Journal of Experimental Psychology:
Human Perception and Performance
2000, Vol. 26, No. 4, 1281–1297

0096-1523/00/$5.00 DOI: 10.1037//0096-1523.26.4.1281

Functional Stabilization of Unstable Fixed Points: Human Pole Balancing Using Time-to-Balance Information

Patrick Foo, J. A. Scott Kelso, and Gonzalo C. de Guzman
Florida Atlantic University

Humans are often faced with tasks that require stabilizing inherently unstable situations. The authors explored the dynamics of human functional stabilization by having participants continually balance a pole until a minimum time criterion was reached. Conditions were manipulated with respect to geometry, mass, and characteristic "fall time" of the pole. Distributions of timing between pole and hand velocities showed strong action–perception coupling. When actions demonstrated a potential for catastrophic failure, the period of hand oscillation correlated well with the perceptual quantity "time to balance" ($\tau_{bal} \equiv \theta/\dot{\theta}$), but not other quantities such as θ and $\dot{\theta}$ alone. This suggests that participants were using available τ_{bal} information during critical conditions, although they may not have been attending to this type of perceptual information during typical, noncritical motions of successful performance. In a model analysis and simulation, the authors showed how discrete τ_{bal} information may be used to adjust the parameters of a controller to perform this task.

Biologically significant activities such as the maintenance of posture (e.g., Jeka & Lackner, 1994, 1995), the development of posture and locomotion (e.g., Lee & Aronson, 1974; Thelen, 1990), and the learning of new motor skills (Zanone & Kelso, 1992, 1997) may be viewed as involving active stabilization of inherently unstable fixed points of a dynamical system. In each of these cases, relevant perceptual information is used by the participant to stabilize the unstable system (Kelso, 1998). In quiet standing posture, light fingertip contact with a touch bar that is too weak to provide physical support can reduce the mean sway amplitude or entrain the motion of the body if the touch bar oscillates (Jeka & Lackner, 1994, 1995; Jeka, Schöner, Dijkstra, Ribeiro, & Lackner, 1997). In the moving room paradigm, postural compensations to changes in optical flow (stemming from the motion of the room) can be demonstrated across a range of ages and motor developmental stages (Bertenthal & Bai, 1989; Bertenthal, Rose, & Bai, 1997; Lee & Aronson, 1974). Information about the relative phase between rhythmically moving limbs may be used to stabilize previously unstable, to-be-learned patterns of coordination (Zanone & Kelso, 1992, 1997). The present work explores functional stabilization through the model task of humans balancing an inverted pendulum along a linear track (Treffner & Kelso, 1995, 1997, 1999). Unlike the foregoing posture and locomotion paradigms where a sensory modality drives or entrains an action system, in pole balancing the system perturbs itself: At each instant, the participant's own actions directly influence the perceptual information that guides action. To date, detailed kinematic studies of human pole balancing have been largely absent (see, however, Treffner & Kelso, 1995), although a similar balancing task has been used as an interference task during studies of hemispheric cerebral function (see Kinsbourne & Hicks, 1978).

The study of stabilization of inherently unstable systems has received much attention within the domain of control theory. Here, the textbook example of balancing an inverted pendulum has played a key role in elucidating basic control concepts as well as motivating various control design techniques (e.g., see Kwakernaak & Sivan, 1972). The inverted pendulum configuration in our experiment (see Figure 1A) is based on the often-used cart–pole design by Barto, Sutton, and Anderson (1983). As in most theoretical considerations, we focus on the nature of the control force F and bypass the details of its delivery. Linear control, in which F is a linear function of the state–space variables, has been widely used to balance an inverted pendulum successfully (see Geva & Sitte, 1993, for a review). Here state–space variables are defined as those quantities such as positions and velocities used to describe the state of the system. We distinguish these from parameters, which are those quantities that are externally determined and typically used to specify the overall strength of the control signal given the form of F. For example, if the force F is linear with respect to the pole angle ($F = k\theta$), then k is a parameter, whereas θ is a state–space variable. By linearizing the equations for the motion of the cartpole system, it is possible to extract the parameter range that results in successful balancing in the region of small pole angles.

Neuromorphic linear controllers whose parameters (sometimes called *weights* in the neural network literature) are determined using artificial neural networks provide not only effective solutions but also possible insights into how biological systems may accomplish the balancing task (e.g. Anderson, 1989). Using actual movement data (the state–space variables) from balancing experiments, neural network controllers can learn how to mimic humans

Patrick Foo, J. A. Scott Kelso, and Gonzalo C. de Guzman, Center for Complex Systems and Brain Sciences, Florida Atlantic University.

This research was supported by National Science Foundation Grant SBR 9511360, National Institute of Mental Health (NIMH) Grants MH42900 and K05MH01386, and NIMH Training Grant MH19116. We thank Betty Tuller for her critical discussions of the manuscript.

Correspondence concerning this article should be addressed to Patrick Foo, Center for Complex Systems and Brain Sciences, Florida Atlantic University, 3848 FAU Blvd., Innovation II, Boca Raton, Florida 33431-0991. Electronic mail may be sent to foo@walt.ccs.fau.edu.

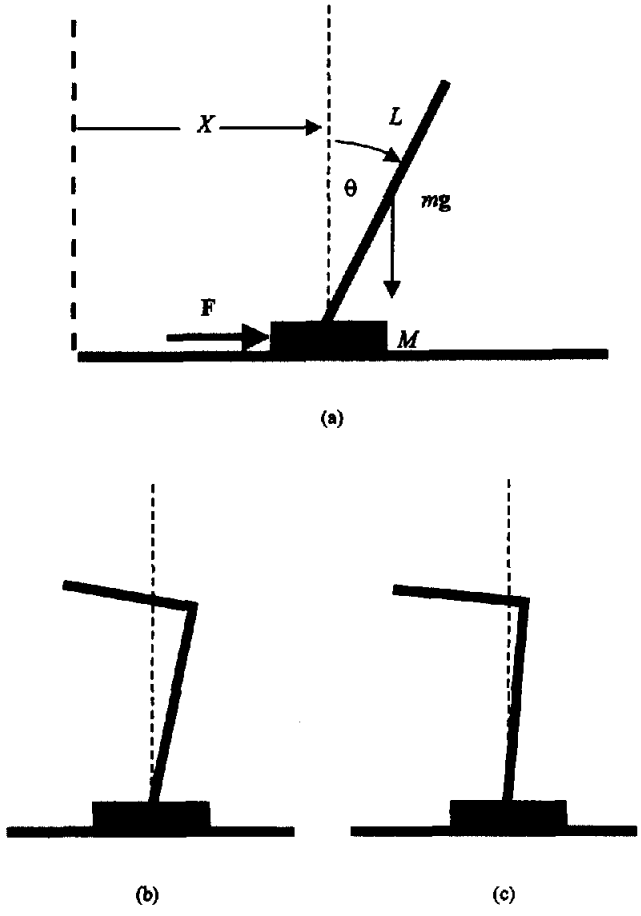

Figure 1. The experimental setup consisting of a pole configuration, a cart, and a horizontal linear track. A: Straight pole configuration and the relevant physical parameters and variables used in the analysis. The convention used is as follows: θ is positive (negative) when the pole is right (left) of the vertical; X = cart position; M = mass of the cart, L = pole length, m = pole mass, F = force applied by the hand to the base; g = acceleration due to gravity. B: Steel L-pole configuration shown at its balance point (the base rod inclined 11.4° from the vertical). C: Wood L-pole configuration at its balance point (the base rod inclined 4° from the vertical). For the L-poles, the angles were adjusted by subtracting the appropriate balance point value.

accurately (Guez & Selinsky, 1988). Along the same line, but using time-varying digitized images of the cart–pole system instead as inputs, Tolat and Widrow (1988) constructed a pattern-recognizing system that mastered the balancing task well enough to keep the pole from falling.

A central problem in neural network design in which no teacher is available (unsupervised learning) is how to adjust future control actions on the basis of the outcome of a sequence of actions. This is known as the credit assignment problem and is usually addressed by imposing a suitable constraint such as maximizing the balance time or minimizing excursions from the vertical. As noted emphatically by Geva and Sitte (1993), a random search in weight space can produce outcomes that in many respects are not significantly different from those resulting from a more systematic search. In the present article, we explore the idea that in addition to the kinematic variables (e.g., hand position and velocity), human controllers also use direct perceptual information to solve the balancing task. Clearly, a key question is which perceptual variables the human controller uses.

In typical situations, humans rely on both haptic and visual information, the latter being very much affected by the relative motion of the human–pole system. This article examines the role of a perceptual quantity, time to balance (τ_{bal}), defined as the ratio of the angle of the pole with the vertical (θ in Figure 1) to its rate of change:

$$\tau_{bal} = \theta / \dot{\theta} . \quad (1)$$

When τ_{bal} is negative, its magnitude is an approximate measure of the time to upright the pole. Only in the case when the angular velocity is constant is τ_{bal} the actual time to upright. In the present work, we show that participants are more attuned to the temporal behavior of the pole and produce more consistent values of τ_{bal} and $\dot{\tau}_{bal}$ when there is a potential for failure compared with when the pole is near balanced conditions. These behaviors suggest that participants are using available τ_{bal} information, especially near critical conditions. During typical (noncritical) motions of successful performance, wider variation in these quantities is observed, suggesting that participants may not be attending to this type of perceptual information. In a model analysis, we show how such information may govern the dynamics of perception–action coupling in this task.

The main goals of the present research are to identify (a) biologically relevant perceptual variables in successful balancing and (b) the nature of the perception–action coupling that supports functional stabilization. Once these variables are identified experimentally, we may then formulate a procedure for adjusting the parameters of the control system. Although this is not exactly a formulation of a control law, it follows the spirit of Warren (e.g., 1988, 1998) in that we aim to specify how perceptually available information constrains action and vice versa.

Method

Participants

Thirty-six right-handed students between the ages 17–30 served as participants (18 men, 18 women). All participants completed an informed consent form prior to data collection and received course credit for participating. All participants were treated in accordance with the ethical standards of the American Psychological Association (APA).

Apparatus

The apparatus consisted of a pole configuration attached via a metal bearing to a cart (mass = 388.0 g) such that the pole only pivoted in the xy plane (see Figure 1). The cart was constrained to slide along a 180-cm linear track oriented parallel to the x-axis. The track rested securely on a table approximately waist-level height for adult participants (92 cm). Three different poles of varying length, mass, and configuration (one straight pole, two L-shaped poles) were used. The L-poles were designed so that the center of mass did not coincide with any portion of the pole: Here the simplistic strategy of aligning the pole with vertical (which stabilizes the straight pole) was ineffective.

In the straight-pole condition (see Figure 1A), the pole was made from a single aluminum dowel (length = 108.0 cm, diameter = 1.0 cm, mass = 207.4 g, and moment of inertia I = 0.07 kg/m^2). In the L-pole conditions,

the configurations had identical appearances but different material compositions, moments of inertia, total masses, and balance points. The steel L-pole (see Figure 1B) had the same moment of inertia as the straight pole (I = 0.07 kg/m^2) but a greater total mass (279.0 g). This pole was constructed by attaching a steel rod (length = 29.4 cm, mass = 163.8 g) perpendicular to and on top of an aluminum rod (length = 60.0 cm, mass = 115.2 g) acting as the base. The 90° bend extended to either the left or right side of the participant. When this pole was balanced so that the center of mass lay directly on top of the pivot point, the base aluminum rod was displaced approximately 11.4° from vertical. The wood L-pole (see Figure 1C) was constructed similarly, but with the steel portion replaced by a 15.9 g wooden rod of the same length. The total mass and moment of inertia of the wood L-pole configuration were 131.1 g and 0.02 kg/m^2, respectively. When balanced, the base aluminum rod was displaced approximately 4.0° from vertical. For the three pole systems, the surfaces of the rods were covered to mask their textures and material compositions.

