that identifies periodicities in the motion of objects in order to identify human and animal gaits. Haritaoglu et al. describe a system that classifies a silhouette of a human into one of four main postures, and then further classify it into one of three view-based appearance. Based on the classification they label the body parts of the subject.

In the model-based category Hunter et al. [13], Rowley and Rehg [22], and Bregler and Malik [9] use models that represent joints and limbs in a human. They estimate the pose of a person by extracting joint angle information from a sequence of images. Wachter and Nagel [25] describe a system that derives pose by tracking individual body parts using constraints imposed by the known kinematic structure. While these approaches use an explicit kinematic model, the model need not be so detailed. Wren et al. [26] track a set of *blobs* that are parts of a person. Their system identifies some body parts based on their expected locations in typical views. They do not strictly enforce a kinematic model. Fujiyoshi and Lipton [11] use a simplified model of a person that maps a simplified *star* skeleton onto the silhouette of a moving person. They analyze periodic changes in the skeleton to recover features used to identify moving objects.

We note that some model-based methods do not really do motion analysis although their input may be a temporal image sequence. For example, Hunter et al. [13] compute a pose estimate of a human figure in each frame of a video sequence. They use an iterative expectation-maximization algorithm to arrive at the estimate. The only sharing of information between frames occurs when the estimate from one frame is used as an initial estimate for the following frame. In other cases, motion may be used to help extract parameters of a model, but the model itself lacks temporal information.

The salient feature of a gait is its periodicity. Other motions are either not periodic or exhibit periodic behavior that is much different that that of a gait, e.g., a person waving an arm. Some of the many approaches to analyzing human motion cited above, are not uniquely aimed a gaits. These methods look at broader categories of motion and so they do not exploit periodicity. In this paper we consider only gait, and therefore consider only that part of human motion that is periodic.

When considering the many gait analysis methods available, it is important to consider which methods are appropriate for answering which questions such as:

1. Does the motion in an image represent a human gait as opposed to an animal gait, a car, a tree blowing in the wind, or some other type of motion?

2. Can you recognize which person is walking by their gait?

3. How is the person moving? Are they running or walking, skipping or shuffling or hopping?

4. Is the person carrying anything? Which arm are they using to carry?

5. Is the gait unusual due to some pathology?

A model-free approach is well-suited to the first three questions while the fourth suggests the need for a model. The fifth question may be approached both ways.

We examine the issues involved with both model-free and model-based gait analysis. Requirements for gait perception identified in the psychology literature provide a common framework for understanding both approaches. The requirements highlight the importance of frequency and phase in the oscillations that are inevitably observed when analyzing a gait. We present tools for dealing with the frequency and phase information. The common framework suggests a path for connecting the two approaches to yield a continuum between model-free and model-based methods.

2. Requirements for Gait Analysis

Bertenthal and Pinto [4] identify the following three important properties in the human perception of gaits.

- *Frequency entrainment*. The various components of the gait must share a common frequency.

- *Phase locking*. The phase relationships among the components of the gait remain approximately constant. The lock varies for different types of locomotion such as walking versus running.

- *Physical plausibility*. The gait must be a stable solution to an equation of motion. Each joint in a gait forms a stable phase plane trajectory, or limit cycle. Bertenthal and Pinto refer to these as periodic attractors.

We refer to these as the FPP requirements (Frequency-Phase-Physics). It is reasonable to assume that machine vision systems must also meet FPP requirements to perceive and analyze gaits.

The FPP requirements form the basis of a framework for understanding gait analysis systems. Oscillations are at the heart of gait analysis. For that reason frequency entrainment and phase locking can be found in both model-based and model-free methods. Physical plausibility requires a model because the physics must be evaluated with respect to a model. We can understand gait analysis methods in terms of how they perform, or perhaps fail to perform, each requirement.

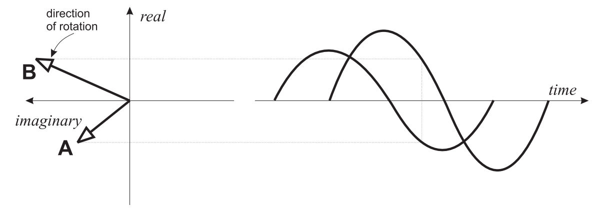

Figure 1. A phasor, or phase vector, is a complex vector that rotates about the origin, generating a sinusoid when projected onto the real axis. The magnitude and direction of the vector gives the amplitude and phase of the sinusoid respectively. Timing is given by the relative phases. Here phasor $\mathbf{A}$ leads phasor $\mathbf{B}$.

3. Tools

3.1. Phasors

The key to interpreting complex motions lies in the timing of events in the motion. When the motion is cyclical, as it is with a gait, then timing reduces to phase relationships among oscillations. Phasors are an elegant tool for dealing with a collection of sinusoids oscillating at a common frequency. For this reason they are well-known to physicists and electrical engineers, and are commonly used in analysis of power distribution systems where voltages and currents oscillate at a common frequency of 60*Hz*.

