Chapter 13

A Dynamical Model of the Coupling between Posture and Gait

Bruce A. Kay and William H. Warren, Jr.

Abstract

Walking requires the coordination of several tasks, including making forward progress (by producing a gait) and maintaining upright posture. In this chapter we discuss how a dynamical model of this coordination can be constructed. Experimentally, subjects walked on a treadmill while viewing large-field visual displays of an oscillating hallway. In our modeling, we focus on two features of the data, mode-locking between posture and gait (in which N cycles of postural sway were produced for M gait cycles) and the amplitude response (as a function of frequency) of the postural component. We discuss how to characterize dynamically posture and gait and their coupling, detailing a coupled-oscillator model of these and other coordination features. Although the model does not capture several other features of posture-gait coordination, we discuss how such a modeling effort can inform theories of how stable walking behavior is produced.

13.1 Introduction

Walking is one of the most ubiquitous activities performed by humans. It is an act that requires the timing and sequencing of many components for a walker to effectively navigate through the environment. Although there are many ways to analyze this complex action, a good starting point for looking at this coordination is to break it down into its component tasks. First, the walker makes forward progress through the world, or locomotes. Second, while doing so, the walker must maintain balance using an upright posture. Given that a biped's base of support is constantly changing and the center of mass is never directly above the stance foot, this involves a complex balancing act. Third, the walker must navigate through the normally cluttered natural environment, steering through openings such as doors and around obstacles such as parked cars—some of which may be moving. Finally, actors often perform other activities while walking (e.g., thinking, talking, gum-chewing), which implies that walking is a stable act, going along by itself if not greatly perturbed. That is, its essential components are robustly sequenced and timed.

In our laboratory, we have been looking at how three of the main component tasks of walking—locomoting, maintaining balance, and steering

—are coordinated. In this chapter, we describe the nature of the coordination between two of these components, posture and locomotion, as revealed by changes in postural sway and gait when probed with visual stimulation. We also describe a dynamical systems model of this behavior, going into the details of the modeling endeavor and showing that such modeling can make definite contributions to a theory of how humans walk. Our working hypothesis is that each component of walking can be understood as a component dynamical system having it own properties, and that the behavior of the walking system as a whole can be understood as the complex interaction—or coupling—of the component dynamics. In short, our first goal is to characterize each component's dynamics, and the second is to understand how the components' dynamics interact to produce the behavior of the entire walking system. Our overall aim in this modeling effort is to produce something "sufficiently definite" that can lead to constructive insights about this task.

Before going on, we should clarify what we mean by dynamical systems and dynamical models. A dynamical system is a system whose state at any instant of time can be characterized, at least in principle, by a set of scalar observables (Thompson and Stewart, 1986). The evolutionary history of the system is given by the time-series of these observables. Any system that evolves in time can be considered a dynamical system. The importance of a dynamical systems analysis for a theory of how a system works comes in the detailed specification of the dynamics of the system. The two main questions asked in a dynamical systems analysis are (1) what is the set of scalar observables that characterize the system, and (2) what are the dynamical laws that give rise to the patterns seen in the observables' time-series? The latter take the form of relationships among the observables, that is, expressions of how the observables interact. Having answers to these two questions can take us a long way toward understanding how a system operates.

For a simple system, such as a planar pendulum, we can identify the relevant observables ahead of time, for example, its angular displacement from the straight-down position and its angular velocity. The dynamical laws can be written down, based on a physical analysis of the forces present in the system, and any structural parameters present (e.g., the length of the pendulum and the force of gravity). These laws often take the form of differential equations, because in many situations they can be stated as functions that relate the rates of change of the various observables. Dynamical systems theory is used to describe the evolution of the system's state over time from a variety of starting conditions, both qualitatively and quantitatively, given the applicable laws. For a pendulum, the theory describes how the bob will move given its initial displacement and velocity (the starting conditions) and the parameterized pendulum equation (the law for this situation), in terms of both the qualitative and the exact time-series behavior of the observables (see, e.g., Thompson and Stewart, 1986).

In principle, walking can be considered a dynamical system because there is presumably a set of time-varying observables that can completely characterize its state. A detailed biomechanical-neural list of all the observables and the laws expressing their evolution over time could be worked out, but the sheer complexity of such a list is mind-boggling (e.g., Hatze, 1980). It is conceivable that from such a set of laws specific predictions could be made about global features such as the coordination of posture and locomotion, but the road to those predictions appears to us to be extremely long and arduous. Very little is known about the behavior of dynamical systems having so many observables (e.g., Guckenheimer and Holmes, 1983). In any attempt to understand such a system, then, we must simplify the system and choose a much smaller set of observables to try to understand. We do not know ahead of time which of the many possible observables are crucial for understanding the behavior we are interested in, so the choices we make are arbitrary, but we must start somewhere. Having chosen a small subset of the entire observable set, we must also make hypotheses about what the laws of interaction for the small subset are, because we can no longer rely on detailed analysis of how the observables in the full set interact.

Therefore, our approach, which is complementary to a detailed biomechanical-neural one, is to choose for study some simple observables that reflect global walking behavior, such as the sway associated with a point on the body or the angle between two leg segments. Recording the time-series of these observables, we look for patterns in the time-series, and try to reproduce these patterns with a candidate dynamical model.

Dynamical systems modeling is dynamical systems theory in reverse—an attempt to determine what kinds of laws relating the chosen observables could have produced similar behavior. From generic properties of known dynamical systems, we choose candidate laws relating our observables, then analyze and simulate these systems to evaluate the parallel between the candidate model and the observed behavior. If there are close parallels, such that the model produces time-series similar to the observed behavior, we say that we have a good model, and proceed to generate more predictions that can be tested experimentally. We can also say that the operation of the model—how the observables are functionally related in the laws—must be similar to the actual system in some manner, and if we understand the model, we have thereby achieved some understanding of the system. If there are distinct dissimilarities between the model and the real data, we keep trying, and hopefully will have learned something, at least what kind of dynamics our behavior is not. In dynamical system

modeling, understanding of a complex task like walking is proportional to how well the dynamical models mimic the observed behavior, and how well we understand the way the laws in the model produce the mimesis.

Again, our hypothesis is that we can do such modeling at the level of the components of the walking task, with separate models for the separate components, and then synthesize the components into a larger model for walking as a whole, choosing the appropriate dynamical laws that couple the components. Before proceeding from the components to the aggregate, a component model must pass the test of similarity to its component's real behavior. Let us now try to characterize in dynamical terms the component systems of posture and gait.

13.2 *The Postural System*

The task of the postural system during both quiet stance and walking is to maintain an upright balance. While standing in place, this task is to maintain the body's center of gravity (cg) over the base of support, which can be thought of as an area on the ground bordered by the feet. The phase space for posture—the space of all the observables of the postural state—can be simplified to the fore-aft/left-right coordinates of the cg with respect to the base of support. For postural stability, only a restricted region of this space is thus allowed. If at any time the cg is above a point outside this area, balance will be lost if no corrective measures (e.g., moving the feet, throwing one's arms out) are available to the actor. Thus, the cg must be kept within a restricted region of postural phase space, and if perturbed away from this region, the postural system's task is to bring it back using various effectors, based primarily upon visual and somatosensory information.

In the language of dynamical systems theory, this restricted region in postural space can be described as an *attractor*, although of a somewhat peculiar kind. The standard dynamical attractor that is its closest analog is the *point attractor*—that is, a preferred single state (consisting of the position, velocity,... of the system), which is returned to following arbitrary perturbations (figure 13.1). The postural attractor, on the other hand, is not a single point in phase space, but a whole set of points in a restricted region of the phase space. The salient feature of the postural task from the dynamical systems perspective is that the observables return to a restricted region of phase space following perturbation. Within that subset of phase space, the motion of the cg over the base of support may be periodic, chaotic, or stochastic (Chow and Collins, 1995), as long as upright posture is maintained.

