Published in final edited form as:

J Mot Behav. 2010 November ; 42(6): 389–399. doi:10.1080/00222895.2010.526451.

Bouncing between Model and Data: Stability, Passivity, and Optimality in Hybrid Dynamics

Renaud Ronsse¹ and Dagmar Sternad²

¹Biorobotics Laboratory, Institute of Bioengineering, École Polytechnique Fédérale de Lausanne, Switzerland ²Department of Biology, Electrical & Computer Engineering, and Physics, Northeastern University, Boston, Massachusetts

Abstract

Rhythmically bouncing a ball with a racket is a seemingly simple task, but it poses all the challenges critical for coordinative behavior: perceiving the ball's trajectory to adapt position and velocity of the racket for the next ball contact. To gain insight into the underlying control strategies, the authors conducted a series of studies that tested models with experimental data, with an emphasis on deriving model-based hypotheses and trying to falsify them. Starting with a simple dynamical model of the racket and ball interactions, stability analyses showed that open-loop dynamics affords dynamical stability, such that small perturbations do not require corrections. To obtain this passive stability, the ball has to be impacted with negative acceleration—a strategy that subjects adopted in a variety of conditions at steady state. However, experimental tests that applied perturbations revealed that after perturbations, subjects applied active perceptually guided corrections to reestablish steady state faster than by relying on the passive model's relaxation alone. Hence, the authors derived a model with active control based on optimality principles that considered each impact as a separate reaching-like movement. This model captured some additional features of the racket trajectory but failed to predict more fine-grained aspects of performance. The authors proceed to present a new model that accounts not only for fine-grained behavior but also reconciles passive and active control approaches with new predictions that will be put to test in the next set of experiments.

Keywords

dynamical system; interactive task; optimal content; rhythmic cooperation; stability

Bouncing a ball with a racket is a task that has served as model in scientific fields as various as nonlinear physics (Bapat, Sankar, & Popplewell, 1986; Guckenheimer & Holmes, 1986; Holmes, 1982; Tufillaro, Abbott, & Reilly, 1992; Tufillaro & Albano 1985; Wiesenfeld & Tufillaro, 1987), robotics (Bühler, Koditschek, & Kindlmann, 1988, 1990, 1994; Reist & D'Andrea, 2009; Ronsse, Lefèvre, & Sepulchre, 2006, 2007; Zavala-Rio & Brogliato, 1999, 2001), and human movement control (de Rugy, Wei, Muller, & Sternad, 2003; Dijkstra, Katsumata, de Rugy, & Sternad, 2004; Morice, Siegler, Bardy, & Warren, 2007; Ronsse et al. 2007; Ronsse, Wei, & Sternad, 2010; Schaal, Atkeson, & Sternad, 1996; Siegler, Bardy, & Warren, 2010; Sternad, Duarte, Katsumata, & Schaal, 2001a, 2001b; Wei, Dijkstra, & Sternad, 2007, 2008). This wide-ranging appeal may be partly due to its apparent simplicity

coupled with an unexpected complexity in behavior that is exemplary for all hybrid dynamics (i.e., continuous dynamics controlled by discrete-time events). Hence, the bouncing ball has been invoked as a model system for studying other hybrid dynamics, such as hopping (Brown & Zeglin, 1998; Cavagna, Franzetti, Heglund, & Willems, 1988) or juggling (Sternad, 1999). Further, the task of ball bouncing involves continuous rhythmic elements (the racket trajectory) and discrete events (the impacts with the ball), two features that may require two different control modes that may play the role of primitives (Dégallier & Ijspeert, 2010; Hogan & Sternad, 2007; Schaal, Sternad, Osu, & Kawato, 2004).

In the present article, we review research by Sternad and colleagues to show how this simple model guided research by generating hypotheses that were experimentally tested. As the hypotheses were quantitatively explicit, experimental results could falsify them. This falsification strategy guided the research through a sequence of models that successively revealed in more detail the principles of control that humans apply in such tasks. This dialogue between model and data followed the spirit of Popper (1959): A model is valid until an experimental finding falsifies it. Platt (1964) advocated this strategy and the strong inference principle—ruling out options of explanation—to make true progress in science. This philosophy motivated us to (a) develop a task permitting a simple model, (b) formulate precise hypotheses about control strategies, (c) falsify the predictions, and (d) formulate an alternative model based on the new insights that is then put it to test again.