To monitor movements of the hand and pole, five infrared-emitting diodes (IREDs) were affixed to strategic locations. IREDs 1 and 2 were attached to the base of the track, and the adjoining line defined the horizontal reference. IRED 3, placed on the pivot point of the cart, was used to monitor the hand displacements that were restricted to horizontal movements. IREDs 4 and 5 were placed on the pole in such a way as to permit measurement of its angular orientation (in either the straight or L-pole conditions). Signals from the five IREDs were sampled at 100 Hz using an OPTOTRAK 3010 system and were later filtered with a low-pass (set at 8 Hz) second-order Butterworth filter. The basic kinematic data thus consisted of horizontal hand position (x) and the angle (θ) from the pole's balance point.

Procedure

The task was to balance the pole by moving the cart with the right hand along the linear track (the x-axis). Haptic information from the cart, but not directly from the pole itself, was thus available to the participant. The goal was to balance the pole for 30 s without allowing it to fall and make contact with the track. Failure to complete 30 s of continuous balancing constituted an unsuccessful trial. Obviously participants received visual feedback about success or failure of their performance on each trial. Each participant performed blocks of 10 experimental trials followed by a 1-min rest period. Moreover, during the rest period, participants were given knowledge of results in the form of total time of balance for each trial in the previous block. In addition to interblock rests, participants were provided 5-min rest periods at Trials 40 and 70.

Design

Participants were placed via matched (by sex) random assignment into one of six experimental groups (n = 6 for each group: 3 men, 3 women) and asked to perform all three testing sessions (acquisition, transfer, and retention; see Table 1) until they could reach a performance criterion. Participants were required to perform a given balancing task until they could complete three successful attempts out of five consecutive trials for each of the three testing sessions.

Each group was assigned a criterion task during the acquisition session that was to balance successfully one of the three pole configurations: straight, L-steel, and L-wood. On completing the acquisition session, the participant attempted the transfer task, which was to balance a different pole. One week later, the participant was asked to balance the original pole in a retention test.

If a participant did not successfully balance the first pole after 99 trials and also did not successfully complete the transfer task after 50 trials, his or her data were excluded from the final data analysis. For each participant who reached criterion, the number of trials executed to meet the criterion level of performance (Nc) was noted.

Table 1
Experimental Design

Group	Day 1: Acquisition, Session 1	Day 1: Transfer, Session 2	Day 8: Retention, Session 3
1	Straight pole	Steel L-pole	Straight pole
2	Straight pole	Wood L-pole	Straight pole
3	Steel L-pole	Straight pole	Steel L-pole
4	Steel L-pole	Wood L-pole	Steel L-pole
5	Wood L-pole	Straight pole	Wood L-pole
6	Wood L-pole	Steel L-pole	Wood L-pole

Results and Discussion

Did participants significantly improve their balancing performance by the end of the experiment for any of the poles? A mixed 6 (groups) × 3 (testing sessions) analysis of variance (ANOVA) with group as a between-subjects factor and testing session as a within-subjects factor was performed on the dependent variable (Nc) to assess differences in performance between acquisition and retention tests. A significant interaction, $F(10, 60) = 3.59$, $p < .01$, revealed that practice led to a significant improvement in performance for the four L-pole groups (Groups 3–6, see Table 1) between acquisition and retention. However, excellent performance of the straight-pole groups (Groups 1–2) in acquisition limited the amount of improvement available on retention testing, and no significant change in performance was seen.

Which pole was easiest to balance during the acquisition session? A Tukey post hoc analysis revealed that during acquisition, the straight pole was significantly easier to balance than the steel L-pole and the wood L-pole ($p < .05$). The significant difference in performance between the straight pole and the steel L-pole (I = 0.07 kg/m^2, for both poles) suggests that the moment of inertia manipulation was not effective. The similar performance of the steel and wood L-poles indicated that the pole configuration manipulation (compared with the straight pole) depressed acquisition performance.

Did practice with different poles during acquisition improve performance (proactive transfer) during the transfer session? A Tukey post hoc analysis of the previous 6 (groups) × 3 (sessions) mixed ANOVA revealed that participants who practiced with either L-pole in acquisition (e.g., Group 4, which balanced the steel L-pole in acquisition and transferred to the wood L-pole) showed a significant performance improvement in the transfer session when compared with participants balancing the same (transfer) pole in acquisition (and did not receive any previous practice; e.g., Groups 5 and 6). Participants who balanced the easier straight pole in acquisition (e.g., Group 2) did not receive the benefit of transfer, and their performance was comparable with participants who did not have any previous practice at all (e.g., Groups 5 and 6 in acquisition). Additionally, no performance improvement was seen in the straight-pole transfer session groups (e.g., Groups 3 and 5) because of the exceptional performance of participants who balanced the straight pole in the acquisition session (e.g., Groups 1 and 2). Thus, when improvement was possible, practice with the more difficult balancing poles during acquisition ameliorated performance in the transfer session. These results differ from those of Bachman (1961), who found no evidence for transfer of learning between two gross balancing tasks.

It should be noted that Bachman's study was designed to discover a generalized motor ability, and the tasks were designed to be distinctive in execution while retaining the common element of balancing.

Kinematic Analysis: Time Series of Unsuccessful and Successful Balancing

Our principal focus is on the quantities that characterize successful balancing behavior, both in acquisition and across differences in the geometrical and physical properties of the balanced object. Data for the hand velocity ($\dot{x}$) and pole angular velocity ($\dot{\theta}$) were obtained by numerically differentiating the time series data for the hand position (x) and the pole angle (θ). Figure 2 shows representative time series of the hand position (solid lines) and pole angle (stippled lines) for one participant balancing a straight pole (see Figure 2, A and B) and one participant balancing an L-pole (see Figure 2, C and D). The plots on the left side show balancing behavior early in acquisition; plots on the right side show successful performance late in retention testing. The straight-pole participant was able to develop and maintain successful balancing in both acquisition and retention. Note that the peaks (valleys) of the hand position coincide with valleys (peaks) of the pole angle, indicating an antiphase coordination between the hand and the pole.

In contrast, early in the acquisition session, the participant using the wood L-pole (see Figure 2C) fails to develop this antiphase pattern, and the pole moves to extreme angles followed by failure. Later (see Figure 2D) this participant successfully balances the pole during the retention phase, and the antiphase coordination between hand and pole is evident.

All trials showing continuous pole balancing lasting longer than 3 s were included in the final data set. This set included 339 trials meeting the original 30-s criterion as well as 1,935 additional trials. For each trial, all cycles of contiguous balancing were included except the last cycle that preceded a failure. Each cycle was defined as one half-period of the continuous hand velocity time series. Thus 76,637 cycles of 2,274 individual trials were included in the analysis of continuous balancing. Note that even for unsuccessful trials, successful cycles of balancing were included in our analysis.

Kinematic Analysis: Coupling Between Hand and Pole

To explore coordinated behavior beyond that induced by the mechanical coupling of the pole to the cart, we analyzed the timing differences between the hand and pole actions for each cycle (half-period) of the hand velocity trajectory. As noted in Figure 2, hand position and angle tend to be coordinated when the pole remains upright. However, further consideration suggests that the velocity variables might provide more sensitive evidence of coupling behavior (e.g., Jeka & Lackner, 1994, 1995; Kelso et al., 1998). A cross-correlation was performed between the time series of the hand and pole velocities for each cycle of the hand velocity. The value of Δ was defined as the time difference between the peak of the normalized cross-correlation and a zero-lag value (as in

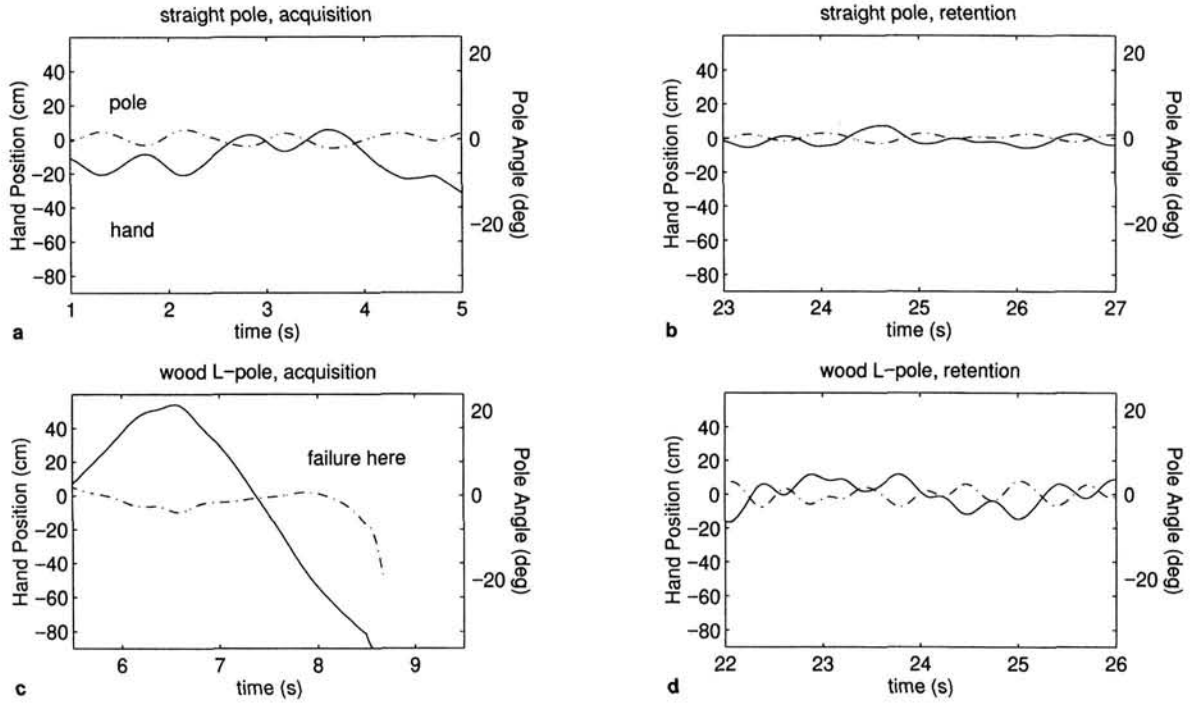

Figure 2. Hand position (solid) and pole angle (stippled) time series for successful and unsuccessful balancing trials. Shown are timeseries of typical balancing behaviors during acquisition and retention sessions using the straight pole (A, B) and the wood L-pole (C, D). In both cases, hand position is antiphase with the pole angle. As seen in Panels C and D, although early balancing using the wood L-pole is unsuccessful, after practice the kinematic behavior takes on the characteristics of the successful trials of the straight pole.

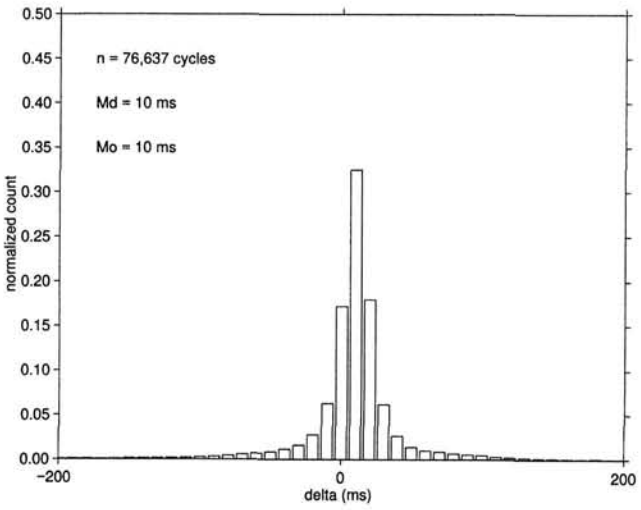

Figure 3. Distribution of Δ values collapsed across learning and pole conditions. Md = median value; Mo = mode.

simultaneous hand and pole velocity peaks). Over the course of successful balancing, participants developed and maintained a high degree of coupling between the hand and the pole as evidenced by an average Δ value of 10 ms (hand lagging the pole) across all learning and pole conditions (see Figure 3). Of course, this value, though impressive, is limited by the sampling rate (± one sample).