The following summary of phasors is adapted from Smith [23]. Consider the function $\mathbf{W}(t)$, a complex function of time, where

$$\mathbf{W}(t) = Ae^{i(\omega t+\phi)},$$

$i = \sqrt{-1}$. The real and imaginary parts of $\mathbf{W}(t)$ are

$$\begin{aligned} W_{real} &= A\cos(\omega t+\phi) \\ W_{imag} &= A\sin(\omega t+\phi) \end{aligned}$$

Therefore, for a sinusoidal function with amplitude A, frequency ω, and phase ϕ such as

$$v(t) = A\cos(\omega t+\phi)$$

we can say that

$$v(t) = \mathrm{Re}\{Ae^{i(\omega t+\phi)}\} = \mathrm{Re}\{Ae^{i\omega t}e^{i\phi}\}, \qquad (1)$$

i.e., a sinusoid is the projection of a rotating complex vector onto the real axis, shown in Figure 1. If we have a set of signals with a common frequency (which is reasonable when analyzing gaits), it is convenient to drop the $\exp(i\omega t)$ term (it acts as a rotating reference for the signals) and deal only with the complex number $Ae^{i\phi}$. We then represent the sinusoid $v(t)$ with the complex *phasor* (for *phase vector*)

$$\mathbf{V} = Ae^{i\phi} = A\angle\phi.$$

A direct consequence of the phasor representation is that we can add two sinusoidal functions simply by adding their phasors in the complex plane, i.e., manipulations of sinusoidal functions are reduced to operations on vectors.

The sinusoids that arise in the analysis of gaits can be viewed as phasors. Since components of a gait share a common frequency, the phasors rotate at a common frequency, ω, due to frequency entrainment. The amplitude, A, varies depending on the signal, and the phase, ϕ, indicates the relative timing of the sinusoids. For example, in Figure 1, phasor **A** leads phase **B**. Although the phasors rotate, the relative phases are constant because all phasors are rotating at the same frequency. This is phase locking. Frequency entrainment and phase locking are, therefore, mutually required. An excellent tool for visualizing phasors can be found at the web site *http://robotics.eecs.berkeley.edu/~cssharp/Phasors/Phasors.html*.

3.2. Directional Statistics

Gaits analysis yields sinusoidal signals, phasors are a good way to represent sinusoids, and a phasor is a vector that indicates a direction. Therefore, we need to deal with directional data in gait analysis. Procrustes shape analysis [14] is a method in directional statistics [18] that can find mean shapes, i.e., patterns of points in two-dimensions, and a measure of the similarity of two shapes. By virtue of using phasors we can represent a phase or timing pattern in a gait as a set of directions, which is mathematically equivalent to a shape, making Procrustes analysis a useful tool for analyzing gaits.

The following is a summary based on Mardia and Jupp [18]. Describe a shape in two-dimensions using a vector of k complex numbers, $\mathbf{z} = [z_1, z_2, \ldots, z_k]^T$, called a configuration. Two configurations, $\mathbf{z}_1$ and $\mathbf{z}_2$, represent the same shape if by a combination of translation, scaling, and rotation, their configurations are equal, i.e.,

$$\begin{aligned} \mathbf{z}_1 &= \alpha\mathbf{1}_k + \beta\mathbf{z}_2, \quad \alpha, \beta \in \mathcal{C} \\ \beta &= |\beta| e^{i\angle\beta}, \end{aligned}$$

as shown in Figure 2(a) and (b). That is, $\alpha\mathbf{1}_k$ translates $\mathbf{z}_2$, and $|\beta|$ and $\angle\beta$ scale and rotate $\mathbf{z}_2$. It is convenient to center shapes by defining the centered configuration $\mathbf{u} = [u_1, u_2, \ldots, u_k]^T$, $u_i = z_i - \bar{z}$, and $\bar{z} = \sum_{i=1}^{k} z_i/k$. From here we can define the *full Procrustes distance* between two configurations, $d_F(\mathbf{u}_1, \mathbf{u}_2)$ as:

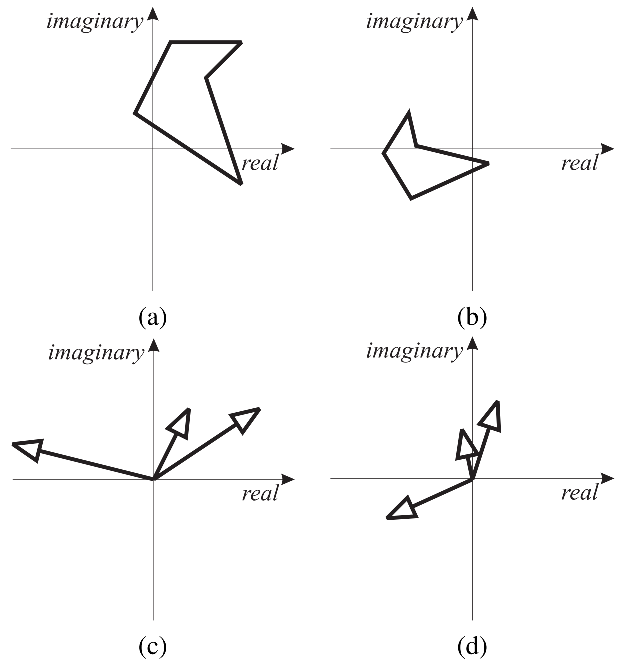

Figure 2. Procrustes analysis applied to shape and phase configurations. In the conventional application, a shape is represented by a vector of complex vertices. The shape in (a), is the same as the shape (b) because one is a translated, scaled, and rotated version of the other. A phasor configuration is also a vector of complex numbers. The configuration in (c) is the same as that in (d) because one is a rotated and scaled version of the other. Rotation is always about the origin so translation is omitted.

$$d_F(\mathbf{u}_1, \mathbf{u}_2) = 1 - \frac{|\mathbf{u}_1^* \mathbf{u}_2|^2}{\|\mathbf{u}_1\|^2 \|\mathbf{u}_2\|^2} \tag{2}$$

which minimizes

$$\left\| \frac{\mathbf{u}_1}{\|\mathbf{u}_1\|} - \alpha\mathbf{1}_k - \beta\frac{\mathbf{u}_2}{\|\mathbf{u}_2\|} \right\|^2. \tag{3}$$

$\mathbf{u}^*$ denotes the transpose of the complex conjugate of $\mathbf{u}$. Note that $0 \leq d_F \leq 1$. The Procrustes distance allows us to compare two shapes independent of position, scale, and rotation. We can find the mean of a set of n shapes by finding the μ that minimizes the objective function

$$\min_{\alpha_j, \beta_j} = \sum_{j=1}^{n} \|\mu - \alpha_j\mathbf{1}_k - \beta_j\mathbf{u}_j\|^2. \tag{4}$$

To find μ we compute the matrix

$$\mathbf{S}_u = \sum_{j=1}^{n} (\mathbf{u}_j \mathbf{u}_j^*)/(\mathbf{u}_j^* \mathbf{u}_j). \tag{5}$$

The *Procrustes mean shape*, $\hat{\mu}$, is the dominant eigenvector of $\mathbf{S}_u$, i.e., the eigenvector that corresponds to the greatest eigenvalue of $\mathbf{S}_u$.

Although Procrustes shape analysis is intended for treating two-dimensional shapes, it is easily adapted to handling vectors phasors [7]. A vector of complex phasors, or a phasor configuration, is equivalent to a shape configuration, as illustrated in Figure 2(c) and (d). Shapes are invariant through translation, scaling, and rotation. Translational invariance is achieved by using the centered configuration $\mathbf{u}$. When using phasors the issue of translation becomes irrelevant. All phasors rotate about the origin, $0 + 0i$, at the entrained frequency. The configurations, $\mathbf{z}$, are already centered, so $\mathbf{z} = \mathbf{u}$.

Procrustes analysis gives a method to compute the mean of a set of phasor configurations and also to compare phasor configurations. In the practice of gait analysis, we form phasor configurations as part of some analysis of a motion sequence, and then use the Procrustes method to perform desired tasks such as describing or recognizing a gait based on the configuration.

4. Phase-Based Gait Analysis

The FPP requirements, phasors, and Procrustes analysis give us a framework for understanding gait analysis systems. FPP tells us where to look for the important information, phasors give us a way to represent it, and Procrustes analysis allows us to manipulate it to perform useful tasks. We now cast both model-free and model-based gait analysis in this framework.

4.1. Model-Free

In a model-free approach to image analysis, oscillations are found in two places:

1. the shape or silhouette of a walking person oscillates, and

2. pixel intensities in an image sequence containing a gait oscillate.

Examples of shape or silhouette approaches include Boyd and Little [16], and Baumberg and Hogg [2, 3]. Pixel intensity approaches are found in Polana and Nelson [21], Liu and Picard [17], and Boyd [7]. The following treats shape and pixel oscillations separately.