As a first simplification, then, we can try modeling the postural attractor as a point attractor. A simple dynamical system that exhibits a point

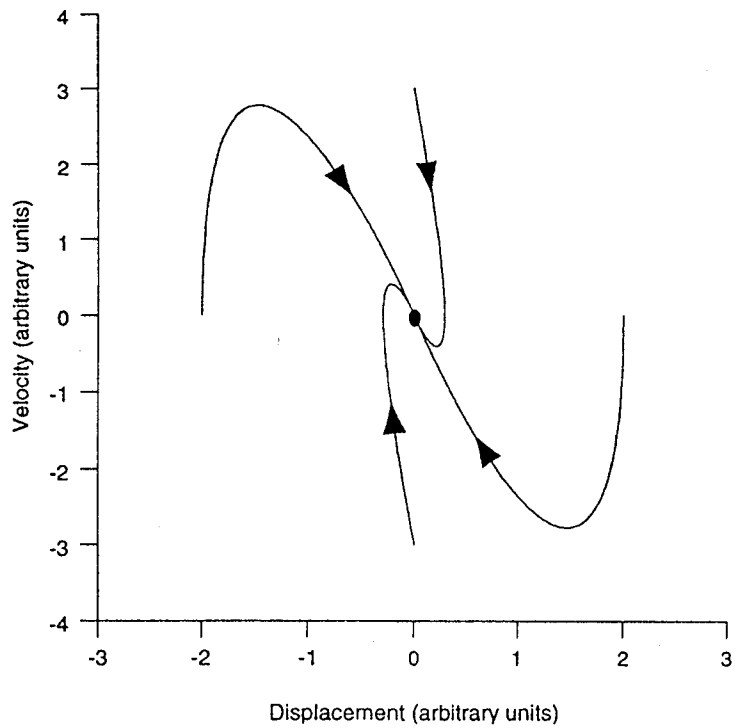

Figure 13.1
Point-attractor dynamics on the phase plane (position versus velocity). Following perturbation, all trajectories are attracted back to the point-attractor. The damped mass-spring used in the model is used here as the example system.

attractor is the following ordinary differential equation, which describes a damped linear mass-spring system:

$$m_P x'' + bx' + k_P x = 0. \tag{13.1}$$

This is an ordinary second-order differential equation that states the law relating the time-varying behaviors of the system's observables. The observables are the displacement x of the mass from its rest or equilibrium position (i.e., when $x = 0$, the spring is neither stretched nor compressed) at any time, and its velocity and acceleration x' and x'' respectively, using the prime notation to denote differentiation with respect to the final observable, time. Time is implicitly present in the time-derivatives of displacement, and is the independent variable in the equation, the other observables being dynamic variables, that is, dependent on time. The equation has three parameters that do not change over time, the mass m_P,

the amount of damping b, and the stiffness of the spring k_P (the subscript "P" denotes "Posture").[1] This equation states that the three observables are linearly related to each other, and given this relationship, the mass exhibits point attractor behavior: from any initial starting point (some definite values of x and x' at some start time), or equivalently, after being perturbed away from the equilibrium position, the mass will return to the equilibrium point after some transient motion (figure 13.1).

The behavior of the linear damped mass-spring system is well understood (French, 1971; Thomson, 1981). In particular, it is well known what will happen if a sinusoidal external force is applied to the mass. This additional force is added to the right side of equation 13.1:

$$m_P x'' + bx' + k_P x = F\cos(2\pi f_D t) \tag{13.2}$$

and introduces an explicit relationship between time and the rest of the system's observables. In this equation, F is the amplitude and f_D is the frequency (in Hz, or cycles per second) of the forcing function (the subscript "D" denotes Driver). After an initial transient, the mass oscillates at the same frequency as the external driver, with a fixed amplitude and phase relationship with respect to the driver. Thus the mass exhibits 1:1 mode-locking with the driver, that is, one cycle of forcing produces one cycle of response.[2] The amplitude and phase of the response, which are functions of all the parameters of the equation, is termed the *frequency response* of the system (Thomson, 1981). If an unknown mass-spring system is forced with an oscillating driver at various frequencies (f_D), the damping and stiffness coefficients b and k_P can be recovered from the observed amplitude-frequency and phase-frequency plots (given a known mass). In particular, for light damping, a peak occurs in the amplitude response near the natural frequency of the mass-spring (determined by the ratio of stiffness to mass). The mass-spring is said to "resonate" to the driver at this frequency, and so this is called a resonance peak (figure 13.2, solid curve). For heavy damping the amplitude response is a monotically decreasing function of the driver frequency (figure 13.2, dashed curve). This is termed "low-pass" behavior, as the mass-spring-damper acts as a filter that allows only the lower driver frequencies to "pass through" in the mass's motion relatively unattenuated (Oppenheim and Schafer, 1975). The amount of damping is measured relative to the values of the other parameters in the equation; for higher values of mass or stiffness, higher values of b are required for the same relative amount of damping to be present.

Applying this methodology to standing posture, several researchers (e.g., Andersen and Dyre, 1989; van Asten, Gielen, and van der Gon, 1988; Yoneda and Tokumasu, 1986) have found that when the postural system is visually driven by a sinusoidally oscillating display over a range

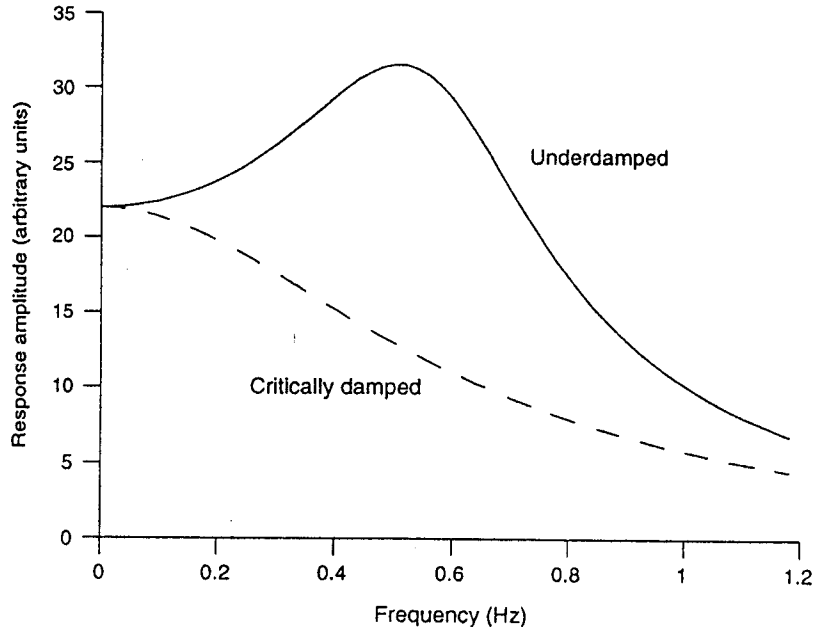

Figure 13.2
Amplitude response functions for an underdamped (solid line) and critically damped (dashed line) linear mass-spring system of figure 13.1.

of frequencies, it exhibits low-pass behavior. That is, the body sways at the same frequency as the visual driver (and so exhibits 1:1 mode-locking with the visual driver), and the amplitude of its response drops monotonically and very rapidly as the frequency of the visual driver is increased. Both of these facts are consistent with the dynamical system of equation 13.2. There is no resonance peak in the amplitude response function; thus the damping coefficient must be large compared with the other coefficients.

Whereas the results just presented pertain to standing, the postural response to visual oscillation during walking is quite different. In our experiments, participants walked on a treadmill, gazing at a large-screen display of a simulated 3-D hallway (subtended visual angles 110° × 95°, horizontal × vertical, figure 13.3), which oscillated sinusoidally in the lateral direction at 14 different frequencies (0.075 Hz to 1.025 Hz, amplitude = 34 cm peak-to-peak) so as to simulate side-to-side translation of the entire hallway.[3] This oscillation was superimposed on a basic radial flow component simulating forward progression down the hallway. Subjects were instructed to follow the side-to-side motions of the hallway so as to remain in the middle of the hallway, using body sway, side-to-side step-

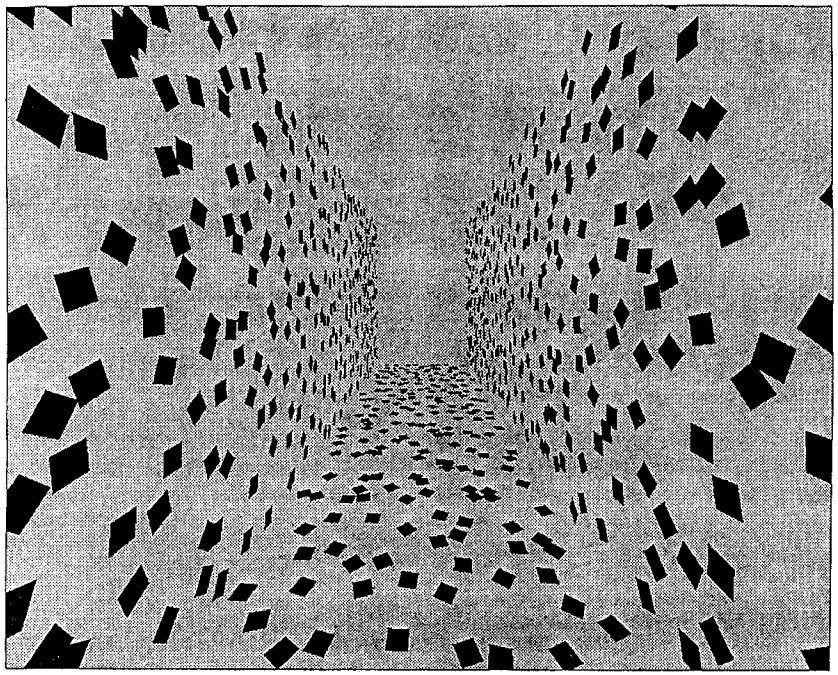

Figure 13.3
Static depiction of the visual hallway (black/white reversed due to printing limitations). The scene was presented on a large rear-projection screen 1 m from the subject.

ping, or whatever motion they felt was appropriate to perform the task. While the subjects walked (during four 40-sec trials at each of the driver frequencies), we recorded their lateral sway using an infrared motion analysis system to record the 3D positions of reflective markers on the side of the neck and the leg (figure 13.4). Four subjects swayed at the same frequency as the visual driver at all 14 frequencies; the other four did so only at the lower frequencies (below about 0.3 Hz). We report data of the former subjects here, although after discussing the model we will return to the latter subjects' behavior.