Kinematics of τ_{bal}

When balancing a pole under the current experimental conditions, information about the pole's state is available primarily through vision. Although, in principle, both the angle and angular velocity are available as measures of this information, they may not be the principal sources through which a human controller regulates action. In the following, we present the kinematic behavior of the hand–pole system in terms of the time to balance, τ_{bal}. Recall that τ_{bal} is defined as the ratio of θ over $\dot{\theta}$ (see Equation 1). Then, using a model analysis, we argue that knowledge of τ_{bal} and the hand position (x) is sufficient to implement a more direct control system for balancing the pole.

To familiarize the reader with the typical properties of τ_{bal} and how it behaves in terms of conventional position and velocity measures, we show in Figure 4 three time series for the successful balancing data of Figure 2D. We examine the behavior of these perceptual variables at the hand velocity extrema ($\ddot{x} = 0$). The rationale for selecting the hand velocity extrema is that they reflect most accurately the onset of interceptive actions such as when the hand starts to reverse its direction of motion or, equivalently, when it starts to recover from a previous action (see also Treffner & Kelso, 1995; Wagner, 1982). Although previous balancing data has shown that the perceptual τ variable demonstrates the smallest coefficient of variation during successful performance at approximately 170 ms prior to the onset of hand deceleration (Treffner & Kelso, 1995), in the current experiment we explore the possibility that sufficient perceptual information may be gleaned by monitoring τ_{bal} and $\dot{\tau}$ at the onset of hand deceleration, because these variables appear to be conserved near these time points in the balancing cycle. Note that around the times of hand velocity extrema ($\ddot{x} = 0$; see solid lines of Figure 4A), the value of τ_{bal} is conserved near $\tau_{bal} = 0$ (see Figure 4B); furthermore, at the corresponding time, the value of $\dot{\tau}_{bal}$ is conserved near $\dot{\tau}_{bal} = 1$ (see Figure 4C). In other words, during the onset of deceleration of

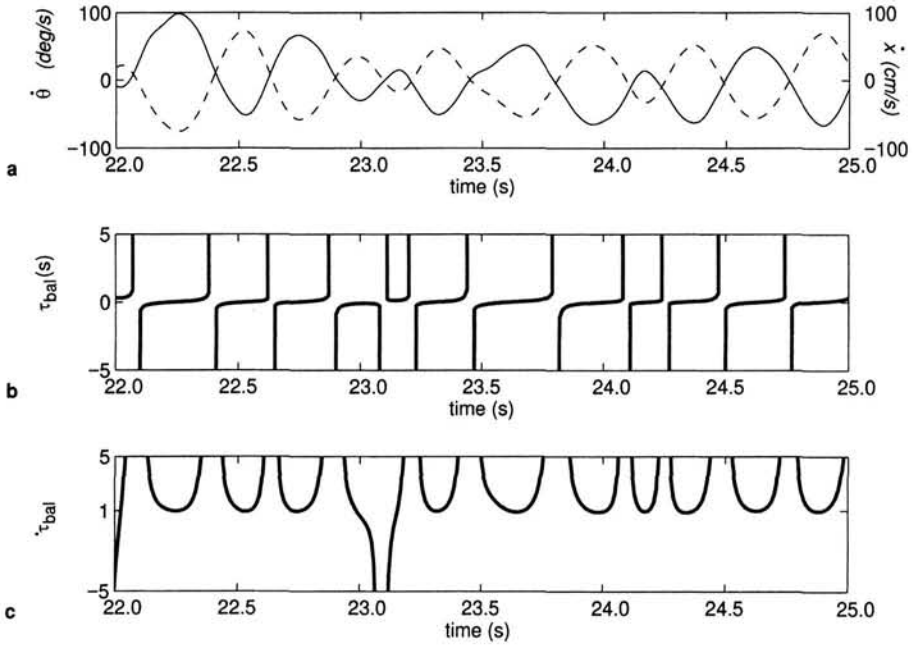

Figure 4. A: Time series plots of the hand velocity (solid lines) and pole velocity (stippled lines) showing tight antiphase coordination. Data are from the successful trial shown in Figure 2D. B and C: Plots of τ_{bal} (time to balance) and $\dot{\tau}_{bal}$ (derivative of τ_{bal} with respect to time) computed from the same data source as in Panel A.

the hand, the pole is moving near the vertical ($\tau_{bal} = 0$) and is typically overshooting the vertical ($\dot{\tau}_{bal} = 1$).

Unlike the continuous and smooth variations of the $\dot{x}$ and $\dot{\theta}$ plots, τ_{bal} exhibits characteristic singularities in an almost regular manner. The divergences occur when $\dot{\theta}$ reaches zero and usually correspond to pole reversals (see Figure 4B). The behavior between reversals is schematically illustrated in Figure 5. To aide in understanding the figure, we note first that because of the symmetry of θ and $\dot{\theta}$ with respect to left–right exchange, τ_{bal} has corresponding sign symmetry. Regardless of which side of vertical the pole is on, $\tau_{bal} > 0$ means that the pole is moving away from the vertical, whereas $\tau_{bal} < 0$ means the pole is moving toward the vertical.

Consider now Figure 5A, which shows a typical movement consisting of the pole crossing the vertical as it goes from one side to the other. This means that τ_{bal} changes from negative (moving toward the vertical) to zero (exactly at the balance point) to positive (moving away from the vertical) values, traversing a sigmoid path. If, during a restoring motion, the pole undershoots the vertical (see Figure 5B), then τ_{bal} is always negative, but its absolute value initially decreases, then diverges. The result is an inverted U-shaped curve. After the undershoot, the pole reverses direction and drifts away from the vertical ($\tau_{bal} > 0$). As the pole falls, its angular velocity increases and τ_{bal} decreases. Subsequent deceleration of the pole then increases τ_{bal} again to ∞. This produces the U-shaped plot in Figure 5C. When participants fail to recover the pole, the positive τ_{bal} value increases without return until failure: This corresponds to an incomplete U-curve (see Figures 5D and 2C). Note that because of the symmetry mentioned earlier, the same scenarios occur irrespective of which side of the vertical the pole resides on. In Figure 6, we show the phase portrait of (τ_{bal}, $\dot{\tau}_{bal}$) for a section of the time series data in Figure 4.

Except for the negative $\dot{\tau}_{bal}$ region, the shape of the curve suggests quadratic dependence of $\dot{\tau}_{bal}$ with τ_{bal} but with variable curvature at the origin. Note that the instances of pole crossover (which correspond to the sigmoid trajectory in Figures 4B and 5A) describe a curve like the two quadratic curves in the inner part of the phase portrait of Figure 6. Instances of pole undershoots, corresponding to inverted U-curves in Figures 4B and 5B, are represented by the two trajectories in Figure 6 where the τ_{bal} values remain negative. Conversely, drifts of the pole away from vertical (see U-shaped curves of Figures 4B and 5C) describe the two trajectories in the phase portrait where τ_{bal} remains positive. Note that $\dot{\tau}_{bal}$ values cluster around 1 for small values of τ_{bal}, indicating that, as noted previously, during the onset of deceleration of the hand the pole is moving near the vertical ($\tau_{bal} = 0$) and is typically overshooting the vertical ($\dot{\tau}_{bal} = 1$). Thus, one expects the variability of $\dot{\tau}_{bal}$ to be minimal at $\tau_{bal} = 0$. Does this prediction hold true in general? An analysis was performed on all successful cycles of balancing motion, across pole and learning conditions to determine the values of τ_{bal} and $\dot{\tau}_{bal}$ at the onset of hand deceleration (see Figures 7 and 8). Participants produced an average τ_{bal} of 0, and a $\dot{\tau}_{bal}$ of 1 across all successful balancing cycles, regardless of pole or learning conditions. The robust conservation of these perceptually based variables across manipulations of pole and learning conditions suggests that τ_{bal} may provide the informational support for successful completion of this task.

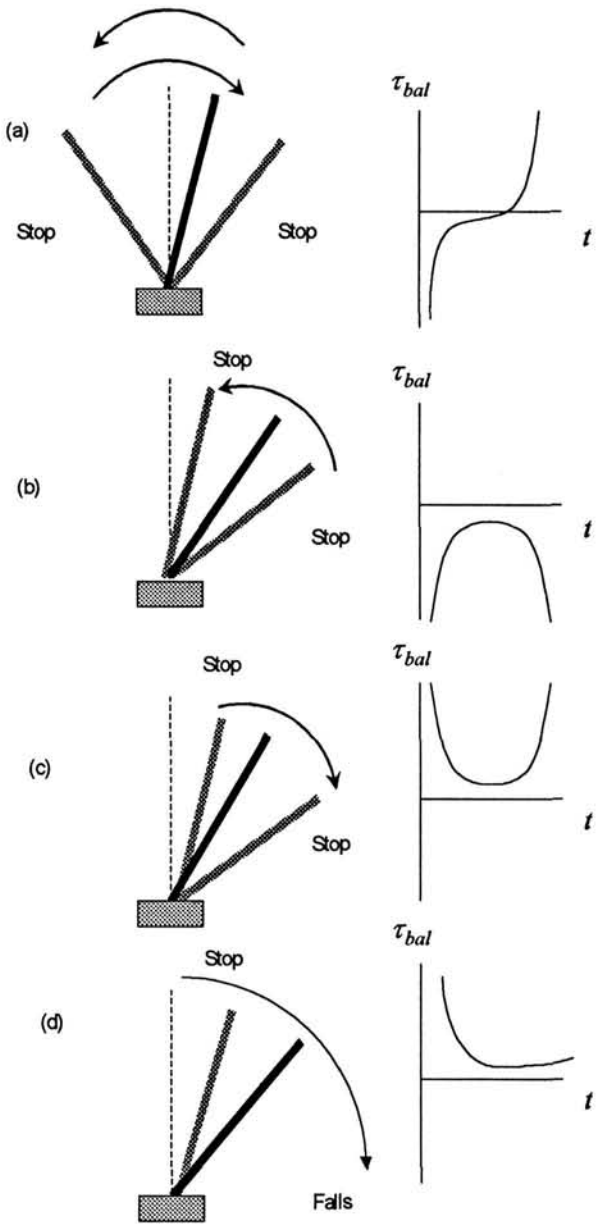

Figure 5. Relationship between the shape of the curve τ_{bal} (time to balance) versus t and the pole trajectory. When the pole is confined to the region $-(\pi/2) < \theta < (\pi/2)$ so that it does not dip below the horizontal, the time series plot of τ_{bal} is composed of four basic shapes: A: sigmoid: τ_{bal} goes through the sequence $-\infty \rightarrow 0 \rightarrow +\infty$ in a reverse S-shape fashion and occurs when the pole starts from one side, overshoots the vertical, and ends up on the other side; B: Inverted-U, which occurs as the pole moves up but undershoots the vertical; C: U-curve, which occurs when the pole drifts away from vertical after an undershoot; and D: Incomplete U-curve, which occurs during a catastrophic fall of the pole. These τ_{bal} versus t curves are the same regardless of which side of vertical the pole resides in. Note that when the pole does not dip below the horizontal, Plots B and C always occur together.