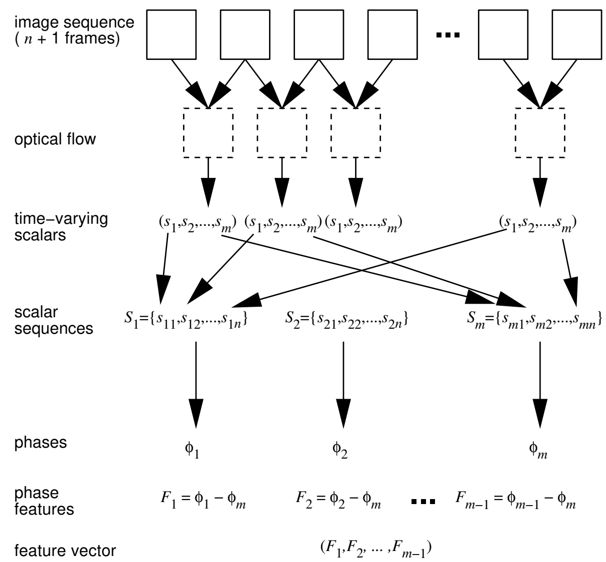

Figure 3. The shape-of-motion gait recognition system.

Shape Oscillations: Consider the shape-of-motion method of Little and Boyd [16] show in Figure 3. For each frame in a gait sequence they find a set of scalars that describes the shape of the moving region. Each of these scalars oscillates throughout the gait cycle at twice the fundamental frequency of the gait, i.e., at the step rate. Thus the various scalar signals are frequency entrained. The combined phases of the oscillating scalars produce a phase feature vector. One element of the vector is used as a phase reference and is therefore subtracted from the others. An independent analysis of variance on each of the remaining phases verifies that phase depends on the individual subject that is walking. The phase feature vectors are also used in a system to recognize individual subjects using phase exemplars. This method effectively uses variations in phase lock to recognize people.

While the shape-of-motion system uses phase locking, the analysis does not exploit the directional nature of the phase data. Subtracting a phase reference effectively adds errors in the reference to the phases that remain in the feature vector. Furthermore, all phase values are treated as numbers on a line. Rotation of phases for some scalars by 180° avoids the problems associated with phase wrapping. Alternatively, treating a phase as a phasor of unit length and using Procrustes analysis avoids these problems. Rotational invariance of the phasor in the Procrustes method obviates the need to subtract a phase reference and propagate the error in the reference. The analysis of variance for directions [18] and Procrustes shape analysis for recognition provide more elegant tools for analyzing the shape-of-motion phase data without requiring special treatment for phase wrapping.

The method of Cutler and Davis [10] is also a shape method. They identify periodicities in a vector of intensities for a tracked object. The process effectively does frequency entrainment. Phase-locking is limited though. The positions of peaks in a similarity image indicate the phase in the cycle of the gait, but does not indicate phase locking between component motions of the gait.

Shape-of-motion is view-dependent. Little and Boyd observed that recognition of gaits was possible from different viewing angles, but that the recognition had to be trained independently for each angle. This is not surprising given their model-free approach. However, the use of phase features results in scale-invariance, provided that the subject is large enough in the image to allow the system to measure the scalar shape descriptors well. Similarly, the method of Cutler and Davis is scale-invariant. They require that the object size remain nearly constant for a single sequence but no scaling between sequences is necessary. Their method is view-invariant since they only require to identify temporal symmetries to distinguish between human and other motions. More complex tasks may reveal view-dependence.

Pixel Oscillations: Boyd [7] uses video phase-locked loops (vPLLs) to analyze the oscillations in pixels due to a gait. A vPLL, shown in Figure 4, synchronizes the oscillations of an internal oscillator at each pixel with the intensity oscillations observed at the pixel. This synchronization performs both frequency entrainment and phase locking. Polana and Nelson [21] and Liu and Picard [17] extract similar information using fast Fourier transforms, but discard the phase information that is crucial for phase locking. In the vPLL system, a phasor represents the oscillations at each pixel. The image of phasors representing the oscillations throughout the image can be arranged as a phasor configuration making it possible to average over multiple frames of video, average over multiple subjects, and to compare configurations using the full Procrustes distance. Boyd illustrates the use of phasor configuration exemplars to compare gaits and other oscillations. For example, Figure 5 shows three-element phasors configurations for two gait sequences. Note that the configurations are both frequency entrained and phase-locked.

The vPLL method is view-dependent too. Unfortunately, the use of phases here does not produce scale-invariance. Correspondence of the positions of the phasors in the phasor configuration is critical and thus requires both the view and scale to be fixed. A close examination of the full Procrustes distance shows why this is the case. Equation (2) is just a simple variation of the rotationally invariant matched spatial filters described by Hsu et al. [12]. Thus, in this context, the pixel-based phasor configurations share some of the limitations of matched spatial filters.

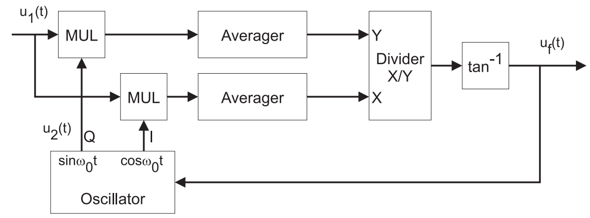

Figure 4. A digital-averaging phase-locked loop. A vPLL uses one loop at each pixel in an image to synchronize the oscillator to the oscillations in pixel intensity.