The article is organized as follows. First, we derive the dynamic model of the bouncing ball. Then, we present hypotheses derived from an open-loop or passive model of control. Although the main prediction of this passive model withstood the tests of falsification, additional findings illustrated that active control mechanisms were also present. Therefore, in the subsequent section, we introduce a closed-loop controller, based on optimality principles. Although this model qualitatively captured active racket control, fine-grained analyses revealed that this optimization algorithm was not adequate to account for human data. Then, we show how the optimal model was adapted to address the experimental findings, and in closing we briefly link the new model with elements of the passive control strategy.

Bouncing Ball Model

One route into understanding how humans control their actions is to develop a mechanical model of the task. The model system was bouncing a ball rhythmically to a target height with a moving surface or paddle that impacts the ball (Figure 1). When the human actor becomes part of the model, replacing or controlling the movements of the surface or paddle, we can analyze how task performance changes with the human controller now as part of the system.

The model for the ball's impact dynamics is based on three assumptions: (a) the movements of the ball follow ballistic flight under the influence of gravity, (b) the impacts are instantaneous with the coefficient of restitution capturing the energy loss at impact (Newton's impact law), and (c) the mass of the ball is significantly larger than the racket such that there is no rebound of the racket due to the ball impact. The first assumption implies that the ball trajectory $b(t)$ between impacts k and $k+1$ traverses a parabola:

$$b(t) = b_k + \dot{b}_k t - 0.5g t^2, \quad (1)$$

in which b_k and $\dot{b}_k$ denote the ball position and velocity immediately after impact k , respectively. Equation 1 holds for $t_k \leq t \leq t_{k+1}$ in which t_k and t_{k+1} denote successive impact

times. From Equation 1, the ball velocity obeys $\dot{b}(t) = \dot{b}_k - gt$. The second assumption is that the racket–ball collisions are governed by the impact law:

$$\begin{aligned} (\dot{b}_k - \dot{r}_k) &= -\alpha (\dot{b}_k^- - \dot{r}_k^-) \\ \iff \dot{b}_k &= -\alpha \dot{b}_k^- + (1+\alpha)\dot{r}_k, \end{aligned} \quad (2)$$

in which $\dot{r}_k$ is the racket velocity at impact k and $\dot{b}_k^- = \dot{b}_{k-1} - g(t_k - t_{k-1})$ is the ball velocity just before the same impact. The coefficient of restitution α captures the energy dissipation at each impact.

In the reviewed studies the authors examined ball bouncing in one vertical dimension in which subjects were instructed to propel the ball to a fixed target height. Experiments and model developments aimed to reveal the control of the racket trajectory $r(t)$ and its impact with the ball to fulfill the task.

A Blind Bouncing Model: Elements of Passive Control

In this section, we develop the hypothesis that dynamic stability of the bouncing ball model presents a mode that human subjects learn to use as a control strategy. The reason for this hypothesis is that if the racket–ball system stabilizes itself, then the actor does not need computational resources to monitor and correct small errors in task performance.

The traditional way to characterize the bouncing ball dynamics is to model the racket movement $r(t)$ as stationary, for example as a sinusoid with fixed amplitude and frequency $r(t) = A \sin(\omega t)$. This approach was used by Holmes (1982) to demonstrate that this simple hybrid dynamics can display a bifurcation route to chaos (see also Bapat et al., 1986; Guckenheimer & Holmes, 1986). Indeed, for a fixed racket frequency ω , it can be shown that small amplitudes produce a period-one attractor (i.e., constant bounce height). Increasing A , this stability bifurcates to a period-two attractor (succession bounce amplitudes alternate between two different heights), then period four, following the period-doubling route to chaos (i.e., showing no longer periodic patterns in bounce amplitudes).