As in standing posture, we observed 1:1 mode-locking between sway and driver during walking, which is a characteristic of the forced linear system of equation 13.2 (figure 13.5). As can be seen in figure 13.6, however, the amplitude response for postural sway during walking not only has a low-pass characteristic like that for stance, but also a second peak around 0.8 Hz, which is close to the preferred stride frequency of 0.9 Hz observed on control trials with no display oscillation. This second peak

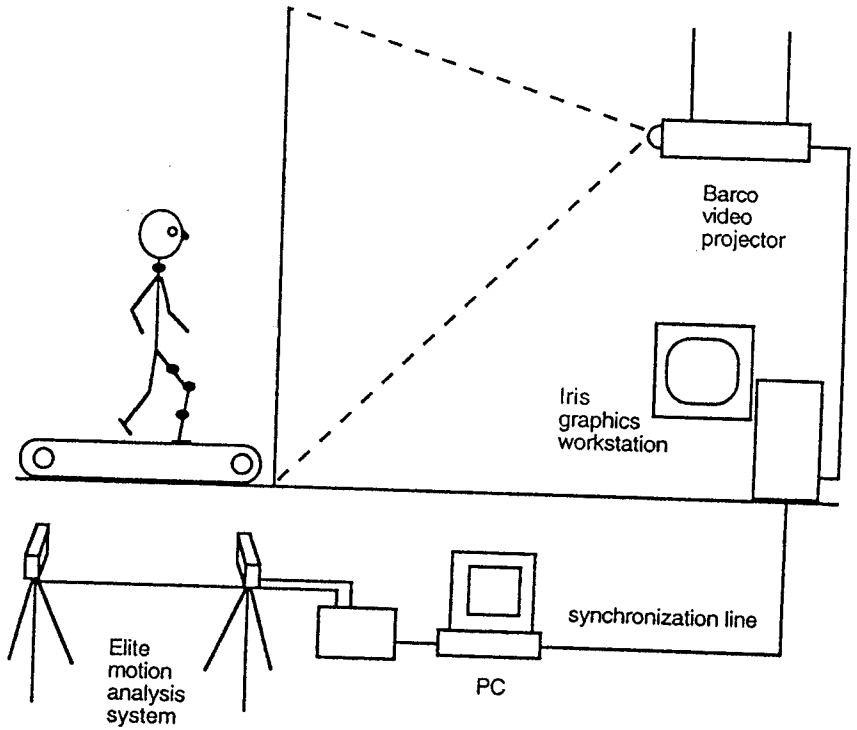

Figure 13.4
A schematic depiction of the experimental setup, showing the location of the passive reflective markers, whose 3D positions are recorded directly to disk by the Elite motion analysis system.

cannot be explained by equation 13.2, which predicts at most a single peak, with reduced sway amplitudes on either side. Also, the second peak is rather broad, such that the sway amplitude is affected at a range of frequencies, not just one.

So, it appears that gait has some influence on postural sway, particularly near the subject's preferred stride frequency. In our dynamical model, this influence takes a particular form of coupling between the two systems. Before turning to that issue, however, we must discuss how to characterize the second task of walking, locomotion, in dynamical terms.

13.3 The Locomotor System

In order to locomote over the ground, evolution has provided bipeds with a gait system in which two legs alternately swing back and forth, so

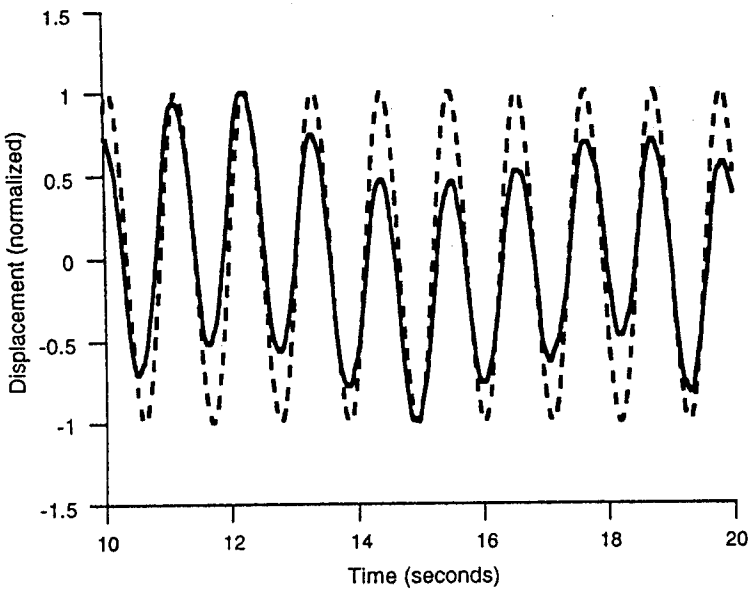

Figure 13.5
A 10-second example time-series of side-to-side sway when visually driven at a frequency of 0.925 Hz (solid line), and a time-shifted version of the sinusoidal driver (the best fit to the data time-series, dotted line). The displacements of the two time-series have been normalized to the interval $[-1, 1]$ for plotting purposes. The correlation between sway and driver was 0.921 for the entire 40-second trial.

that the feet are placed on the ground at the appropriate times and places. Dynamically, one of the most salient features of this behavior is the oscillatory component of gait. Walkers typically choose a preferred frequency and amplitude of leg motion, and if interrupted (by having to step over or around an obstacle, for example), they return to roughly these same two parameters of walking (Belanger and Patla, 1987). The closest dynamical analog is the *limit-cycle* attractor, an oscillation that has stable frequency, amplitude, and waveform in the face of perturbation. The reason this is called a limit-cycle can be shown by representing such a stable oscillation on the phase plane, a plot of the instantaneous position and velocity of the oscillating element. On the phase plane, there is one path that is stable for the limit-cycle, and this is a closed orbit. Following an arbitrary initial condition (again, x and x' at start time) or any arbitrary perturbation away from the limit-cycle, the system will evolve back to the limit-cycle during a transient process (figure 13.7).

The gait system's behavior is not quite a limit-cycle attractor. There is not a single closed cycle in any phase space representing the system's

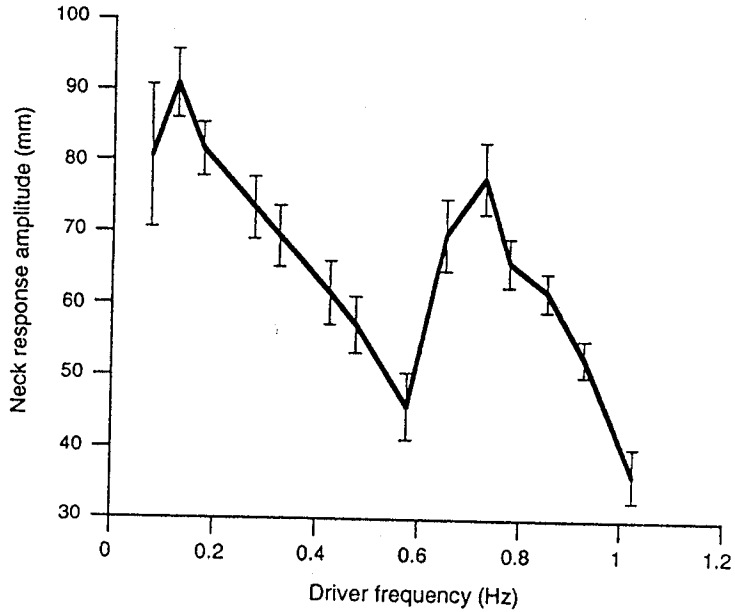

Figure 13.6
Amplitude response of body sway in response to visual oscillation during walking, for our successful walkers, with standard error bars at each driver frequency. Note the basic low-pass characteristic plus a broad peak around 0.8 Hz.

state, but rather a collection of such cycles into a band of attraction (Kay, 1988). Not exactly the same amplitude or frequency is used in each stride cycle, so the same states are not visited again and again, because gait is modulated slightly to maintain balance and adapt to environmental conditions. For now we will simplify and assume that we can model the salient oscillatory nature of gait as a limit-cycle attractor, for the same reasons that we simplified our model of the postural attractor.