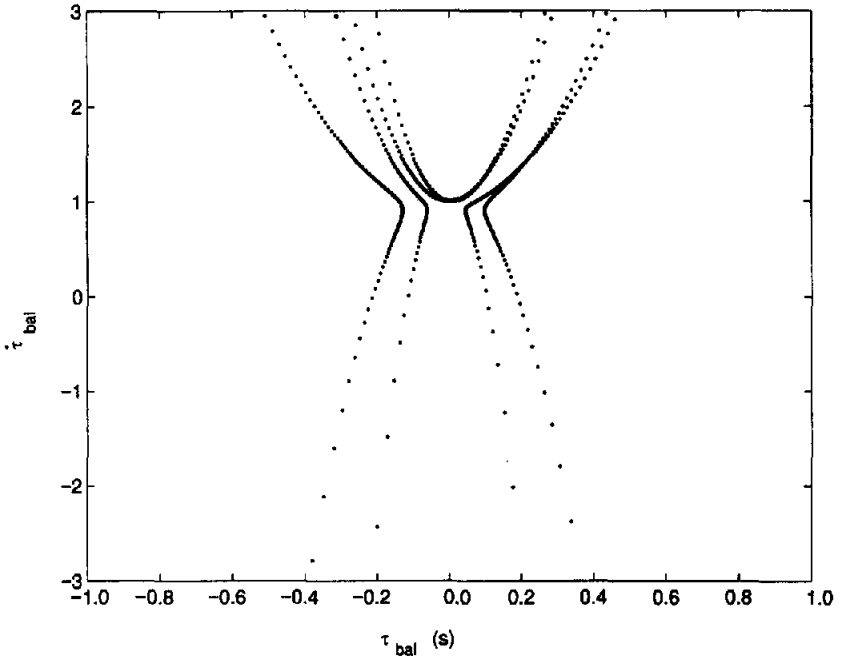

Figure 6. An example of a phase portrait, $\dot{\tau}_{bal}$ versus τ_{bal}, using the data from the successful trial of Figure 2D. $\dot{\tau}_{bal}$ = derivative of τ_{bal} with respect to time.

Classification of Balancing Strategies: Routes to Failure

There were four classes of failures in the present data set. The first class of failures occurred when the pole was balanced for less than 3 s. These "short trials" consisted of an immediate catastrophic fall and did not consist of enough kinematic data to be reliably analyzed. In the second type of failure the participant allowed the base of the cart to contact the edges of our linear track, followed by a catastrophic fall. In these "spatial boundary errors" the participant simply ran out of room with which to maneuver the base and change the direction of the pole. In contrast to these types of errors, participants also produced two classes of failures characterized by a loss of perception–action coupling after some successful balancing had been completed. In order to discriminate the specific routes to failure for the third and fourth failure types, the balancing behavior of the pole was classified based on a local examination of two successive cycles of τ_{bal} behaviors.

On close examination of the experimental time series, it was apparent that the different pole motions (crossing the vertical, undershooting, and drifting) occurred in definite sequences in successful balancing. The observed sequences of τ_{bal} behaviors are shown in Figure 9. From the three pole behaviors seen in successful balancing, here denoted by the boxes, the six paths (shown by arrows) describe the observed sequences between these behaviors. Beginning with the lower box, which symbolizes a crossover of vertical (and a sigmoid τ_{bal} curve), it is apparent that three possible paths exist to the next successful cycle. The participant may continue with another crossover as in Path 1, or if the participant follows Path 2, then the crossover of the current cycle will be followed by a drift (the middle box and a U-shaped τ_{bal} curve). When the crossover is followed by an undershoot (the inverted-U τ_{bal} curve in the upper box), the participant traverses Path 3.

What happened after a participant performed an undershoot? In all the trials of the experimental data set, undershoot was followed by drift, shown here as Path 4. Once a successful drift was reversed, the participant either returned to a crossover (Path 5) or undershot again (Path 6). How did pole falls fit into the classification scheme of τ_{bal} behaviors? The fall of the pole resulted in an incomplete U-shaped τ_{bal} curve (see Figure 5D) or a drift that was not reversed. If one substitutes this unreversed drift into the present τ_{bal} behavior classification, making a diagram of unsuccessful balancing cycles, it is apparent that there exist two different routes to failure. Replace the successful drift (middle box) of Figure 9 with the incomplete U-curve of a falling pole in Figure 10.

What transpired when the participant allowed the pole to fall after a crossover movement (analogous to Path 2 in the successful balancing)? Here the participant performed a successful crossover and decelerated the hand at the base of the pole such that the velocity of the pole approached a zero value. The pole had not reversed directions, however, and the pole moved away from vertical in the same direction in which it crossed the vertical. In order to continue with successful balancing, the participant had to accelerate the pole in the direction opposite to the pole's falling motion and in the opposite direction from the previous hand movement (see Figure 2C). If the participant failed to make this successful reversal of the hand, the pole fell catastrophically. We characterized these cases of failure as a *failure to reverse* the hand.

Conversely, consider the participant failing to balance the pole immediately following an undershoot (analogous to Path 4 in the successful balancing). Here, the participant undershot the vertical and then moved his or her hand in the direction of the pole's fall enough to decelerate the pole to a near-zero velocity. As the pole moved away from vertical the participant did not make a success-

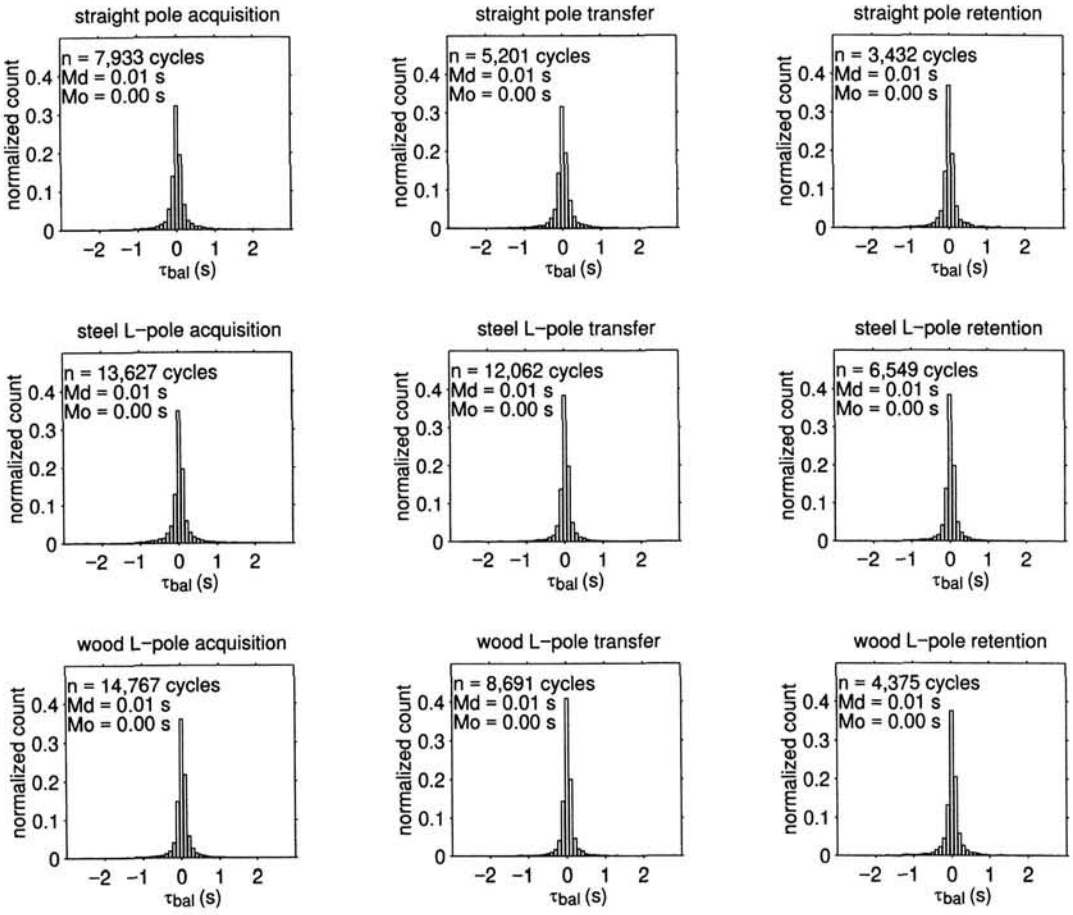

Figure 7. Distribution of τ_{bal} (time to balance) values according to experimental conditions. Md = median; Mo = mode.

ful acceleration of the hand in the same direction as the fall of the pole (and the current direction of the hand). We classified this type of failure as a *failure to continue* the hand.

From this classification, it is clear that while traversing Paths 2 and 4 (see Figure 9) the participant had to perform an active intervention to prevent a failure. One may thus classify Paths 2 and 4 as instances of successfully preventing a failure to reverse and failure to continue, respectively. Similar to the approach taken by others (e.g., Savelsbergh, Whiting, & Bootsma, 1991), we focus on critical situations where an active intervention is needed by the participant to prevent the irreversible fall of the pole. During these two pole sequences (Paths 2 and 4) it is postulated that the participant must pay special attention to the motions of the pole in order to sustain successful balancing. It was especially during these situations that we chose to examine the relation between perceptual variables and pole motions in greater detail.

Correlational Analysis of Hand and Pole Kinematic Variables

Can one find a strong relationship between a measure of the oscillations of the pole with the perceptual variables identified in this experiment? A correlation analysis was performed on all successful cycles of balancing behavior (see also Wagner, 1982). The half-period of hand velocity ($\dot{x}$, an approximation of the "time to upright" the pole) was correlated with the magnitudes of τ_{bal}, $\dot{\tau}_{bal}$, $\dot{x}$, θ, and $\dot{\theta}$, sampled at the onset of hand deceleration. Also included in this analysis were the complements of τ_{bal} and $\dot{\tau}_{bal}$, specifically, τ_{fall} and $\dot{\tau}_{fall}$, which are defined in the following discussion as well as in Treffner and Kelso (1995).

The correlation coefficients are reported in Table 2. The first sets of analyses (first column of Table 2) were performed on the kinematic data collapsed across all cycles. No significant relationships were found. In contrast, when partitioned with respect to the τ_{bal} path classification scheme (see Figure 9, and columns 2–7 of Table 2), a highly significant inverse correlation ($r = -.96$) was found between the τ_{bal} values and $\dot{x}$ period during the cycles of Path 2 (crossover to drift). Furthermore, during the same pole sequences, relationships that approached significance were seen between the $\dot{x}$ period, $\dot{\tau}_{bal}$, τ_{fall}, and $\dot{\tau}_{fall}$. No similar strength relationship between τ_{bal} and $\dot{x}$ period was found in the undershoot to drift (Path 4) cycles (note, however, the relationships with $\dot{\tau}_{bal}$ and $\dot{\tau}_{fall}$). The presence of the significant τ_{bal} and $\dot{x}$ period correlation suggests that participants may be sensitive to the perceptual variable, τ_{bal}, especially during the critical situation of avoiding a