For the investigation of human behavior, the stability region of the period-one attractor is most interesting because it illustrates how the simple rhythmic bouncing task with a single target amplitude can be stable without explicit error corrections of the racket movement. Sternad, Schaal, and colleagues demonstrated that the parametric stability region, identified by Holmes (1982) for perfectly sinusoidal racket movements, can be reformulated for a local parameter independent of the sinusoid assumption (Schaal et al., 1996; Sternad et al., 2001a, 2001b). Specifically, the period-one solution is passively stable if the racket acceleration at impact $\ddot{r}_k$ satisfies

$$-2g \frac{1+\alpha^2}{(1+\alpha)^2} < \ddot{r}_k < 0, \quad (3)$$

in which g is the constant of gravity and α the coefficient of restitution. Equation 3 establishes that passively stable period-one solutions require that the ball is impacted with the racket decelerating before impact, giving rise to a self-stabilizing interactions between ball and racket: If the ball bounces higher than the steady-state target, it will fly longer before contacting the racket again. Because the racket decelerates with this delay, it impacts the ball with a smaller velocity than required at steady state, automatically driving the ball back to the steady-state target height. For example, for a coefficient of restitution at $\alpha = .71$, the region of stability was equal to $\ddot{r}_* \in [-10.1, 0] \text{m/s}^2$.

The first experimental investigations tested this model-based hypothesis: Humans indeed contacted the ball with negative accelerations of the racket in this range and thereby used this passively stable regime (Sternad et al., 2001a, 2001b). Notably, this strategy runs counter to what the biomechanically grounded hypothesis of energy minimization would predict: The energetically optimal contact would be at maximum racket velocity (acceleration zero; Ronsse, Thonnard, Lefèvre, & Sepulchre, 2008). Rather, subjects performed remarkably consistent with the stability hypothesis. Figure 2 shows the data of subjects that performed 40 practice trials, where each of the 40 trial values are averages over 60–80 successive bounces within one 40-s trial. As can be seen, acceleration values of all trials were negative and asymptoted to a value of -4m/s^2 .

Not only were the average racket accelerations negative in many additional task variations, this variable was also consistent with more fine-grained predictions using a Lyapunov stability analysis that predicted acceleration values to be in the approximate range between -5 and -2 m/s^2 (Schaal et al., 1996; Sternad et al., 2001a, 2001b). In addition, results also showed that ball height variability decreased with practice with a time course that covaried with the impact acceleration value: the closer to the optimal acceleration value, the lower the variability of ball height. Note that these results were obtained both using a physical setup and a virtual environment where different coefficients of restitution α could be tested (see Figure 1B; see de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2007, 2008).

The virtual experimental setup also allowed Sternad and colleagues to apply controlled perturbations to test the model's predictions about relaxation behavior (de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2007). If the ball was unexpectedly perturbed following the impact, for instance by manipulating the post-impact velocity of the ball, then the racket trajectory had to transfer more or less energy at the next ball impact to compensate and reestablish the original target height. If humans relied on passive stability alone, then small perturbations would not be actively compensated but rather the unchanged sinusoidal racket movements with negative impact accelerations at contact would stabilize the ball height. However, experimental results clearly rejected this hypothesis: (a) The racket trajectories deviated from sinusoidal behavior and the return to steady state after a perturbation was significantly faster than predicted by the passive model: most of the perturbations were compensated after about two cycles, whereas the passive model predicted relaxation times longer than five cycles; and (b) for large perturbations that kicked the ball out of the basin of attraction, participants easily retrieved the steady state, similar to smaller perturbations. No obvious differences in control strategies were observed in response to perturbations that were in and outside the basin of attraction. These two effects are illustrated in Figure 3.

For large perturbations, Wei et al. (2007) showed that the racket trajectory clearly departed from the quasi-sinusoidal steady-state trajectory after the perturbation—evidence that active control mechanisms were in place. These active modulations of the racket trajectory were mostly visible in the cycle duration, whereas the movement amplitude stayed approximately constant (de Rugy et al., 2003).