A simple dynamical equation that exhibits a limit-cycle attractor is the van der Pol oscillator:

$$m_L y'' + \varepsilon(y^2 - 1)y' + k_L y = 0 \tag{13.3}$$

(Jordan and Smith, 1977; Nayfeh and Mook, 1979; Thompson and Stewart, 1986). Here, y represents the motion of the oscillating component away from the rest position ($y = 0$) and m_L is the component's mass (the subscript "L" denotes Locomotion). The van der Pol oscillator has a nonlinear damping term (the middle term), so the system observables are *nonlinearly* related. This nonlinearity allows the presence of a limit-cycle attractor. Roughly, it causes the oscillator's energy losses and gains to be

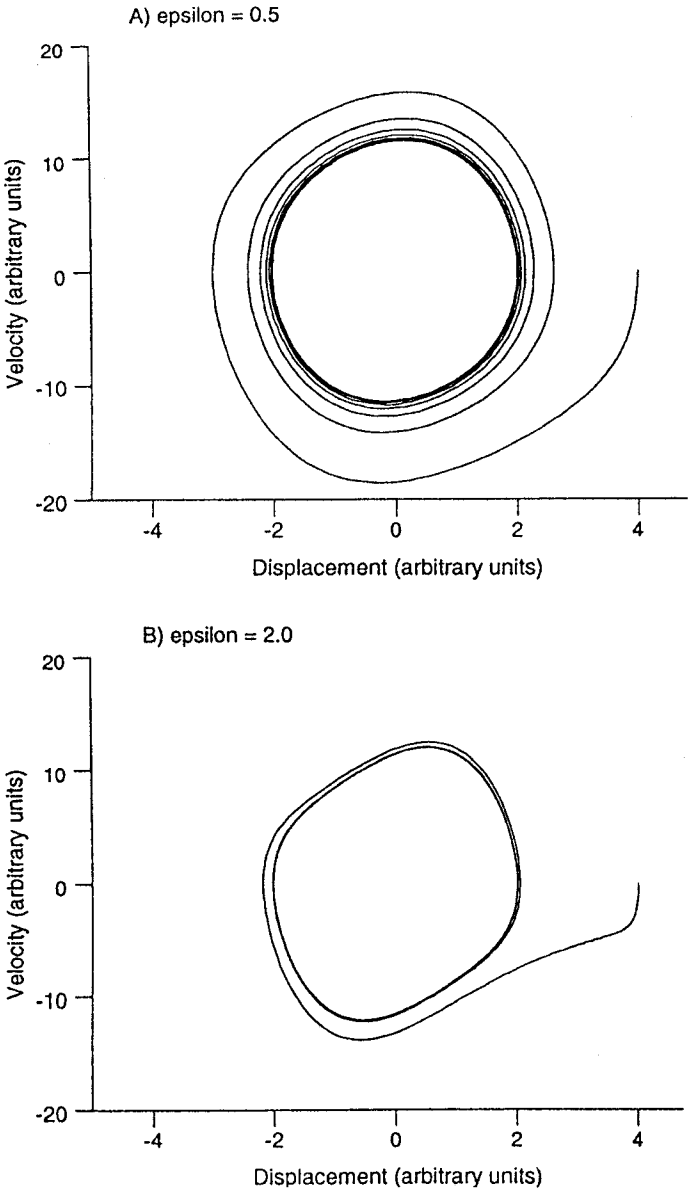

Figure 13.7
Phase-plane plots of the behavior of a van der Pol oscillator. *A*, with nonlinearity (ε) set to 0.5 and *B*, ε set to 2.0. Notice that it takes several cycles for the oscillation to settle down to its limit cycle in *A* following the start, whereas the settling time is much shorter in *B*. On the other hand, the oscillation is much more sinusoidal in *A* than in *B*, since the limit cycle is much more circular in the former.

exactly balanced on the limit cycle. The size of the parameter ε relative to the other parameters is a measure of how nonlinear it is; with larger ε's, return to the limit-cycle is faster than with smaller values, and the oscillations on the limit-cycle are less sinusoidal (figure 13.7).

We assume that the oscillatory gait component can be estimated by measuring any of the leg's motions that oscillate stably over time, such as the angle of the knee, which also happens to be rather easy to measure and analyze. In this case, we are not concerned about modeling either the amplitude of knee oscillation or how it changes over walking frequencies; in any event, this change is small (Diedrich and Warren, 1995). Because it is a salient feature of gait, however, we do want to incorporate a preferred frequency into our model. For fixed parameters, equation 13.3 has a fixed frequency of oscillation, which is approximately the square root of $(k_L/m)/(2\pi)$ for small ε. When not forced by some other term, the van der Pol will oscillate at this frequency regardless of the initial conditions or perturbations.

To anticipate, visually induced oscillations in the postural component appear to have an effect on the limit cycle of the locomotory component. To model this influence, we might try forcing the van der Pol oscillator (Nayfeh and Mook, 1979; Thompson and Stewart, 1986):

$$m_L y'' + \varepsilon(y^2 - 1)y' + k_L y = F\cos(2\pi f_D t). \qquad (13.4)$$

This dynamical system exhibits very complex behavior. Unlike the linear point attractor dynamic, the forced van der Pol will oscillate at the same frequency as the external driver only under a restricted set of circumstances. The stability of the van der Pol's intrinsic oscillation means that its observed frequency usually deviates only slightly from its preferred frequency when forced. Thus, the oscillator (M) does not usually exhibit 1:1 mode-locking with the external driver (N). Either N:M mode-locking, where N and M are (small) integers, occurs (such as 1:1, 2:3, 1:2, etc.), or the two frequencies are unrelated, or chaotic behavior ensues, all depending on the parameter values used in equation 13.4, especially the amplitude of forcing F and the frequency of forcing f_D (Hayashi, 1964; Jordan and Smith, 1977; Nayfeh and Mook, 1979; Thompson and Stewart, 1986). Holding F fixed and decreasing f_D, one observes a particular integer mode-locking for an interval of f_D values, followed by a usually larger interval of unrelated behavior between the two oscillations, followed by a smaller interval of mode-locking, and so forth.

In our experiments we observed both 1:1 and N:M mode-locking between postural sway and the stride cycle when our walkers were swaying 1:1 with the hallway. Because postural sway oscillated at the same frequency as the display on all trials, we computed the frequency ratio between the visual driver (N) and the stride cycle (M, as estimated by the

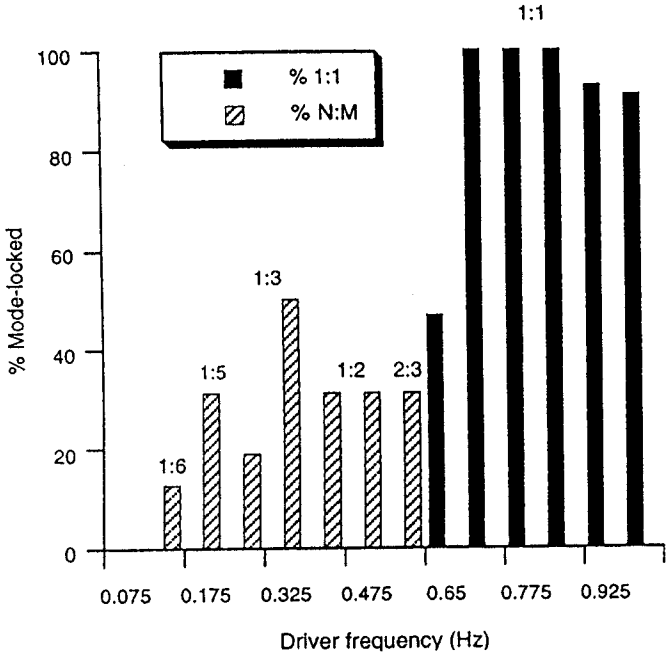

Figure 13.8
Percentages of trials that exhibited mode-locking between the visual display and the knee angle motion (out of 28 trials at each frequency), for the walkers who successfully swayed in synchrony with the visual display. Where no mode-locking occurred, the knees and display had incommensurate frequencies to the resolution of a trial. For visual clarity, the 1:1 mode-locks are indicated by solid bars, the others by striped bars.

knee angle), using several measures of inter-oscillator phase to confirm mode-locking for an entire 40 s trial. Figure 13.8 depicts the percentage of trials in which integer mode-locking was observed at each visual driver frequency. Note that 1:1 mode-locking—in which each cycle of visual oscillation was accompanied by one cycle of knee motion, over an entire trial of 40 seconds—occurred for all trials at or near the preferred stride frequency. The range of driver frequencies for which this mode is stable extends from 0.725 Hz to 1.025 Hz, the highest driver frequency studied. Other $N:M$ modes, such as 2:3 (i.e., two cycles of visual oscillation were accompanied by three cycles of knee motion), 1:2, and 1:3, occurred for lower visual driver frequencies, although less often, with many trials showing no stable integer mode-locking between posture and locomotion.[4] The ranges of driver frequencies over which these modes are stable are much smaller than the range for 1:1 mode-locking. Apparently, the locomotory component exhibits some of the same qualitative features as the van der Pol oscillator when sinusoidally forced.

13.4 Parameterizing the Component Dynamics

We have described two component dynamical systems to serve as candidate models for the postural and locomotor tasks. As described so far, we have a fairly generic form for each of them, with the *qualitative* properties of model and task being made to be in good correspondence. We would like to see if more *quantitative* properties of our observables can be mimicked by the models, which means that we would like to pick reasonable values of the parameters in the candidate laws (equations 13.2 and 13.4). After all, the forms of the models may be correct, but if there are systematic differences between the models and the data, we want to know about it and make corrections if possible. Only by choosing specific values for the parameters can we take this next step. We take each component in turn again.