4.2. Model-Based

Oscillations in a kinematic model occur in the joint angles and also in the positions of the limbs. Often the goal of model-based analysis is to recover the joint angles or limb positions. This is adequate where the goal of the analysis is to capture motion. When the motion of interest is a gait, a simple record of joint angles or limb positions is not adequate to identify an individual gait, nor to answer subjective questions about the style of a gait. For, this one must consider the FPP requirements for gait perception.

Boyd and Little [8] show a spectral analysis of the spatial coordinates of some joints in a walking figure. All joints share the fundamental frequency of the gait, the frequency of steps with the same foot. Other frequencies at integer multiples of the fundamental are also observed. It is possible to observe frequency entrainment by finding the frequencies of oscillations in the model parameters.

Phase-locking can be observed with limb motions. Boyd and Little [8] show variations in phase lock for subjectively different gaits generated by Poser [19], a human modeling software package. For example, Figure 6(a) shows the phasor configuration for a normal gait while (b) shows the configuration for a person that is shuffling. In the normal gait, each hand is nearly in phase with its opposite foot. In contrast, the hand motion is shifted by about 90° and is smaller in magnitude. Variations in phase locking are certainly used within the graphics community to model different gaits [24, 1]. Although we do not know of an example of phase locking in a kinematic model used for gait analysis or recognition, there is ample evidence to suggest that it is possible.

Bertenthal and Pinto [4] describe the limit cycles variations that occur between walking and running, highlighting the use of physical plausibility in gait analysis. Limit cycles have been used to synthesize realistic gaits [15], and it should be possible to use them for analysis and recognition tasks. Since model-free methods cannot derive this information, this is an area where model-based methods can

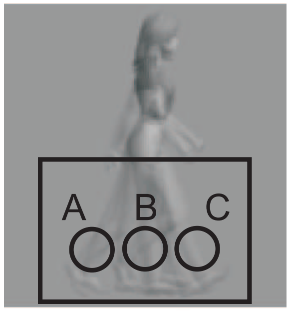

(a)

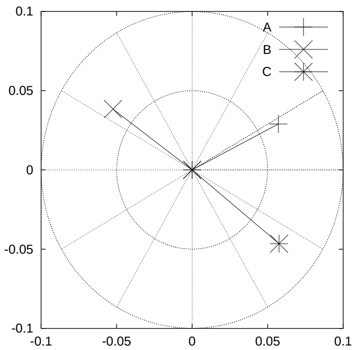

(b)

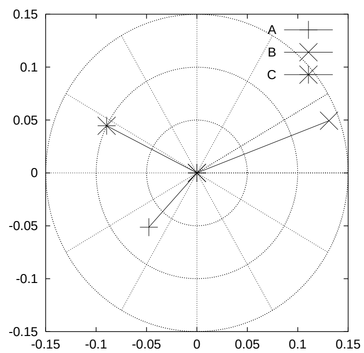

(c)

Figure 5. Comparison of pixel oscillation phasor configurations for two gait sequences: (a) the positions in the image from which the 3-element configurations are extracted, (b) configuration for a synthetic gait, and (c) configuration for a real person walking on a treadmill. The full Procrustes distance between the two configurations is 0.19.

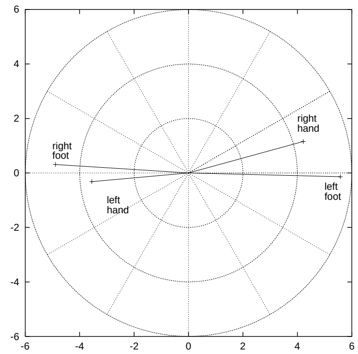

(a)

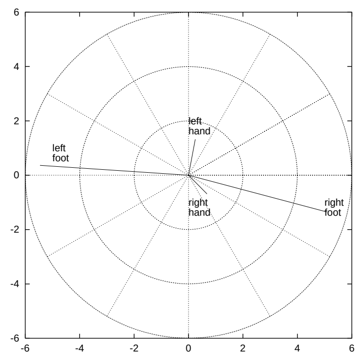

(b)

Figure 6. Example phasor configurations of joint position for subjectively different motion for (a) a normal gait, and (b) a shuffle.

excel.

5. Connecting Exemplars to Kinematic Models

We can answer many useful questions about gait without a kinematic model, but it is clear that for a complete analysis, the model is required at some point. However, connecting model-free and model-based interpretations of a gait is difficult. The path between the two representations has not been clear.

The FPP requirements of gait perception suggest a mechanism to connect the model-free to the model-based realm. The oscillations measured in gait analysis are all derived from the subject. Therefore, the frequency entrainment, phase locking, and physical plausibility are all properties of the walker. The measurements made by both model-based and model-free systems of these properties are all different views of the same underlying phenomenon.