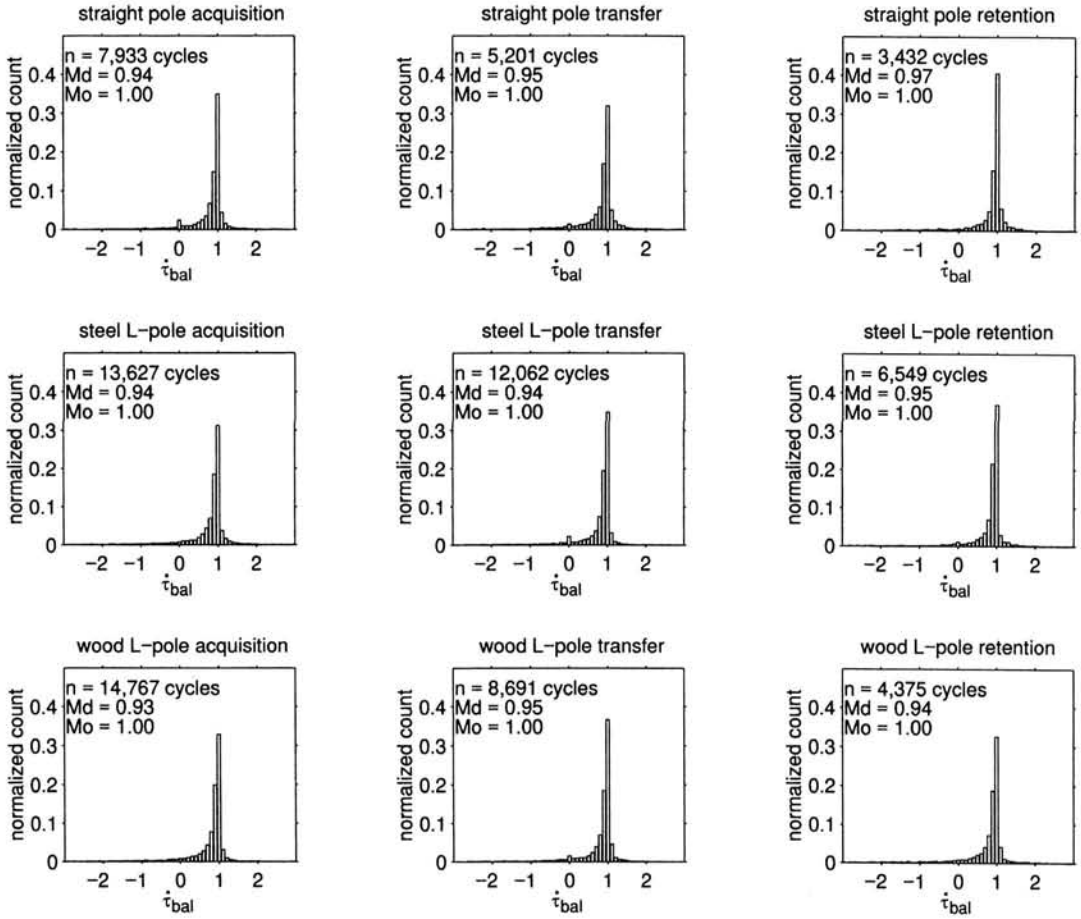

Figure 8. Distribution of $\dot{\tau}_{bal}$ values according to experimental conditions. $\dot{\tau}_{bal}$ = derivative of τ_{bal} with respect to time; Md = median; Mo = mode.

failure to reverse. Note that our experimental data do not show a relationship between τ_{bal} and the period of the pole during typical motions of successful performance. Participants may not be attending to this type of perceptual information during these noncritical situations.

$\dot{\tau}_{bal}$ Values Across τ_{bal} Sequence Classification

What are the typical values of $\dot{\tau}_{bal}$ during the onset of hand deceleration? Although $\dot{\tau}_{bal}$ was found to be relatively invariant across learning and pole conditions in the previously presented kinematic analysis (see Figure 8), the present classification system allows us to perform an analysis based on local cycle kinematics that describe qualitatively distinct pole motions. Specifically, given the six different sequences used in successful balancing (see Figure 9), and two additional sequences seen during failure (see Figure 10), one may partition the kinematic data among these eight categories and examine the values of the perceptual variable $\dot{\tau}_{bal}$ at the onset of hand deceleration. A $0 < \dot{\tau}_{bal} < 0.5$ value would predict a decreasing deceleration and an undershoot of the target, and $0.5 < \dot{\tau}_{bal} < 1$ would predict an overshoot of the target in a "hard collision."

An 8 (τ_{bal} paths) $\times$ 36 (participants) ANOVA with $\dot{\tau}_{bal}$ as the dependent variable was performed on all successful balancing cycles as well as those final cycles that denoted a failure. The significant main effect of τ_{bal} paths, $F(7, 184) = 6.20$, $p < 0.05$, and post hoc Tukey tests revealed that the cycles of crossover to crossover (Path 1) had a statistically similar $\dot{\tau}_{bal}$ mean as the crossover-to-drift cycles, the crossover-to-undershoot cycles, and the failure-to-reverse cycles. The undershoot-to-drift cycles and the drift to crossover cycles had significantly different $\dot{\tau}_{bal}$ means from each other and from the rest of the categories, as did the drift-to-undershoot and failure-to-continue cycles. Importantly, when the different pole sequences were compared according to whether the pole overshot the balance point (second column of Table 3) or undershot the balance point (third column of Table 3), good agreement was found in accordance with previous predictions (see Lee, 1976; Lee, Young, & Rewt, 1992). Note that the mean value in the failure-to-continue sequence may be due to outliers in these particularly noisy data, and the median value is $\dot{\tau}_{bal} = 0.12$. These results suggest that functional stabilization may be considered as a collision problem with respect to the balance point on a cycle-to-cycle time scale.

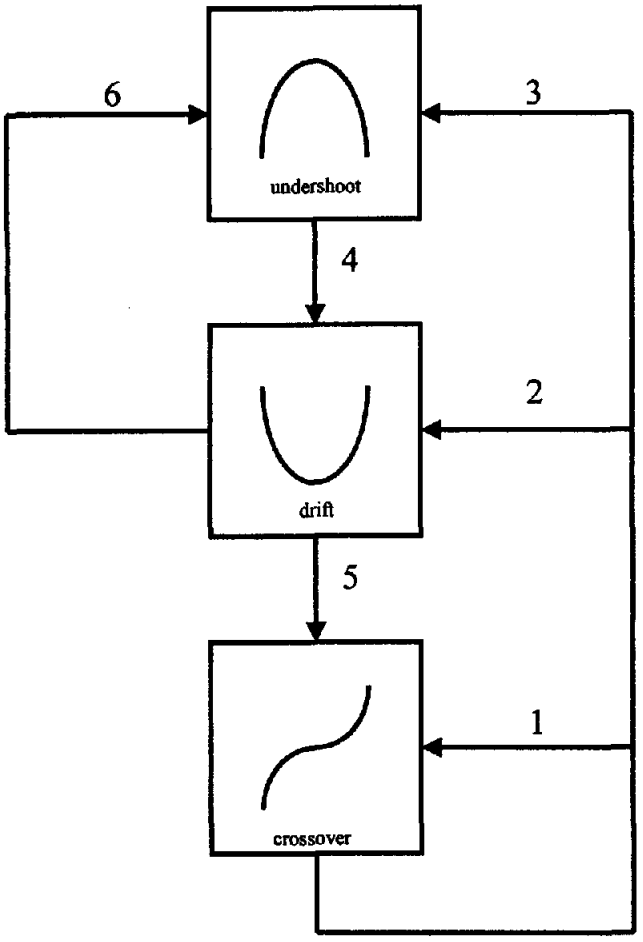

Figure 9. Classification system based on two successive cycles of *successful* balancing behavior. Here the six different paths that are seen in the experimental data and their corresponding τ_{bal} versus t curves (see Figure 5, A–C) are enumerated.

Functional Stabilization: Modeling Considerations

Figure 1 shows a schematic of the cart–pole system used in our analysis of functional stabilization. For the straight-pole condition, the mass m is distributed uniformly over the pole length. The angle of the pole is indirectly controlled through an external applied force F acting horizontally on a cart of mass M. For a given F, the equations of motion for cart position X and the pole angle θ may be derived using physical principles and are given in the Appendix (e.g., see Elgerd, 1967; Ogata, 1978). When formulated as such, the problem of pole balancing often reduces to resolving two key questions: (a) What is the "state" dependence of the function F? and (b) How are the parameters in F modulated to effect a suitable controlled condition. The first question concerns the choice of available kinematic measures (e.g., pole angle and hand position) used to assess the state of the pole. The second question involves the algorithm (and operationally, its physical instantiation) used to implement control. From a control theoretic viewpoint, the variables of choice have always been the canonical ones of pole angle, hand position, and their velocities. For the control strategy, linear control in these variables has been successfully applied even outside the linearized regime of the cart–pole system. In explicit form, linear control means

$$F = \alpha_1\theta + \alpha_2\dot{\theta} + \beta_1 x + \beta_2\dot{x}, \tag{2}$$

where α_1, α_2, β_1, and β_2 are constants. If the unstable system to be controlled has known mechanics, these coefficients may be determined without much difficulty and are usually specified as a range of parameter values. For systems whose intrinsic dynamics are not a priori well established, neuromorphic controllers (controllers based on artificial neural networks) are often used. The determination of the coefficients (*weights*, in neural network parlance) then defines the control problem. Consider now the specific cart–pole system of Figure 1, in which F is implemented by a human controller. In Figure 11 we plot the experimental time series of the hand and pole velocities (Panel A), pole angle (Panel B), and the inferred external force F for a representative time interval (Panel C).

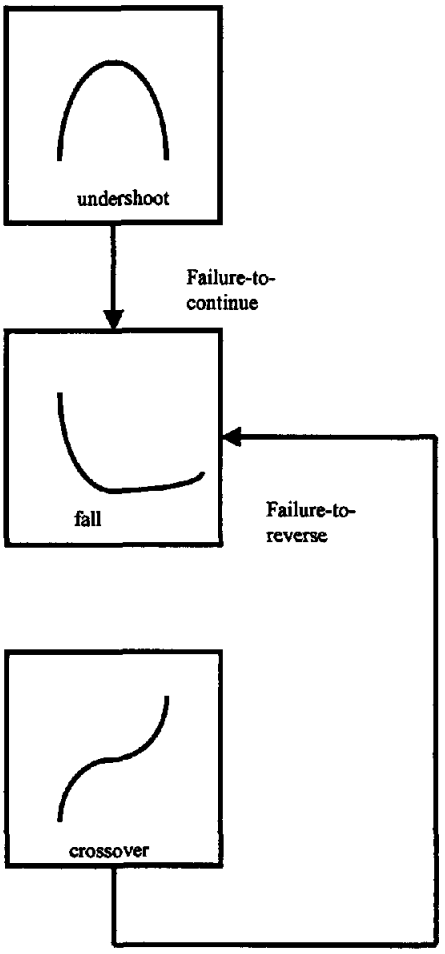

Figure 10. Classification system based on two successive cycles of pole trajectories immediately preceding a catastrophic fall of the pole. Here the incomplete U-curve (τ_{bal} vs. t; see Figure 5D) of a pole fall replaces the U-curve found during successful balancing. Two different routes to failure, the failure to reverse and failure to continue, are seen (compare with Paths 2 and 4 in Figure 9).

Table 2
Correlations Between Hand Motions and Perceptual Variables

Perceptual variable	All cycles	Categorized by τ_{bal} path					
		Path 1	Path 2	Path 3	Path 4	Path 5	Path 6
τ_{bal}	−.01	−.33	**−.96**	−.40	.13	−.03	−.04
$\dot{\tau}_{bal}$	.09	−.06	.38	−.18	.26	−.14	−.08
τ_{fall}	.01	−.18	−.29	−.26	.07	−.02	−.02
$\dot{\tau}_{fall}$	.05	−.13	.26	−.17	.25	−.17	−.12
$\dot{x}$	.04	.06	.14	.07	.08	−.02	−.02
θ	.03	.10	.19	.02	.13	−.01	−.06
$\dot{\theta}$	−.04	−.05	−.11	−.07	−.07	.01	−.01

Note. The correlation value in boldface represents a highly significant inverse relationship between τ_{bal} and the oscillation of the pole when participants successfully avoid a failure to reverse. bal = balance.