The postural point attractor dynamic of equation 13.2 can be used to model the low-pass portion of the amplitude response curve we obtained (figure 13.6), ignoring the broad peak around the stride frequency. For simplicity, in the remaining we assume the mass of both the postural and locomotory components to be equal to 1.0 (see note 1). Assuming that the system is critically damped, that is, there is just enough damping to eliminate oscillations in any unforced transient and produce a monotonic amplitude response, the only remaining free parameter is the stiffness k_P. A value of $k_P = 14.21$ closely mimics the shape of the amplitude response, which implies a natural frequency of $f_0 = \text{sqrt}(k_P)/(2\pi) = 0.6$ Hz, and a damping of $b = 2m\omega_0 = (2)(1)(2\pi f_0) = 7.54$. To scale it to the actual amplitudes (in mm) we observed, the forcing term is $F = 312.61$.

For the gait component, there are only two parameters to set, stiffness k_L and nonlinearity ε (having set the mass to 1.0). We can determine k_L from the preferred stride frequency we observed in our subjects when they are not being forced by a visual oscillation. The average stride frequency was 0.9 Hz, which corresponds to a $k_L = 31.98 = (2\pi^*0.9)^2$. We chose ε to be 0.5, as this gives moderately fast return to the limit-cycle while still producing fairly sinusoidal oscillations, both of which make for a much simpler analysis and simulation effort.

13.5 Coupling the Postural and Locomotor Systems

We have now identified two separate dynamical models for two components of walking: a linear point attractor dynamic for posture, and a nonlinear limit-cycle attractor dynamic for gait. How do these two components interact? We made some initial hypotheses about what the coupling might be like, and then observed the total model's behavior via numerical simulation. Validation of the model comes in judging how well the simulation maps onto the observed behavior of our subjects.

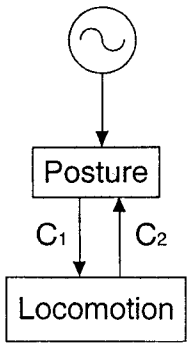

Figure 13.9
Schematic depiction of what is coupled to what in our walking situation, as we see it. C_1 and C_2 refer to the coupling terms of equation 13.5.

The first set of hypotheses is about the overall structure of the situation (figure 13.9). First, suppose that the visual driver has a direct influence on the postural system as in equation 13.2, such that the driver is coupled unidirectionally to the postural system, but postural sway does not reciprocally influence the visual driver. This was the situation in our experiment because we used open-loop displays, so the perspective in the display on the screen was not updated on the basis of the observer's head position, but this is distinctly not the case in natural situations, where the view of a hallway, for example, changes with every move of the head. Second, suppose that some form of coupling is present between posture and gait. Our data suggest that this coupling is bidirectional, for not only does gait affect postural sway, but also vice versa. The peak in the postural amplitude response near the preferred stride frequency indicates that posture is affected by gait. Conversely, the mode-locking behavior indicates that the visual driver acts to modulate the frequency of the stride cycle. Given the assumption that the visual driver is directly influencing only the postural component, gait must be influenced indirectly by the visual driver through the postural component.

Now, what is the exact form of these couplings? That is, given the parameterized system:

$$x'' + 7.54x' + 14.21x = 312.61\cos(2\pi f_D t) + C_1 \tag{13.5a}$$

$$y'' + 0.5(y^2 - 1)y' + 31.98y = C_2, \tag{13.5b}$$

what are the coupling functions C_1 and C_2? One of the simplest forms the coupling functions C_1 and C_2 could take is linear functions of the other component's observables; for example,

$$C_1 = a_1 y + b_1 y' \tag{13.6a}$$

$$C_2 = a_2 x + b_2 x', \tag{13.6b}$$

where the as and bs are constants.

Consider first the coupling function C_2, from posture to gait. Suppose coefficient a_2 is set to some nonzero value, with all the rest of the coefficients set to zero. With these settings, the postural component's displacement directly affects the motion of the gait component, but gait has no effect back on posture. In this case, the visual oscillation would drive the postural component at the driver frequency (but with an amplitude and phase depending on its frequency response), and the postural component would in turn drive the locomotor component sinusoidally. This would also be the case if coefficient b_2 were set to a nonzero value, so that the velocity of the postural component would be coupled to the gait component (with all other coefficients in equation 13.6 equal to zero). Either of these coupling functions might mimic the mode-locking behavior we observed, since a sinusoidally driven van der Pol exhibits mode-locking.

The sinusoidally forced van der Pol and similar oscillators exhibit exceedingly narrow mode-locking regions for the case where $N < M$, that is, when the driver frequency (N) is lower than the natural frequency of the van der Pol (M) (Hayashi, 1964; Nayfeh and Mook, 1979). This frequency relationship, which is the one we found in our experiment, is called *superharmonic entrainment*, because the response is at a higher frequency than the driver. In the forced van der Pol, much broader mode-locking regions exist for *subharmonic entrainment*, where the response is at a lower frequency than the driver. The superharmonic mode-locking regions for this way of forcing the oscillator are so narrow that it ought to have been very difficult for us to find them in our experiments. In other words, no mode-locking occurs at all for nearly any value of a_2 with driving frequencies below the natural frequency of the van der Pol, yet we observed some form of mode-locking at almost all driving frequencies we tested (figure 13.8). Putting it another way, this way of forcing the van der Pol oscillator is structurally unstable for superharmonic mode-locks, in that very slight variations in parameter values lead to very large changes in the behavior of the system (Thompson and Stewart, 1986). Consequently, this way of coupling from posture to gait doesn't predict the observed behavior.

Larger superharmonic mode-locking regions, observable over a wider range of frequencies, are produced by exciting the van der Pol in a very different way. Rather than directly forcing the van der Pol's observable state variables (y, y', y''), the driver varies one of its parameters, stiffness

k_L, which was previously held constant:

$$y'' + 0.5(y^2 - 1)y' + (31.98 + \beta x')y = 0. \tag{13.7}$$

Here, the coefficient of y is the sum of two terms, but can still be considered as the stiffness of the van der Pol. Thus, its stiffness is now continuously affected by the velocity x' of the postural component. The equation can be rewritten algebraically to show that $C_2 = -\beta x'y$, so another way to describe this is that this coupling function introduces a nonlinear relationship among the observables of the two components. This type of coupling—called *parametric excitation* (Cartmell, 1990; Nayfeh and Mook, 1979)—continuously modulates the original relationship among the van der Pol's observables themselves. In the case where there is no coupling back to the postural component, x' is a sinusoid and thus excites the van der Pol's stiffness sinusoidally. That is, the van der Pol's stiffness is now a function of time: $k'_P = k_P + \beta V \cos(2\pi f_D t)$, where V is the peak velocity of the postural component, which depends on the driver frequency according to the postural component's amplitude response.

In addition to the wider frequency ranges that exhibit superharmonic mode-locking, this type of mode-locking exists for much larger ranges of values of β (the amount of coupling from the postural component to the locomotory component) than it did for the linear coupling parameters a_2 and b_2 in equation 13.6b. The presence of wide regions of stable superharmonic mode-locking indicates that such mode-locking is now structurally stable. We have numerically computed these regions of stability in the two-dimensional parameter space of driver frequency and amplitude for the parametrically forced van der Pol in equation 13.7, and they appear in figure 13.10. The tongue-like shape of the stable 1:1 parameter region is apparent; that is, the range of frequencies that exhibit this mode-locking increases as the forcing amplitude is increased from the tongue's "tip" to its "blade." In this regard these regions are similar to the *Arnold tongues* observed in some closely related dynamical systems (named for the Russian mathematician V. I. Arnold; Thompson and Stewart, 1986). Note also that some minimum amount of forcing is required for mode-locking at any forcing frequency, below which the two oscillations are unrelated. In rough form, then, we have some idea of how to couple the postural component to the gait component.

Now consider the reciprocal coupling function C_1 back from gait to posture. What we want to account for is an increase in postural sway around the stride frequency. As it turns out, direct forcing of the postural state produces this effect, once the parametric excitation from the postural component to the gait component has been introduced. Specifically, we set b_1 in equation 13.6a to a nonzero value; the greater its value, the