Consider the frequency entrainment requirement. Typically a model-free method entrains to the step rate, twice the fundamental frequency, of the gait. Thus the model-free frequency and the model-based frequency are the same, within a factor of two. In fact the frequencies measured can only differ by an integer factor. No other relationship is possible for a gait.

Phase locking is more complicated. With a kinematic model, phase locking appears in the joint and limb oscillations. In a model-free approach the phase is in the oscilla-

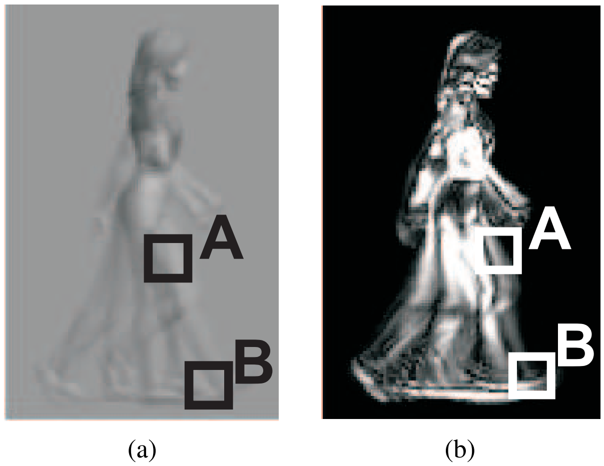

Figure 7. Example vPLL analysis of a synthetic gait: (a) superposition of frames from the input sequence, and (b) the magnitude output of the vPLL.

tions of shape descriptions or pixel oscillations. The path through the imaging system from joint and limb motion to images and ultimately shape descriptions is complex and non-linear. Boyd and Little [8] demonstrate that it is unlikely that a simple relationship between phase in the model and phase in the shape exists.

Connecting pixel oscillations to a kinematic model offers more hope. Figure 7 illustrates why. The oscillations viewed at position **A** are due to the left and right thighs alternately swinging forward to occlude the background. Therefore, the oscillations in these pixels must be synchronized with the oscillations of the thighs. Although the pixel oscillations do not distinguish between left and right thighs, it is reasonable to assume that left and right limbs are phase locked at $180°$ and that each oscillates at half the frequency of the pixel oscillations. The oscillations at position **B** are due to the left and right feet alternately swinging forward. The phase at **B** with respect to the phase at **A** indicates the timing of the shin movement with respect to the thigh. Thus it is possible to extract the frequency entrainment and phase locking information required for a kinematic model from the model-free pixel oscillations. The difficulty in using this method lies in the finding the correspondence between pixels and body parts, i.e., the selection of the points **A** and **B**. The pixel positions for limbs are both view and scale dependent. Nevertheless, once the correspondence has been established, the connection between model-free and model-based representations is also established.

In addition to the timing information that can be transferred from the model-free to the model-based domain, it should be possible to derive relative sizes of body parts from model-free analysis. For example, one could determine which portions of a pixel oscillation phasor configuration belong to the torso, and which belong to the legs. From this information one can derive the relative length of the legs with respect to the torso. Similarly, the relative size of the thighs and shins could be derived. This size information combined with timing can yield a kinematic model that represents both joint angles and (within a multiplicative constant) limb sizes.

6. Discussion

Table 1 summarizes model-free and model-based approaches to gait analysis. Some work cited in the table does not refer exclusively to gaits, but analyzes human motion in general. These are included in the table because the work is aimed at analyzing human motion in general, and therefore has the potential for application to gaits.

The table differentiates between the methods based on their advantages and disadvantages. While this shows some important distinctions, we feel a more profound understanding can be gained by considering what questions can be answered as a result of a particular analysis. The following discussion lists some advantages and disadvantages, then focuses on a set of questions and how the various methods can or cannot answer them.

6.1. Advantages and Disadvantages

After the segmentation of the shape of a moving region, computation of shape features is fast, and so is the analysis of the resulting signals. If segmentation is done using background subtraction or a fast implementation of optical flow, then model-free recognition based on shape oscillations should be able to run in real-time on current hardware. Computation of magnitude and phase of pixel oscillations is also fast either by FFT [21, 17], or using vPLLs [7]. In contrast, model-based methods that recover a full kinematic model tend to be slow. For example, the method of Hunter et al. [13] requires solution of an iterative expectation-maximization problem off-line. The process benefits from using solutions to previous frames as a starting point for iterations, but the method is still far from real-time frame-rate performance. Methods that use a less sophisticated model such as the star model of Fujiyoshi and Lipton [11] can run faster.