To compute F, kinematic information and experimental pole–cart parameters were used as inputs to the equations of motion, which were then inverted to compute the force. Note that pole angle (and to a certain degree, pole and hand velocities) follows F in time, which suggests a strategy based on proportionate or linear control. Note that the control scheme we are proposing asserts the functional form $F = \alpha\theta$, where the determination of the effective coefficient α (possibly nonconstant) is the problem. We propose that the perceptual variables τ_{bal} and $\dot{\tau}_{bal}$ are used to evaluate and sometimes adjust the weightings of the linear control function, on a cycle-by-cycle basis, during critical actions where the participant must actively intervene in order to prevent a failure. To see how this may be done, we assume the following equations for the pole angle and cart position:

$$\ddot{\theta} = -k_1\theta \text{ and } \ddot{x} = +k_2\theta, \tag{3}$$

where $x = L^{-1}X$ is the hand position normalized with respect to the pole length, and the coefficients k_1 and k_2 are nonlinear and possibly discontinuous functions of x, $\dot{x}$, θ, $\dot{\theta}$. Note that this is a perfectly legitimate formulation because no restrictions on k_1 and k_2 are made at this point.

It is helpful to discuss the special case when these functions (k_1 and k_2) are constant. If the controller maintains constant and positive k_1 and k_2 at all times, then the pole oscillates about the vertical with a frequency $\Omega = \sqrt{k_1}$ and an amplitude that depends on the initial conditions. From Equation (3) one can see that the motion of the cart also oscillates at the same frequency. In addition, depending on the initial conditions, the center of oscillation of the cart moves at a constant velocity. On the other hand, if $k_1 < 0$ and constant (k_2 constant, any sign), then, except under very special initial conditions, the pole falls at an exponential rate away from its starting position. The hand executes similar (exponential) behavior plus some constant velocity motion. Plots of the experimental values of k_1 (computed from $-(\ddot{\theta}/\theta)$) over time, however, show that k_1 may be positive and negative depending on the situation and varies smoothly except at certain discrete points (see Figure 12A for a representative experimental time series). These singularities occur when the θ value approaches zero.

Table 3
Comparison of Mean $\dot{\tau}_{bal}$ Values Across τ_{bal} Path Classification

Sequence	Path	Undershoot ($0 < \dot{\tau}_{bal} < 0.5$)	Overshoot ($0.5 < \dot{\tau}_{bal} < 1.0$)
Crossover to crossover	1		0.96 (0.07)
Crossover to drift	2		0.87 (0.12)
Crossover to undershoot	3		0.97 (0.08)
Undershoot to drift	4	0.24 (1.77)	
Drift to crossover	5	0.18 (1.66)	
Drift to undershoot	6	0.36 (1.82)	
Failure to reverse	7		0.94 (0.12)
Failure to continue	8	−2.26 (7.72)	

Note. Path refers to τ_{bal} versus t trajectories shown in Figures 9 and 10. The $\dot{\tau}_{bal}$ mean of −2.26 in the failure-to-continue sequence may be due to outlier values. The median value in this sequence is 0.12. Standard deviations are enclosed in parentheses.

We now consider how the k_1 coefficient relates to the perceptual variables. Differentiating the expression for τ_{bal} with respect to time,

$$\dot{\tau}_{bal} = 1 - \frac{\theta\ddot{\theta}}{\dot{\theta}^2}, \tag{4}$$

which together with Equation (3), can be rewritten as $\dot{\tau}_{bal} = 1 + k_1\tau_{bal}^2$. On solving for k_1, this yields

$$k_1 = \frac{\dot{\tau}_{bal} - 1}{\tau_{bal}^2}, \tag{5}$$

the relationship we are looking for. Note that the sign of k_1 depends only on $\dot{\tau}_{bal}$:

$$k_1 = \begin{cases} < 0, & \dot{\tau}_{bal} < 1 \\ > 0, & \dot{\tau}_{bal} > 1. \end{cases} \tag{6}$$

In the preceding discussions, we showed that for $k_1 > 0$ and constant, the pole oscillates with a fixed frequency and amplitude. Relaxing now this constraint of regular motion, a simple strategy to keep the pole oscillating about the vertical is to keep $k_1 > 0$ at all times. This means maintaining the condition $\dot{\tau}_{bal} > 1$ regardless of the state of the pole and hand. The action of the controller is such that it aims to always overshoot the vertical without considering the consequence for the next cycle. Because of its open-ended nature, this strategy suffers from the undesirable consequences that even though the restoring feature (because $k_1 > 0$) may control the initial oscillations, subsequent motion may lead to

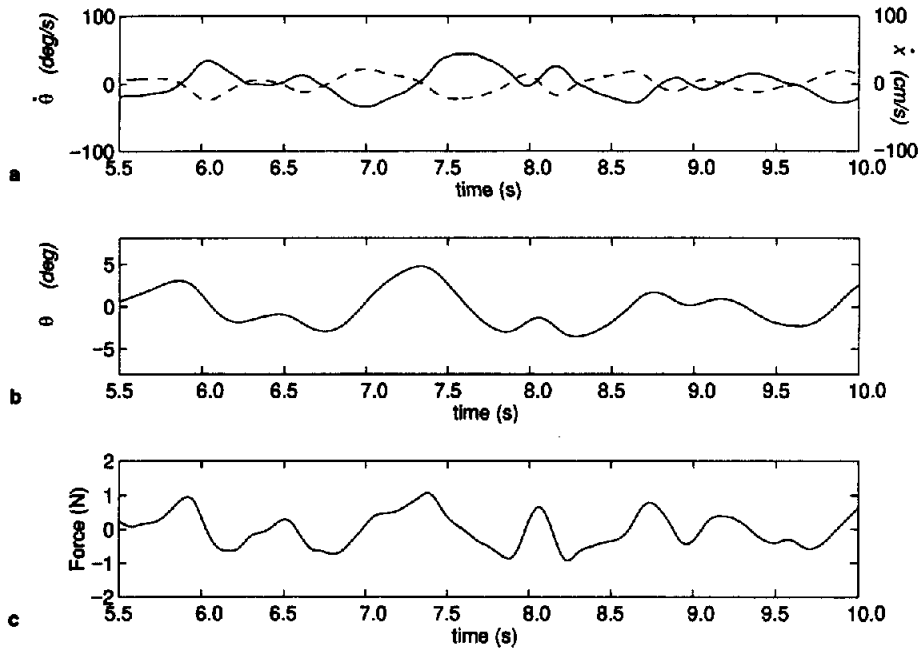

Figure 11. Time series of hand velocity (solid lines) versus pole (stippled lines) velocity (A), pole angle (B), and force (C), as computed from Equations A1 and A2. Note the close relationship between the force and the angle during nonperiodic pole motions. This implies a k value that is not constant and that exhibits a complex dynamics (see text for discussion).

increasing excursions in the pole angle and therefore a possible failure.

Possible Control Strategy

In the previous discussion, we explored a simple strategy for pole balancing without specifying how the internal mechanisms of the controller are modified contingent on the perceived state of the pole. To be able to do more, one must incorporate some physical aspects of the problem as well as a model of the control mechanism. For the moment, we disregard the effect of the horizontal velocity and assume that the force is dependent only on θ, $\dot{\theta}$, and x. Just as we have expressed the motion of the pole angle using a harmonic-like equation, we now assume that the physical force applied by the hand to the cart can be expressed in the form $f = \alpha_1\theta + \alpha_2\dot{\theta} + \beta x$, where β is a constant and α_1 and α_2 are functions whose variations will be specified shortly. Here we have used the reduced form of the force, $f = L^{-1}m^{-1}(m + M)^{-1}F$ as given in the Appendix. The control force can be rewritten as $f = (\alpha_1 + \alpha_2\tau_{bal}^{-1})\theta + \beta x = \alpha(\tau_{bal})\theta + \beta x$, where τ_{bal} is as defined in Equation 1. For the purpose of easing the discussion, we consider the linearized form of the full equations of motion (Equations A3 and A4 of the Appendix). In the linear region, the functions k_1 and k_2 in Equation 3 are given by

$$k_1 = \frac{\alpha - \omega^2 + \beta\dfrac{x}{\theta}}{\dfrac{4}{3} - \mu}, \text{ and} \tag{7}$$

$$k_2 = \frac{\dfrac{4}{3}\alpha - \mu\omega^2 + \dfrac{4}{3}\beta\dfrac{x}{\theta}}{\dfrac{4}{3} - \mu} \tag{8}$$

(see Appendix). Because the principal stimulus is the visually specified pole angle, we assume that the variation of the parameters is implemented principally through the α, and keeps the x coefficient β constant (note that μ is the normalized mass). Recall from a previous discussion that keeping $\tau_{bal} > 1$ at all times means $k_1 > 0$, so that although there is always a restoring force on the angle, there is a possibility for failure due to uncontrolled oscillation amplitude. Instead of this open-ended condition, consider the case in which the controller tries to maintain the condition $\tau_{bal} = 1$. When imposed at all times (t), this also leads to the unrealistic situation of no oscillation at all (i.e., the frequency $\Omega = \sqrt{k_1} = 0$, identically). For successful balancing, τ_{bal} needs only to be kept within the range from 0.5 to 1. Additionally, τ_{bal} needs not be restricted as such all the time, but only at time points T_n of peak hand velocity or onset of deceleration (see Table 3). In contrast, movement cycles that lead to pole undershoots and sometimes failure yield τ_{bal} below 0.5 at the same time points T_n. (Note that by Lee's, 1976, analysis, a pole moving upright and maintaining $\tau_{bal} = 0.5$ at all times reaches the vertical with zero velocity and acceleration.) This suggests a control strategy that keeps τ_{bal} between 0.5 and 1.0. Assume that before the application of the control action, $\tau_{bal} < 0.5$, and thus there is a potential for undershooting the vertical. Then, at the next time instant, in order to

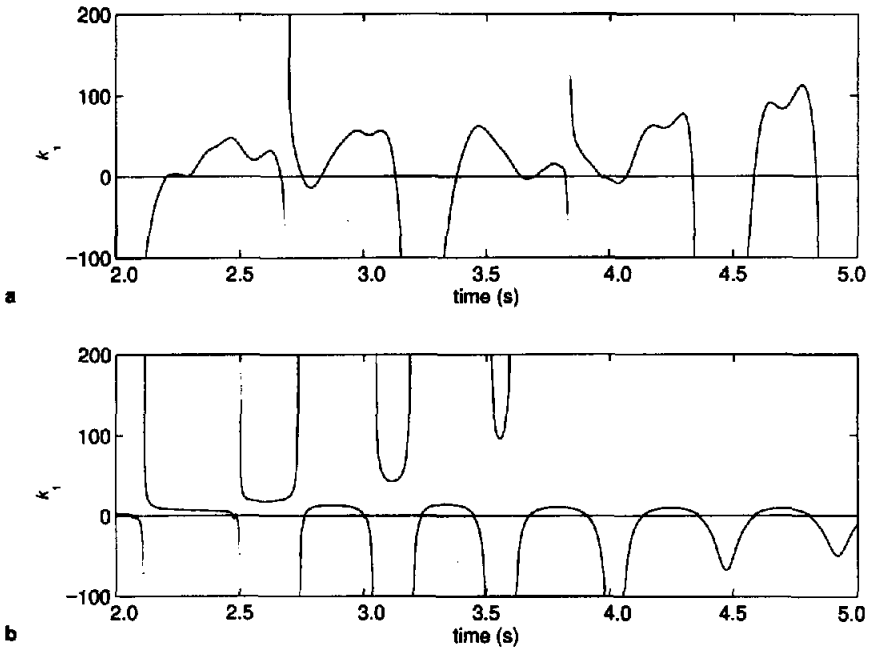

Figure 12. Representative time series of experimental (A) and simulated (B) k_1 values obtained from the ratio $-\dot{\theta}/\theta$ (see Equation 3). θ = the angle the pole makes with the balance point; $\dot{\theta}$ = angular velocity.

increase $\dot{\tau}_{bal}$ to 0.5, a controller must increase k_1 at least by an amount, δk_1, where

$$\delta k_1 = \left| \frac{\dot{\tau}_{bal} - 0.5}{\tau_{bal}^2} \right|. \quad (9)$$