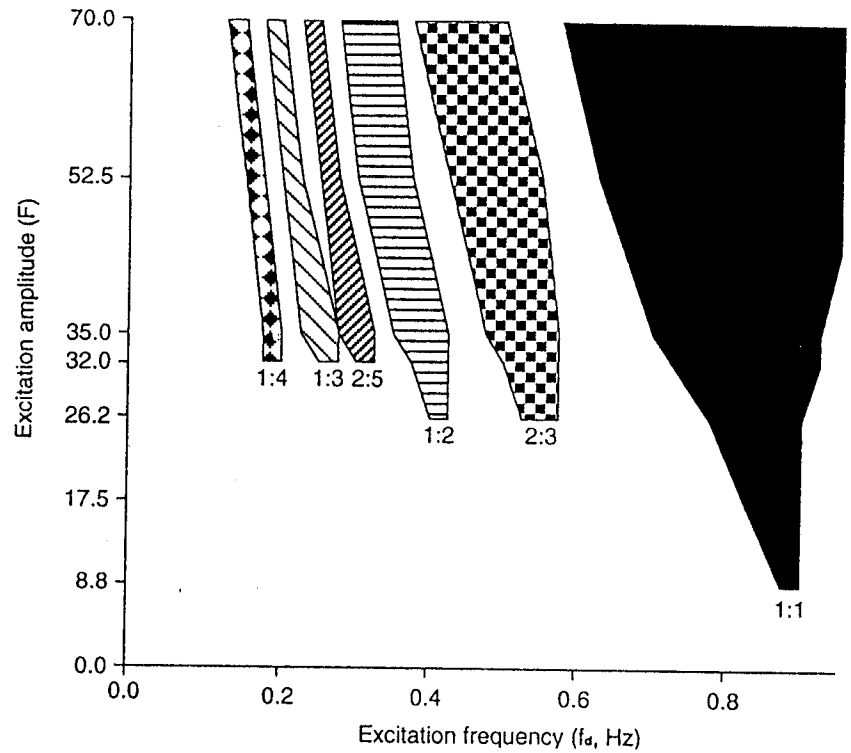

Figure 13.10
Mode-locking regions for the parametrically excited van der Pol oscillator, with the parameters of equation 13.7, except that the term bx' is replaced by $F\cos(2\pi f_d t)$, where F is the excitation amplitude and f_d is the excitation frequency. Inside the shaded regions of this parameter space, the variously labeled mode-locks ($N : M$ = #excitation cycles : #observed oscillator cycles) are stable and easily observed; outside, the relationship between the excitation and the van der Pol is quasiperiodic (N and M are not rationally related), that is, no mode-locking occurs there. This figure summarizes the results of simulations at 630 combinations of the two parameters, that is, at seven excitation amplitudes (see ordinate for the values) and 90 excitation frequencies (from 0.10 to 1.00 Hz in 0.01 Hz steps).

stronger the coupling from gait to posture and the larger the resonance peak.

Concerning exact values for β and b_1, we chose β so that the width of the 1:1 mode-locking regime was equal to what we observed in our data, that is, from 0.725 Hz to 0.925 Hz, when b_1 was set to zero; β turned out to be 1.0. Then, we increased b_1 until a large resonance peak appeared around the gait frequency; somewhat arbitrarily, we settled on $b_1 = 10.0$. We note that the 1:1 mode-locking regime maintained the same width, as did the other modes depicted in figure 13.10, when the C_1 coupling from gait to posture was added to the model.

13.6 The Complete Model and Its Successes

We now have a detailed, if preliminary, dynamical model of the interaction of posture and gait in our experimental setting:

$$x'' + 7.54x' + 14.21x = 312.61\cos(2\pi f_D t) + 10.0y' \qquad (13.8a)$$

$$y'' + 0.5(y^2 - 1)y' + (31.98 + 1.0x')y = 0 \qquad (13.8b)$$

Let us summarize how the terms in these equations correspond to the two components, and how the model operates. A sinusoidal visual driver directly forces the postural component, which has a linear point-attractor dynamic. In turn, the postural component parametrically excites the stiffness of the gait component, which has a nonlinear limit-cycle dynamic; reciprocally, the gait component directly forces the postural component's state. Thus, the two components are coupled bidirectionally but in two fundamentally different ways, using parametric coupling in one case and direct coupling in the other. Reflecting the open-loop conditions in our experiments, the visual display is only unidirectionally coupled to the walker's dynamics, although we will soon close the loop in our apparatus and extend the model to this case. To simulate the system's response to different frequency drivers, the only parameter that is changed is the driver frequency (f_d), while the structure of the equations and all the remaining parameters are held constant.

The model successfully reproduces key features of the behavior we observed in our subjects. The first two features were built into the model, and so we will review them briefly, but the remaining three features were discovered after the whole model was constructed.

First, the amplitude response of the postural component is very similar to that of the neck marker, with a low-pass region below about 0.6 Hz and a resonance peak around the preferred locomotor frequency (figure 13.11, solid curve). The low-pass feature is due to the intrinsic dynamics of the model's postural component (dashed curve), and the resonance peak

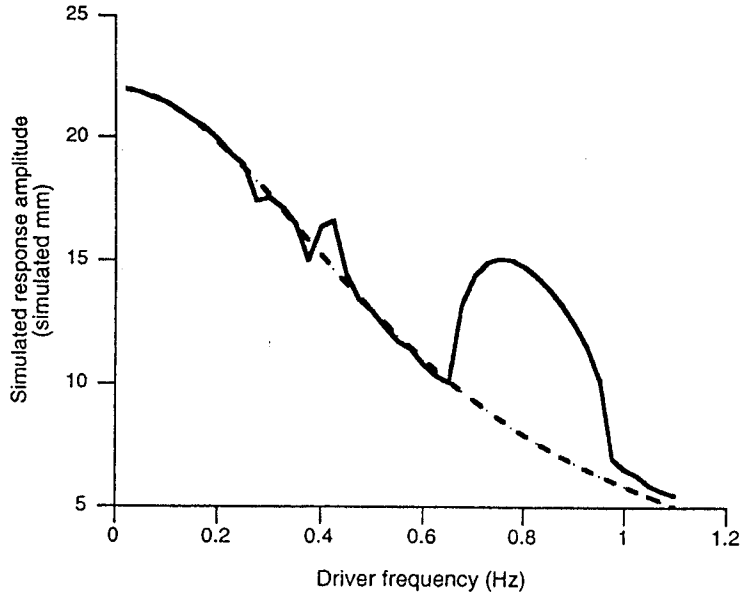

Figure 13.11
Amplitude response of the complete model (solid line) and the amplitude response of the forced linear mass-spring without the coupling from the locomotory component (dashed line).

is due to the combination of parametric coupling from the postural component to the gait component and the direct coupling in the opposite direction. Without both, no such resonance peak occurs. The model postural component also locks to the visual driver in the 1:1 mode across all driver frequencies used in the experiment. Second, there are a number of regions in the parameter space of the model where observable superharmonic mode-locking occurs between its two dynamical components (figure 13.10), similar to the observed walking behavior. In fact, the widths of the superharmonic regions decrease for the higher N:M modes, as does the percentage of mode-locked trials in our data (compare figure 13.8). In the model, the width of the 1:2 region is one-quarter that of the 1:1 region, which compares closely to the relative frequencies of occurrence of the two modes in walking. This feature of the model's behavior is due to the parametric coupling from the postural to the gait component.

Third, and unexpectedly, no stable mode-locks between the visual driver and the gait component occur in the model much above the locomotor frequency (0.9 Hz). That is, the model appears to have very narrow subharmonic entrainment regions (and so in a sense is the opposite of a

directly driven system). Nor was the model's postural component able to track the visual driver at these frequencies. Similarly, the postural sway in our subjects failed to lock to any driver faster than about 1.0 Hz. In the model, the gait component appears to dominate the behavior of the postural component at the higher frequencies.

Fourth, the model may be able to account for individual differences among our subjects. Recall that some subjects swayed 1:1 with the driver at all frequencies, whereas others only tracked the display at frequencies below about 0.3 Hz. With the preceding parameter settings, the model mimics the former subjects' amplitude response and mode-locking behavior rather closely. Interestingly, the unsuccessful subjects' behavior can be qualitatively mimicked by changing one parameter in the model—the coupling strength between the visual driver and the postural component, F. If F is reduced from 312.61 to 50.0, say, the postural response is dominated by the gait component at a much lower frequency, resulting in a 1:1 mode-locking region that extends out to about 0.3 Hz (figure 13.12).

Finally, when the visual driver is turned off, which corresponds to $F = 0.0$ in model, the gait component causes the postural component to oscillate at the stride frequency, as we observe in control trials that had no display oscillation.

13.7 *The Model's Failures*

The model is not a complete success, however. Simulations behave differently from the observed behavior when we turn from the rather gross features of mode-locking and amplitude resonance to finer time-series details. In particular, there are two important differences between the actual and modeled kinematics of the knee. First, the amplitude of knee oscillation remains rather constant from cycle to cycle in the real walkers, but does not in the model. For example, in the 1:2 mode the simulated knee amplitude varies systematically, being large on one cycle, small in the next, in a strictly alternating pattern (figure 13.13b, middle waveforms). Amplitude modulation of the model's "knee" is even more apparent when mode-locking is absent, whereas there is very little cycle-by-cycle knee variation in amplitude in the walking data (compare figure 13.13a and b, bottom waveforms). Second, the timing of the model's peak displacements at 1:2 has the same alternating pattern, first a long, then a short cycle in time, whereas the observed knee timings do not vary in this systematic manner. In real walking, it appears that the overall stride frequency has been altered to match the postural oscillations, but that individual cycles are not modified in order to do so. In the model, on the other hand, the continuous modulation of the gait stiffness parameter produces cycle-by-cycle changes in the oscillator's frequency. We shall return to this point later.