Model-free methods tend to be view-dependent because the shapes and correspondence between pixels and body parts changes with view. These methods can be made to work well for a set of predefined viewing angles. Shape methods can be scale-invariant by either using scale-invariant shape descriptors, or by looking at phases which are inherently scale-invariant. On the other hand, model-based methods tend to be view- and scale-invariant. Typically the model is manipulated to match the scale and view

	oscillators		
	Model-Free (exemplars)		*Model-Based*
	shape	pixel intensity	joint angles
Examples	• Baumberg and Hogg [2, 3] • Little and Boyd [16] • Cutler and Davis [10] • Bobick and Davis [5, 6]	• Polana and Nelson [21] • Liu and Picard [17] • Boyd [7]	• Fujiyoshi and Lipton [11] • Hunter et al. [13] • Rowley and Rehg [22] • Wachter and Nagel [25] • Wren et al. [26] • Bregler and Malik [9]
Frequency Entrainment	Shape descriptors oscillate at same frequency	Pixel intensities oscillate at same frequency	Joint angles oscillate at same frequency
Phase Locking	Phasor configuration of shape descriptors fixed	Phasor configuration of pixel pixel intensities fixed	Phasor configuration of joint angles fixed
Physical Plausibility	NA	NA	joint angles move in a physically plausible manner
Advantages	• Computation of shape descriptions is fast. • Analysis of shape signals is fast. • Some shape extraction methods are fast. • Scale-invariant. • Little or no camera calibration	• Both FFT- and PLL-based methods are fast. • Keeps detailed spatial information. • Little or no camera calibration	• Can answer questions related to kinematic model. • Speed is acceptable for simple models. • View- and scale-invariant.
Disadvantages	• Can be slow for some optical flow methods. • Answers appearance-based questions only. • View-dependent. • Detailed spatial information lost in shape descriptors. • Requires training examples.	• Answers appearance-based questions only. • View and scale-dependent. • Requires training examples.	• Computation is slow for full kinematic models. • Usually requires camera calibration. • Assumes correct model is known.
Questions	• Human locomotion versus animal or other motion?		• Is the person carrying anything? In which arm are they carrying it?
	• If human, which human? • Is the subject walking normally, skipping, shuffling? (subjective description of gait) • Is the gait unusual due to some pathology?		

Table 1. Summary of gait analysis approaches in FPP requirement framework. NA means not applicable. Examples include some methods that do not apply strictly to gaits.

of the data. The penalty for this is generally slower algorithms and often the need for camera calibration.

Model-based approaches assume that the appropriate model is known a priori. It makes little sense to match the model of a horse to data for a human. In order for a model-based method to apply to motion in general, it requires a library of models and a mechanism to select the correct model. Model-free methods do not have this requirement. However, in order for a model-free method to recognize a motion, it must be trained with examples of similar motions. Model-free methods cannot recognize motions that they are not trained for.

6.2. Questions

Is a moving object human, animal, or something else? Answering this question is clearly the domain of model-free methods since the question must be answered without knowing which model to use. A system that answers this question has been demonstrated [10]. It may be argued that a model-based method could test a set of models and answer this question by determining which model best fits the data. While this may work, the computational requirements make such an approach impractical.

If it is a human, which human? Since this question assumes that we already know which model to use, it is in both the model-free and model-based domains. Recognition of individuals has been demonstrated using a shape-based model-free approach [16] relying on unique phasor configurations of shape oscillations. Although we do not know of an example, it should also be possible to recognize individuals based on phasor configurations of joint angle oscillations.

Is the subject walking normally, skipping, or shuffling? This question calls for a subjective description of a gait. Clearly it can be answered using a model-based approach since variations in joint angle data have been used to synthesize subjectively different gaits in computer graphics. However, if it is possible to recognize individual gaits without a model, it should also be possible to recognize subjectively different gaits without a model. This has not been demonstrated though.

Is the gait unusual due to some pathology? If a gait varies due to some pathology of the subject we certainly expect to see it in the joint angle model. Evidence of the pathology should be in asymmetry between left and right body, in timing of motions as the subject alters their gait to minimize pain or favor an injured muscle, or in the magnitude of oscillations. Variations due to pathology should also be evident in the model-free phasor configurations as shown by Murray [20] who relates variations in a subjects centroid (a shape descriptor) with patient pathology.

Is the person carrying anything? In which arm are they carrying it? If a subject is carrying an object we expect the position of their arms to differ from their normal gait. This change in posture indicates that a model-based approach is well-suited to answer the first question. Note that it may be possible to answer the question without a model. The act of carrying an object will likely alter the observed shape of a subject and the oscillations in pixels occupied by the subject. This may be sufficient to determine whether or not a subject is carrying an object. However, if we combine all the factors that we believe can alter the model-free perception of a gait (who is it, are they tired, are they carrying something, are they injured, and so on) we can expect to get a feature space with large overlapping regions. For example, we may conclude from a model-free analysis of gait data that we are seeing person A, or person B limping, person C carrying a briefcase, or some other combination of factors. To increase the resolution of the feature space will likely require movement to some sort of model, although not necessarily a full kinematic model. The second question requires a kinematic model since we must identify an individual arm to answer it.