From Equation 7, this can be implemented by incrementing α by an amount $\delta\alpha = (4/3 - \mu)\ \delta k_1$. In the opposite case, of $\dot{\tau}_{bal} > 1$, the controller must decrease k_1 by the same amount above to reduce $\dot{\tau}_{bal}$ to 0.5, or alternatively, by an amount

$$\delta k_1^* = \left| \frac{\dot{\tau}_{bal} - 1}{\tau_{bal}^2} \right|, \quad (10)$$

that is, $\delta\alpha = (4/3 - \mu)\ \delta k_1^*$, to bring down the value of $\dot{\tau}_{bal}$ to 1. To test the suitability of this strategy, we simulated the hand and pole motion using a simple update rule. The parameter α was incremented or decremented by a small amount $\Delta\alpha = \epsilon\,(4/3 - \mu)\ \delta k_1$ or $\Delta\alpha = \epsilon\,(4/3 - \mu)\ \delta k_1^*$, $0 < \epsilon < 1$, depending on whether $\dot{\tau}_{bal}$ was less than 0.5 or greater than 1.0 at time points when the hand velocity reached an extremum. After changing α, its value was fixed for the duration of the movement cycle (i.e., until the next peak hand velocity was again encountered). After a few iterations, $\dot{\tau}_{bal}$ was expected to stabilize to either 0.5 or 1.0 at peak hand velocity positions. Note that k_1 is a function not just of α but also of the other state variables. Setting $k_1 = 0$ or, equivalently, $\dot{\tau}_{bal}$ to 1 at peak velocity maxima, does not mean $k_1 = 0$ at all times. In fact, k_1 may be positive or negative depending on the state of the pole. Figures 12B and 13 show the results of a simulation using the earlier mentioned strategy. Plots of the time series of $\dot{x}$ (solid lines) and $\dot{\theta}$ (stippled lines) are shown in Figure 13A. The corresponding simulation for τ_{bal} and $\dot{\tau}_{bal}$ are shown in Figure 13, B and C. Notice that these plots correspond to time series of the same variables from the experimental data shown in Figure 4, A–C. Initially, the pole exhibits essentially periodic oscillations. After some iteration however, undershoots occur in which the pole does not quite reach the vertical. The simulation shows qualitative agreement with our experimental data (compare with Figure 4). The phase portrait τ_{bal} versus $\dot{\tau}_{bal}$, plotted in Figure 14, also compares quite favorably with the experimental phase portrait of Figure 6. Parabolic curves correspond to successful crossing of the vertical. The outermost curves result from undershoots and drifts (see discussion of Figure 6). Apparent gaps in the trajectories are due to the discrete nature of the updates of the α parameter in our model.

Further Discussion

Similar optical variables (to τ_{bal}) have been shown to be available to human observers through changes in retinal flow fields. Early works showed that ambient optic arrays that produce retinal expansion patterns provide important information about approaching objects and can elicit avoidance or defensive behaviors in animals and infants (Schiff, 1965; Schiff & Detwiler, 1979). Manipulations of the incident optic array have also provided evidence that when catching oncoming objects, participants gear their actions to optic τ rather than to separate distance and velocity information. For example, Savelsbergh et al. (1991) showed that the time of appearance of the maximal closing velocity (and opening velocity; see Savelsbergh, 1995) of the hand was significantly later for an oncoming ball presented as progressively deflating than for a ball of constant size. This suggests that participants geared their actions to the characteristics of the deflating ball and not solely to its time of arrival based on computations involving position and velocity attributes.

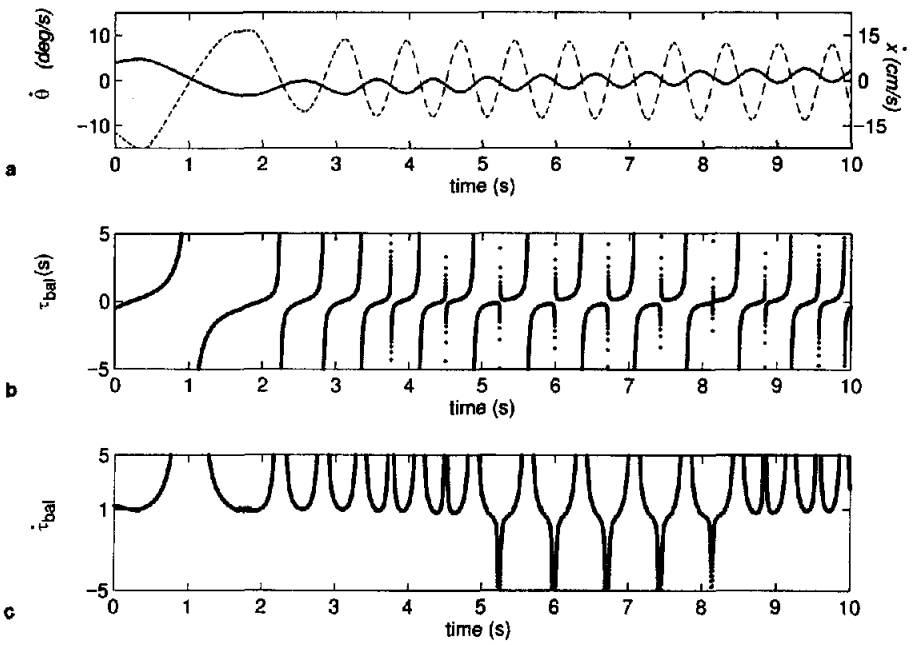

Figure 13. Simulation of (A) hand (solid lines) and pole velocities (stippled lines), (B) τ_{bal} (time to balance), and (C) $\dot{\tau}_{bal}$. The parameters used were: $\beta = 4.6$, $\epsilon = .75$, $\alpha_0 = 2.8$, for the initial value of α, $\mu = 0.4$, $\omega = \sqrt{9.8}$, noise $SD = 0.001$, and a simulation run time of 100 s. Initial conditions were $\theta_0 = 28°$, $\dot{\theta}_0 = -11.5°/s$, $x_0 = -50$ cm, $\dot{x}_0 = 4.1$ cm/s. $\dot{\tau}_{bal}$ = derivative of τ_{bal} with respect to time. θ = the angle the pole makes with the balance point; $\dot{\theta}$ = angular velocity.

For the current experimental research, τ_{bal} specifies the relative rate of constriction of the angle the pole makes with the vertical. There has been controversy in the literature concerning the types of perceptual information (e.g., visual and auditory) that are available to observers (e.g. Lee & Reddish, 1981; Shaw, McGowan & Turvey, 1991; Tresilian, 1994; Wagner, 1982), how these different modalities may be combined (Savelsbergh, Whiting, & Pepers, 1992; Tresilian, 1994), and to what extent the empirical data support the optic τ versus alternative possibilities (Bootsma, Fayt, Zaal, & Laurent, 1997; Bootsma & Oudejans, 1993; Savelsbergh, 1995; Tresilian, 1994, 1995; Wann, 1996).

Our results show that although there are differences in the initial performance between balancing a straight versus an L-pole (e.g., see Figure 2, A and C) these differences become insignificant as the participants become more skilled at the task (e.g., see Figure 2, B and D). This indicates that participants eventually learn the balance position of the L-poles and adapt such that the resulting behavior is indistinguishable from that of a straight-pole task. One crucial component of the process of learning to balance the L-pole seems to be in the initial act of setting up the pole. Because the balance position does not manifest itself until after the participant has moved the pole or manipulated the initial conditions, this suggests that the perception of this physical attribute (equilibrium position) is acquired dynamically. If the participants had aligned the L-shaped pole with an absolute reference (the vertical position) instead of the dynamically acquired balance point, the pole would fall in the same direction of the open end of the L-shaped pole (e.g., to the left for the L-poles in Figure 1, B and C). In fact, the participants who were excluded from this experiment demonstrated significantly more falls in this direction during the failed acquisition trials of the L-shaped poles, $\chi^2(1, N = 28) = 182.08$, $p < .01$. These participants did not typically set the initial conditions of the task at this balance point (in contrast to successful participants, especially later in practice) and were not able to identify the balance point of the L-shaped poles during exit interviews. Evidence of perception of physical attributes that are not directly available is also found in other studies. For example, by actively wagging a nonuniform stick, a person acquires information on essential (with respect to translational and rotational motions) attributes of the object, such as its inertia tensor (Pagano & Turvey, 1992; Turvey, Burton, Pagano, Solomon, & Runeson, 1992). Because the present analysis centers on the kinematics of successful balancing rather than on the participants' acquisition of this skill, the question of how one can adjust his or her action to fit different physical constraints (i.e., different balance points) was not explored in detail. This will be an interesting topic for future research.

The postulated control strategy in balancing may be compared with that used in controlled braking. Although the essence of the controls in both cases is the same (i.e., both use $\dot{\tau} = c$ at some point), the resulting dynamics are, of course, very different. In Lee (1976), it can be shown that the regulated motion is one of constant deceleration caused by a frictional force. Because the effect of friction ends the moment the stop position is reached, there is no need for analysis beyond the goal position. In pole balancing, the destabilizing factor of gravity acts like a reverse spring (at least for small angles) and is balanced by the spring-like mechanism of the controller. Another difference between the two systems is the way τ information appears to be used. Whereas Lee's (1976) strategy was to continuously modulate action on the basis of the current

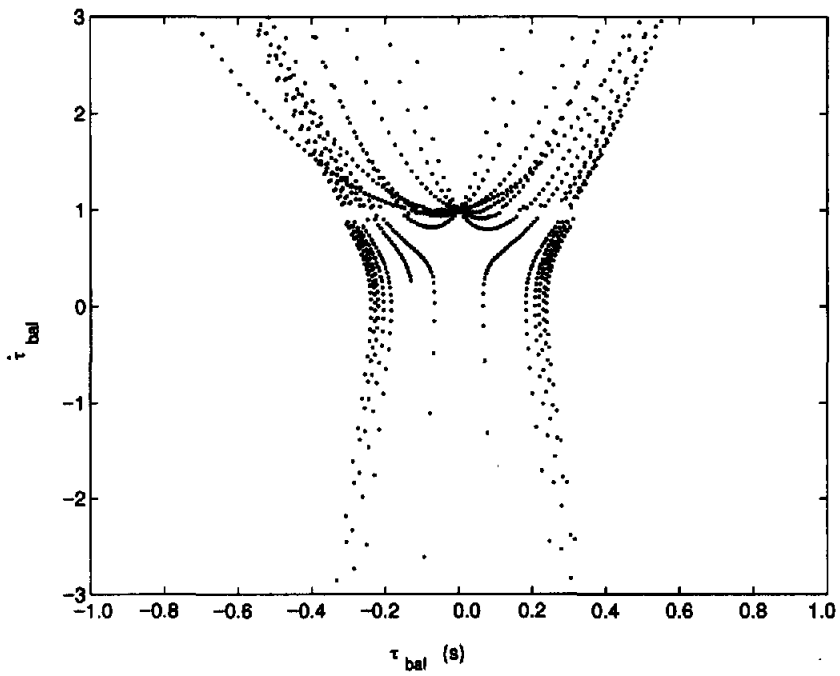

Figure 14. Phase portrait of $\dot{\tau}_{bal}$ versus τ_{bal} (time to balance) using the simulation of Figure 13. $\dot{\tau}_{bal}$ = derivative of τ_{bal} with respect to time.

value of $\dot{\tau}$, here the stiffness parameter is updated only near peaks of hand velocity. Although our data indicate that $\dot{\tau}_{bal}$ was conserved most of the time near such points, it is still unclear whether this is the strategy itself, or a part of (even a consequence of) some other strategy. Previous experiments by Treffner and Kelso (1995) suggested that a more complete description may have to include another τ variable, the inverse of the rate of change of the angle with respect to the horizontal. In Treffner and Kelso (1995) this was called τ_{fall} and was also found to be nearly conserved at points of maximum hand velocity. It can be shown that the pair, τ_{fall} and τ_{bal}, can be derived from θ and $\dot{\theta}$ and vice versa. Thus, no information is lost. They are just represented and perceived differently.