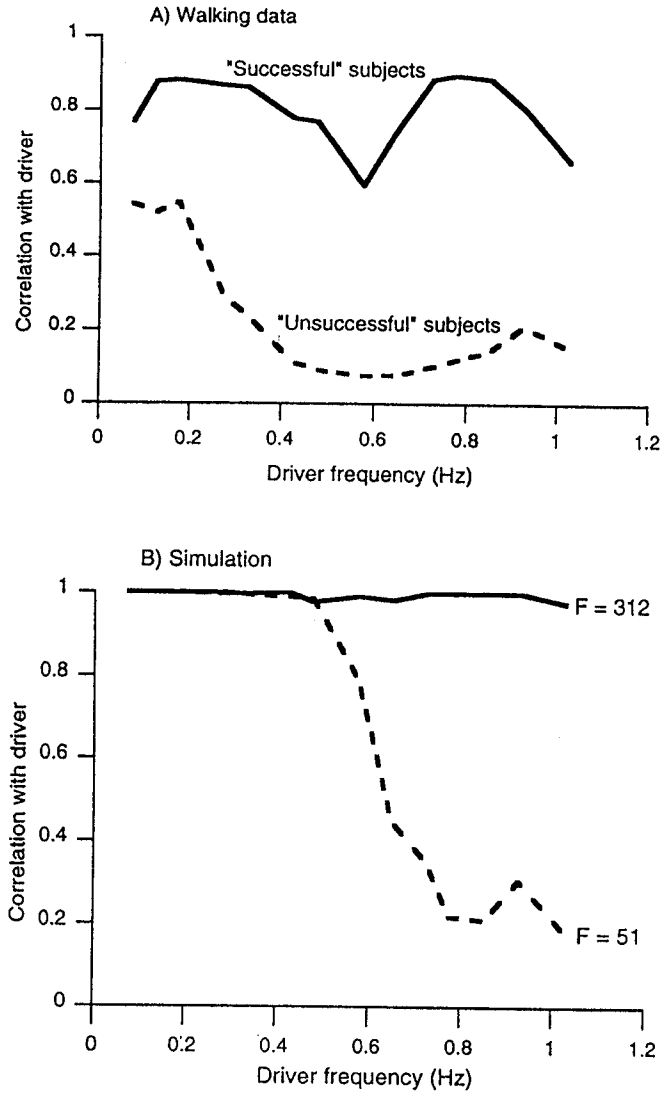

Figure 13.12
A) Mean correlation between the visual display and the neck marker motion for successful (solid line) and unsuccessful (dashed line) subjects. B) Correlation between the visual driver and the locomotory component of the model for $F = 312.61$ (solid line) and $F = 51.0$ (dashed line). For 1:1 mode-locking between the driver and locomotion, the correlation should be near 1.0.

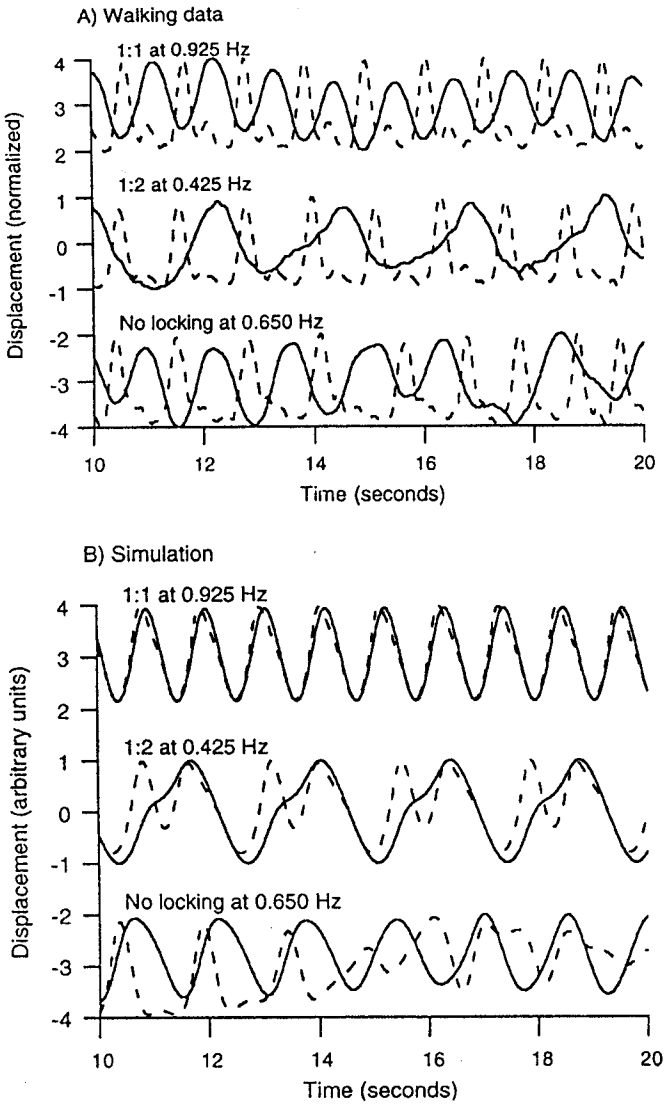

Figure 13.13
Ten-second example time-series of lateral sway motion (solid lines) and knee angle motion (dashed lines) in A) the walking data and B) numerical simulations of the model. From top to bottom in each sub-plot the visual driver was at 0.925, 0.425, and 0.650 Hz.

Besides the preceding kinematic details, the model is unsuccessful in duplicating an aspect of postural sway behavior we observed in another experiment, in which we used sum-of-sines displays. An important distinguishing feature of nonlinear systems is that they lack linear superposition, that is, the response to a sum of input sinusoids, for example, is not equal to the sum of the responses to the sinusoids when input individually. To test whether the postural component is nonlinear, we created a set of composite waveforms, each consisting of the sum of four sinusoids from the same set of fourteen frequencies that were used one at a time in the single-sine experiment discussed previously. In order to make the analysis easier, the four frequencies that were summed in any one trial were incommensurate to the base sampling rate of a trial. This produced waveforms with overall patterns that repeated only once every 40 seconds (the duration of a trial), and appeared very random and unpredictable to the walkers. We then measured the amplitude response of the neck sway at each of the four frequencies in each trial. Across the set of composite waveforms we thus had a measure of the response at each of the fourteen original frequencies, which we then compared to the response when each frequency was presented individually. In the sum-of-sines response, we found the same low-pass characteristic and the resonance peak around the preferred locomotor frequency as we obtained for single-sine displays, but in addition we found a second peak between the low-pass and locomotor peak, at around 0.5 Hz (figure 13.14, solid curve). Some type of nonlinearity in visually-induced postural sway is signaled by this difference.

On the other hand, our model is patently nonlinear, due to the van der Pol's damping term and the parametric excitation term, so we were hopeful that it might produce this result. However, when we forced the model (equation 13.8) with the same sum-of-sines waveforms as we used in the experiment (i.e., by replacing the right side of equation 13.8a with the sums of four cosine terms we used in the experiment), we obtained a fairly flat low-pass amplitude response (figure 13.14, dashed curve). Although there is a difference between the single- and sum-of-sines responses of the model, which reflects its nonlinearity, the form of the difference is not the same as what we observed in our walkers. One possibility is that the model's sum-of-sines response depends upon the precise frequency combinations used, and so slightly different frequency combinations may result in a frequency response that more closely mimics the form of the observed data. Even if that were true, it would imply that the frequency response function is not robust in the face of such changes, and so would not be a good model of this rather consistent finding, which was observed in a number of subjects.

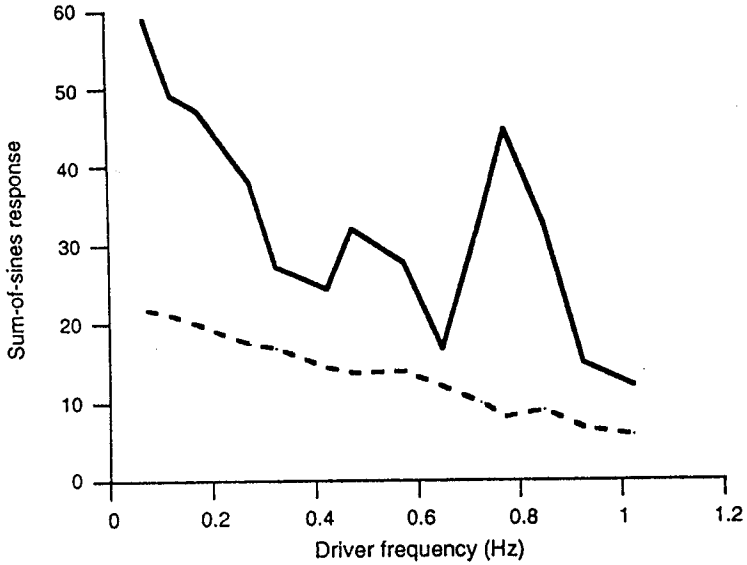

Figure 13.14
Amplitude response function for walking, when sum-of-sines waveforms were used in each trial (solid line), and the corresponding response of the dynamical model to sum-of-sines drivers (dashed line). Compare with figures 13.5 and 13.9.

13.8 *What Have We Learned?*

As we have shown, our model of the interaction between posture and gait is more successful in some respects and less so in others. It could be argued that some of the model's successes are due to our having built them into the model; some of its failures, as a consequence, are things we did not choose to include. But there are also some fundamental differences between the observed walking data and the model's behavior at the level of the cycle-by-cycle time series, and in the fundamental nature of its nonlinearity as revealed by the sum-of-sines displays. Consequently, the model must be modified or replaced to capture more accurately the interaction between posture and gait in walking.