7. Conclusions

Gaits are unique among human motions in producing a complex pattern of oscillations. In spite of the complexity, the FPP requirements provide a mechanism for understanding gaits and comparing gait analysis methods. Phasors and Procrustes analysis give tools that are well-suited to gait analysis. Using these tools, an FPP framework describes the relationship among the various gait analysis methods. It also gives us a mechanism that has the potential to connect model-free exemplar methods with model-based methods.

References

[1] K. Amaya, A. Bruderlin, and T. Calvert. Emotion from motion. In *Graphics Interface 96*, 1996.

[2] A. M. Baumberg and D. C. Hogg. Learning flexible models from image sequences. Technical Report 93.36, University of Leeds School of Computer Studies, October 1993.

[3] A. M. Baumberg and D. C. Hogg. Learning spatiotemporal models from training examples. Technical Report 95.9, University of Leeds School of Computer Studies, March 1995.

[4] B. I. Bertenthal and J. Pinto. Complementary processes in the perception and production of human movements. In L. B. Smith and E. Thelen, editors, *A Dynamic Systems Approach to Development: Applications*, pages 209–239. MIT Press, Cambridge, MA, 1993.

[5] A. F. Bobick and J. W. Davis. An appearance-based representation of action. In *13th International Conference on Pattern Recognition*, Vienna, Austria, August 1996.

[6] A. F. Bobick and J. W. Davis. Real-time recognition of activity using temporal templates. In *Third International Workshop on Applications of Computer Vision*, Sarasota, Florida, December 1996.

[7] J. E. Boyd. Video phase-locked loops in gait recognition. In *International Conference on Computer Vision*, Vancouver, BC, Canada, July 2001.

[8] J. E. Boyd and J. J. Little. Phase in model-free perception of gait. In *IEEE Workshop on Human Motion*, Austin, TX, December 2000.

[9] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *Computer Vision and Pattern Recognition*, Santa Barbara, June 1998.

[10] R. Cutler and L. Davis. Robust periodic motion and motion symmetry detection. In *Computer Vision and Pattern Recognition 2000*, 2000.

[11] H. Fujiyoshi and A. J. Lipton. Real-time human motion analysis by image skeletonization. In *DARPA Image Understanding Workshop*, Monterey, California, November 1998.

[12] Y.-N. Hsu, H. H. Arsenault, and G. April. Rotation-invariant digital pattern recognition using circular harmonic expansion. *Applied Optics*, 21(22):4012–4015, 1982.

[13] E. A. Hunter, P. H. Kelly, and R. C. Jain. Estimation of articulated motion using kinematically constrained mixture densities. In *Nonrigid and Articulated Motion Workshop*, San Juan, Peurto Rico, June 1997.

[14] J. T. Kent. New directions in shape analysis. In K. V. Mardia, editor, *The art of statistical science: a tribute to G. S. Watson*, pages 115–127. Wiley, Chichester, 1992.

[15] J. Laszlo, M. van de Panne, and E. Fiume. Limit cycle control and its application to the animation of balancing and walking. In *SIGGRAPH 96*, pages 155–162, 1996.

[16] J. J. Little and J. E. Boyd. Recognizing people by their gait: the shape of motion. *Videre*, 1(2):1–32, 1998.

[17] F. Liu and R. W. Picard. Finding periodicity in space and time. In *International Conference on Computer Vision*, 1998.

[18] K. V. Mardia and P. E. Jupp. *Directional statistics*. Wiley, Chichester, 2000.

[19] MetaCreations Corporation. *Poser 4 user guide*, 1999.

[20] M. P. Murray. Gait as a total pattern of movement. *American Journal of Physical Medicine*, 16(1):290–332, 1967.

[21] R. Polana and R. Nelson. Detection and recognition of periodic, nonrigid motion. *International Journal of Computer Vision*, 23(3):261–282, June 1997.

[22] H. A. Rowley and J. M. Rehg. Analyzing articulated motion using expectation-maximization. In *Computer Vision and Pattern Recognition 97*, San Juan, Peurto Rico, June 1997.

[23] R. J. Smith. *Circuits, Devices, and Systems*. John Wiley and Sons, Inc., New York, 3 edition, 1976.

[24] M. Unuma, K. Anjyo, and R. Takeuchi. Fourier principles for emotion-based human figure animation. In *SIGGRAPH 95*, pages 91–96, 1995.

[25] S. Wachter and H.-H. Nagel. Tracking of persons in monocular image sequences. In *Nonrigid and Articulated Motion Workshop*, San Juan, Peurto Rico, June 1997.

[26] C. Wren, A. Azarbayenjani, T. Darrell, and A. P. Pentland. Pfinder: real-time tracking of the human body. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):780–785, July 1997.