Conclusion

To survive in a changing environment, biological systems often must adopt means for quickly and efficiently regulating their actions. An important, but largely unstudied subset of these tasks is functional stabilization in inherently unstable situations, addressed here using the prototypical task of pole balancing. We showed how the task may be accomplished by using visual perceptual information to adjust a parameter of a linear controller. To a certain extent, the control strategy we have presented parallels that used in some artificial neural networks. One key feature of a neuromorphic linear controller is the adjustment of the weights using an update rule. In our case, we used constraints on a perceptual variable (τ_{bal}, a temporal measure) to adjust a parameter (here, also a temporal measure related to the fall time of the pole) of a linear controller. We propose that the perceptual variables τ_{bal} and $\dot{\tau}_{bal}$ are used to monitor and sometimes adjust the weightings of this linear control function, on a cycle-by-cycle basis, during critical actions in which the participant must actively intervene in order to prevent a failure. We performed a learning and transfer experiment to explore a hypothesized balancing strategy based on the variable τ_{bal}. For successful balancing, regardless of the general configuration and inertial properties of the pole, participants coupled their hand movements to the movements of the pole. These actions were most predictable near regions of hand velocity extrema where participants exhibited quite consistent values of $\dot{\tau}_{bal}$. Successful balancing, in which the pole is maintained upright with some degree of oscillation, yields $\dot{\tau}_{bal}$ ranging from 0.5 to 1.0 at the onset of interceptive hand action. Perceived deviations from this range may be used prospectively to adjust a future action.

References

Anderson, C. W. (1989, April). Learning to control an inverted pendulum using neural networks. *IEEE Control Systems Magazine, 4,* 31–36.

Bachman, J. C. (1961). Specificity vs. generality in learning and performing two large muscle motor tasks. *Research Quarterly, 32,* 3–11.

Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983). Neuron-like adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics, 5,* 834–846.

Bertenthal, B. I., & Bai, D. L. (1989). Infants' sensitivity to optical flow for controlling posture. *Developmental Psychology, 25,* 936–945.

Bertenthal, B. I., Rose, J. L., & Bai, D. L. (1997). Perception–action coupling in the development of visual control of posture. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1631–1643.

Bootsma, R. J., Fayt, V., Zaal, F. T. J. M., & Laurent, M. (1997). On the information-based regulation of movement: What Wann (1996) may want to consider. *Human Movement Science, 13,* 1282–1289.

Bootsma, R. J., & Oudejans, R. R. D. (1993). Visual information about time to collision between two objects. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1041–1052.

Elgerd, O. I. (1967). *Control systems theory.* New York: McGraw-Hill.

Geva, S., & Sitte, J. (1993, October). A cartpole experiment benchmark for trainable controllers. *IEEE Control Systems, 13,* 40–51.

Guez, A., & Selinsky, J. (1988). A trainable neuromorphic controller. *Journal of Robotic Systems, 5,* 363–388.

Jeka, J. J., & Lackner, J. R. (1994). Fingertip contact influences human postural control. *Experimental Brain Research, 100,* 495–502.

Jeka, J. J., & Lackner, J. R. (1995). The role of haptic cues from rough and slippery surfaces in human postural control. *Experimental Brain Research, 103,* 267–276.

Jeka, J. J., Schöner, G., Dijkstra, T., Ribeiro, P., & Lackner, J. R. (1997). Coupling of fingertip somatosensory information to head and body sway. *Experimental Brain Research, 113,* 475–483.

Kelso, J. A. S. (1998). From Bernstein's physiology of activity to coordination dynamics. In M. L. Latash (Ed.), *Progress in motor control* (pp. 203–219). Champaign, IL: Human Kinetics.

Kelso, J. A. S., Fuchs, A., Lancaster, R., Holroyd, T., Cheyne, D., & Weinberg, H. (1998). Dynamic cortical activity in the human brain reveals motor equivalence. *Nature, 392,* 814–818.

Kinsbourne, M., & Hicks, J. (1978). Functional cerebral space: A model for overflow, transfer, and interference effects in human performance: A tutorial review. In J. Requin (Ed.), *Attention and Performance VII* (pp. 345–362). Hillsdale, NJ: Erlbaum.

Kwakernaak, H., & Sivan, R. (1972). *Linear optimal control systems.* New York: Wiley.

Lee, D. N. (1976). A theory of visual control of braking based on information about time-to-collision. *Perception, 5,* 437–459.

Lee, D. N., & Aronson, E. (1974). Visual proprioceptive control of standing in human infants. *Perception and Psychophysics, 15,* 529–532.

Lee, D. N., & Reddish, P. E. (1981). Plummeting gannets: A paradigm of ecological optics. *Nature, 293,* 293–294.

Lee, D. N., Young, D. S., & Rewt, D. (1992). How do somersaulters land on their feet? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1195–1202.

Ogata, K. (1978). *System dynamics* (pp. 531–536). Englewood Cliffs, NJ: Prentice Hall.

Pagano, C. C., & Turvey, M. T. (1992). Eigenvectors of the interia tensor and perceiving the orientation of a hand-held object by dynamic touch. *Perception and Psychophysics, 52,* 617–624.

Savelsbergh, G. J. P. (1995). Catching "grasping tau." Comments on J. R. Tresilian (1994). *Human Movement Science, 14,* 125–127.

Savelsbergh, G. J. P., Whiting, H. T. A., & Bootsma, R. J. (1991). Grasping tau. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 315–322.

Savelsbergh, G. J. P., Whiting, H. T. A., & Pepers, J. R. (1992). The control of catching. In J. J. Summers (Ed.), *Approaches to the study of motor control and learning.* Amsterdam: Elsevier Science.

Schiff, W. (1965). Perception of impending collision: A study of visually directed avoidance behavior. *Psychological Monographs: General and Applied, 79,* (11, Whole No. 604).

Schiff, W., & Detwiler, M. L. (1979). Information used in judging impending collision. *Perception, 8,* 647–658.

Shaw, B. K., McGowan, R. S., & Turvey, M. T. (1991). An acoustic variable specifying time-to-contact. *Ecological Psychology, 3,* 253–261.

Thelen, E. (1990). Coupling perception and action in the development of skill: A dynamic approach. In H. Bloch & B. I. Bertenthal (Eds.), *Sensory-motor organization and development in infancy and early childhood* (pp. 39–56). Dordrecht, The Netherlands: Kluwer.

Tolat, V. V., & Widrow, B. (1988). An adaptive "broom balancer" with visual inputs. *Proceedings of the International Conference on Neural Networks,* San Diego, CA, July, II-641–II-647.

Treffner, P. J., & Kelso, J. A. S. (1995). Functional stabilization of unstable fixed points. In B. Bardy, R. Bootsma, & Y. Guiard (Eds.), *Studies in perception and action:* III (pp. 83–86). Hillsdale, NJ: Erlbaum.

Treffner, P. J., & Kelso, J. A. S. (1997). Scale-invariant memory during functional stabilization. In M. A. Schmuckler & J. M. Kennedy (Eds.), *Studies in perception and action:* IV (pp. 275–279). Hillsdale, NJ: Erlbaum.

Treffner, P. J., & Kelso, J. A. S. (1999). Dynamical encounters: Long-memory during functional stabilization. *Ecological Psychology, 11* (2), 103–137.

Tresilian, J. R. (1994). Perceptual and motor processes in interceptive timing. *Human Movement Science, 13,* 335–373.

Tresilian, J. R. (1995). Visual modulation of interceptive action. A reply to Savelsbergh. *Human Movement Science, 14,* 129–132.

Turvey, M. T., Burton, G., Pagano, C. C., Solomon, H. Y., & Runeson, S. (1992). Role of the inertia tensor in perceiving object orientation by dynamic touch. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 714–727.

Wagner, H. (1982). Flow-field variables trigger landing in flies. *Nature, 297,* 147–148.

Wann, J. P. (1996). Anticipating arrival: is the tau-margin a specious theory? *Journal of Experimental Psychology: Human Perception and Performance, 22, 4,* 1031–1048.

Warren, W. H., Jr. (1988). The visual guidance of action. In O. G. Meijer & K. Roth (Eds.), *Complex movement behavior* (pp. 339–380). Amsterdam: North-Holland.

Warren, W. H., Jr. (1998). Visually controlled locomotion: 40 years later. *Ecological Psychology, 10* (3–4), 177–220.

Zanone, P. G., & Kelso, J. A. S. (1992). The evolution of behavioral attractors with learning: Nonequilibrium phase transitions. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 402–421.

Zanone, P. G., & Kelso, J. A. S. (1997). The coordination dynamics of learning and transfer: Collective and component levels. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1–27.

Appendix

Physical Equations of Motion for the Cart–Pole System

The configuration of the pole balancing apparatus is shown in Figure 1A. The horizontal position of the cart is given by X, and the pole angle measured from the vertical reference line is θ. The equations for X and θ for this straight pole are given by

$$\frac{4}{3}(m + M)L\ddot{\theta} + (m + M)\ddot{X}\cos\theta = (m + M)g\sin\theta, \text{ and} \quad \text{(A1)}$$

$$(m + M)\ddot{X} - mL\dot{\theta}^2\sin\theta + mL\cos\theta\ddot{\theta} = F. \quad \text{(A2)}$$

To reduce the number of parameters, we normalized the variables and parameters as follows: $\mu = m(m + M)^{-1}$, $\omega^2 = gL^{-1}$, $x = L^{-1}X$, $f = (m + M)^{-1}L^{-1}F$. The relevant physical parameters are therefore the frequency ω, the reduced mass μ; and those needed to specify force f. The parameter ω is a measure of how fast the pole will fall when started with an infinitesimal velocity from an upright position if the cart is fixed at the bottom. Using these transformations, we express Equations A1 and A2 in a standard format:

$$\left(\frac{4}{3} - \mu\cos^2\theta\right)\ddot{\theta} = \omega^2\sin\theta - \mu\dot{\theta}^2\sin\theta\cos\theta - f\cos\theta, \text{ and} \quad \text{(A3)}$$

$$\left(\frac{4}{3} - \mu\cos^2\theta\right)\ddot{x} = -\mu\omega^2\cos\theta\sin\theta + \frac{4}{3}\mu\dot{\theta}^2\sin\theta + \frac{4}{3}f. \quad \text{(A4)}$$

Motion Around Balanced Condition

Near balanced condition, $\sin\theta \approx \theta$, $\cos\theta \approx 1$, and $\dot{\theta} \approx 0$, and the previous equations reduce to

$$\left(\frac{4}{3} - \mu\right)\ddot{\theta} = \omega^2\theta - f, \text{ and} \quad \text{(A5)}$$

$$\left(\frac{4}{3} - \mu\right)\ddot{x} = -\mu\omega^2\theta + \frac{4}{3}f. \quad \text{(A6)}$$

Because $\mu < 1$, we divide both sides by $4/3 - \mu$ and rewrite Equations A5 and A6 in the form

$$\ddot{\theta} = -k_1\theta, \ddot{x} = k_2\theta, \quad \text{(A7)}$$

where k_1 and k_2 are given by

$$k_1 \equiv \left(\frac{4}{3} - \mu\right)^{-1}\left(\frac{f}{\theta} - \omega^2\right), \text{ and} \quad \text{(A8)}$$

$$k_2 \equiv \left(\frac{4}{3} - \mu\right)^{-1}\left(\frac{4}{3}\frac{f}{\theta} - \mu\omega^2\right). \quad \text{(A9)}$$

Received July 7, 1998
Revision received June 7, 1999
Accepted July 8, 1999 ■