Nevertheless, both the successes and the failures of the model inform us about the nature of that interaction, and allow us to draw some general conclusions.

A. Posture and Gait Are (Probably) Bidirectionally Coupled. This is not so trivial a point as it may at first appear. When we first looked at the knee data, we were puzzled to find many trials in which it appeared that the postural oscillations had no effect on the locomotor motions of the knee.

In fact, in only about 30% of the low-frequency trials (that is, well away from the preferred stride frequency) did we observe mode-locking. The remaining trials showed an independence of the two oscillations. Thus, our first idea was that at many postural frequencies, the two components act independently of each other for all intents and purposes. Even with fixed parameters, however, the model also shows independence at the level of its observed kinematics. Except in the mode-locking regions, the two components oscillate relatively independently of one another. The model's behavior thus also looks as though the two components are uncoupled at some driving frequencies, yet the coupling terms are always present, and the two components are therefore always influencing each other. The lesson to be learned, then, is that posture and gait may be bidirectionally *coupled* at all times, though the coupling may not produce tightly-coupled *behavior* in all circumstances.

B. The Coupling Between Posture and Gait Is in Terms Not Only of the Components' Observables but also in terms of the Components' Parameters. As we have shown, the observed mode-locking behavior can be modeled using parametric, and not direct, coupling, from the postural to the locomotor component. This tells us that the overall system involves not only interactions in observables, but also interactions at the level of parameters.

C. The Interaction of Posture and Gait may be at an Altogether Different Level of Dynamics than that Described in Our Model. Although the model can produce superharmonic mode-locking behavior similar to what we observed in our data, we noted above that the details of how that behavior is produced differ drastically between model and data, with small cycle-by-cycle variation in both amplitude and timing in the real knee behavior, but systematic cycle-by-cycle amplitude and timing variations in the modeled knee.

Apparently, mode-locking in actual walking is accomplished in a very different manner from how it is accomplished in the model. Parametric excitation—the basis of the coupling from the postural component to the locomotor component in the model—is a form of coupling that entails continuous variations in one component's parameters (gait stiffness). By virtue of that modulation, gait stiffness has effectively become another time-varying observable, evolving in time at the same rate as the other observables in the model (i.e., it evolves at exactly the same rate as the velocity of the postural component), rather than being a fixed parameter (Farmer, 1990; Saltzman and Munhall, 1992). At each moment in time, then, stiffness takes on a new value, thus changing the observed frequency of the oscillator.

Thus, this continuous variation of gait stiffness in the model says that this crucial locomotor parameter is adjusted continuously, which does not

seem to be a reasonable biological solution to the problem. It may make more sense for a walker to reset the value of his or her locomotor stiffness to a new value in order to mode-lock with the postural oscillations, and then stick with that new value for as long as the postural oscillations are required. This calls for a completely different coupling scheme, in which both the postural and locomotory components' parameters are adjusted by some other process (Dijkstra et al., 1994), rather than by virtue of couplings that are invariant across visual driver frequency, as in the current model. This new way of modulating the locomotor component has the virtue of retaining parameters as parameters, that is, gait stiffness would be re-set by the other process and left there, and so would be constant in time, not evolving at the same rate as the observables. This would also produce the observed phenomenon of fixed cycle-by-cycle knee frequency and amplitude.

The main drawback of the latter approach is that the interaction is taken to a very different dynamical level. It implies the presence of some other process that occurs at a time scale very different from that of the evolution of the two components' observables. Furthermore, it would no longer be the case that it is the dynamics of the interaction of these two components' observables that provides the mode-locking phenomena. It is hard to see how the limited mode-locking regions we observe in both the present data and model would be generated in such an approach. If the locomotor frequency could be adjusted to any frequency, we might expect that some stable form of mode-locking would be present for any driver frequency, because some small-integer frequency ratio could be chosen for any arbitrary driver frequency. The interaction between these two components could not then serve as an explanation of how the preferred mode-lockings arise. As stated, this approach does not have all the answers either.

13.9 *In Sum*

We have outlined a dynamical model of the interaction of posture and gait when the postural component of walking is required to follow a visual oscillation. The model has some strong points as well as some shortcomings. Despite the model's limitations, we have shown that such a modeling effort can lead to some interesting and definite questions about how two important components of a complex task interact. For example, we have tried to ask (and answer), how do we characterize the individual components dynamically? And how are they linked? In other task settings, such as a more natural closed-loop situation in which the subject's movement influences the visual display, we may be able to answer the question, how is vision coupled to the rest of the system? Such questions, we think,

are best answered when definite hypotheses of how the couplings work can be stated, and our choice is to state those hypotheses in the language of dynamical systems.

Acknowledgments

We thank David Rosenbaum and an anonymous reviewer for their very helpful comments on an earlier draft of this chapter. We take full responsibility for any remaining obfuscation. The work reported herein was supported by NIH Grant EY10923.

Notes

1. Because the type of analysis we are pursuing is a solely kinematic one, where motions and not forces are being analyzed, strictly speaking mass should not be in this equation or the following ones. In terms of a dynamical analysis, this is of little importance, since mass plays the role of just another parameter, a linear weight on the acceleration in the present equation. As we discuss later, we arbitrarily set mass to 1.0, but could be altered for modeling purposes if necessary.
2. We have chosen the term "mode-locking" as opposed to frequency- or phase-locking because of its generality. Both frequency- and phase-locking imply that the two oscillations are operating at the same frequency, so that the phase between them is constant. Although this does apply to the case of a 1:1 frequency ratio between driver and response in our experiments, it does not apply to the case of N:M frequency ratios where N and M are unequal.
3. Thus the amount of oscillation with respect to the retina varied as a function of location in the scene, with minimal motion at the center of the screen, that is, for "far" portions of the hallway, and maximal motion at the screen edges, where the hallway was "close" to the walker. For more details on the experimental apparatus, setup, and data analysis, see Warren, Kay, and Yilmaz, 1996.
4. Even chaos may be present in our walking data, but it is extremely difficult to demonstrate in systems as noisy as biological ones, so we will not try here.

References

Andersen, G. J., & Dyre, B. P. (1989). Spatial orientation from optic flow in the central visual field. *Perception & Psychophysics, 45,* 453–458.

Belanger, M., & Patla, A. E. (1987). Phase-dependent compensatory responses to perturbation applied during walking in humans. *Journal of Motor Behavior, 19,* 434–453.

Cartmell, M. (1990). *Introduction to linear, parametric and nonlinear vibrations.* London: Chapman and Hall.

Chow, C. C., & Collins, J. J. (1995). Pinned polymer model of posture control. *Physical Review E, 52,* 907–912.

Diedrich, F., & Warren, W. H. (1995). Why change gaits? Dynamics of the walk-run transition. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 183–202.

Dijkstra, T. M. H., Schöner, G., Giese, M. A., & Gielen, C. C. A. M. (1994). Frequency dependence of the action-perception cycle for postural control in a moving visual environment: relative phase dynamics. *Biological Cybernetics, 71,* 489–501.

Farmer, J. D. (1990). A Rosetta stone for connectionism. *Physica D, 42,* 153–187.

French, A. P. (1971). *Vibrations and Waves.* New York: Norton.

Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields.* New York: Springer-Verlag.

Hatze, H. (1980). Neuromusculoskeletal control systems modeling—a critical survey of recent developments. *IEEE Transactions on Automatic Control, AC-25,* 375–385.

Hayashi, C. (1964). *Nonlinear oscillations in physical systems.* New York: McGraw-Hill.

Jordan, D. W., & Smith, P. (1977). *Nonlinear ordinary differential equations.* Oxford, UK: Clarendon.

Kay, B. A. (1988). The dimensionality of movement trajectories and the degrees of freedom problem: A tutorial. *Human Movement Science, 7,* 343–364.

Nayfeh, A. H., & Mook, D. T. (1979). *Nonlinear Oscillations.* New York: Wiley-Interscience.

Oppenheim, A. V., & Schafer, R. W. (1975). *Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall.

Saltzman, E. L., & Munhall, K. G. (1992). Skill acquisition and development: The roles of state-, parameter-, and graph-dynamics. *Journal of Motor Behavior, 24,* 49–57.

Thompson, J. M. T., & Stewart, H. B. (1986). *Nonlinear dynamics and chaos: Geometrical methods for engineers and scientists.* New York: Wiley.

Thomson, W. T. (1981). *Theory of Vibrations with Applications, 2d ed.* Englewood Cliffs, NJ: Prentice-Hall.

van Asten, W. N. J. C., Gielen, C. C. A. M., & van der Gon, J. J. D. (1988). Postural movements induced by rotation of visual scenes. *Journal of the Optical Society of America, A, 5,* 1781–1789.

Warren, W. H., Kay, B. A., & Yilmaz, E. H. (1996). Visual control of posture during walking: Functional specificity. *Journal of Experiment Psychology: Human Perception and Performance, 22,* 818–838.

Yoneda, S., & Tokumasu, K. (1986). Frequency analysis of body sway in the upright posture: Statistical study in cases of peripheral vestibular disease. *Acta Otolaryngolica* (Stockholm), *102,* 87–92.