Mismatch between model and data was not only evident in the response following large artificial perturbations. Departure from the purely passive regime was also found during steady-state performance (i.e., continuous bouncing without perturbation). To explicitly include the performance variability, the purely deterministic model Equations 1 and 2 were extended with stochastic elements (Dijkstra et al., 2004; Wei et al., 2008). The variability of the experimental data was then compared with the variability of the stochastic model. More precisely, we examined the relation between the (perceived) ball height error (HE) and the subsequent racket velocity at impact. To quantify whether participants tuned the racket velocity at impact as a function of HE during the preceding cycle, the ball height error

during cycle k (HE) was regressed against the racket velocity at the next impact (V). Model-based simulations predicted this regression slope to be -0.4 , as a direct consequence of the stabilizing effect when the racket impacts the ball in the decelerating branch of the racket trajectory. This resulted in lower bounce height with shorter flight duration, and therefore in higher subsequent impact velocity. However, this regression slope can also be interpreted as a proportional feedback gain (V/HE gain), mapping a perceived output (the ball height error) to an updated action input (the racket velocity; Åström & Murray, 2008). The data showed that this regression slope was negative with a magnitude higher than -0.4 in both experienced and nonexperienced participants, especially for higher α conditions (see Figure 4). Given that higher α values were shown to be less stable in the model, this result was interpreted as more active feedback control in these conditions leading to a reduction of the relaxation time of the passive dynamics, as predicted by the model.

Another signature of active control mechanisms appeared through the comparison of the autocovariance structure of the ball state in both model and data. We showed that the lag-1 autocorrelation of the ball height error was near zero during steady-state bouncing, whereas the passive model predicted values between 0 and 1. This result illustrated active correction mechanisms faster than predicted by the passive model (Ronsse et al., 2010). Evidence for translating visual feedback into corresponding action was also found in studies by Morice et al. (2007) and Siegler et al. (2010). This actuation strategy changed when time delays were artificially introduced between the physical racket and its virtual counterpart (i.e., the one actually hitting the ball).

In conclusion, the model based on a sinusoidally shaped actuation of the racket showed that performance without feedback-driven modulations of the racket trajectory is possible. The strongest support that human subjects are sensitive to the stability afforded by the mechanical model was that the average acceleration of the racket at impact was negative. Note that negative acceleration at impact is not only necessary for passive stability. Ronsse et al. (2007) and Ronsse, Lefèvre, et al. (2008) proved that negative acceleration also maximizes robustness in a closed-loop control system. In particular, they showed that the optimal acceleration value—maximizing the stability margin of the controller—lies in the middle of the range of accelerations in Equation 3, drawing a link between passive stability and robust closed-loop control. However, fine-grained analyses of the control strategy, both in steady-state and perturbed performance, illustrated that humans were more efficient than the passive model alone and shortened the system's relaxation time. Moreover, the racket movement after large perturbations clearly deviated from the steady-state trajectory emphasizing the contribution of an active controller (Wei et al., 2007). These results led de Rugy et al. (2003) and Ronsse et al. (2010) to the conclusion that participants tuned into passive stability using active control.

Although a first attempt to capture this mixed control strategy was provided by de Rugy et al. (2003) using an oscillator-based closed-loop model, in the next section we review recent attempts to capture the active component of the bouncing ball control with a new optimality-based exploration.

An Optimal Model: Bounce Like You Reach

The closed-loop model proposed by de Rugy et al. (2003) was based on an oscillator generating the racket movement. The parameters of the oscillator, in particular the time constant setting the period of the cycle, were tuned by the perceived ball height and the expected duration until contact. This worked well to account for the perceived period adaptations following small perturbations. Although this model explained the data, the racket trajectories were approximately sinusoidal for all parameter modulations. Another

finding challenged this relatively stereotyped racket trajectory: In a bimanual version of the task, Ronsse, Thonnard, et al. (2008) found that for increasing interbounce periods the racket trajectory introduced an interval of near-zero velocity (dwell time) in the middle of trajectory, and this dwell time increased with the interbounce period. This finding was put to test with the one-dimensional bouncing ball task, and motivated the design of a new closed-loop model, which no longer relied on a prespecified racket trajectory shape.

A new active model of the controller was proposed in Ronsse et al. (2010) using the theoretical framework of optimal feedback control (OFC). This approach has gained some prominence to account for the control strategies in upper-limb movements such as reaching (Diedrichsen, 2007; Diedrichsen & Gush, 2009; Liu & Todorov, 2007; Todorov, 2005). OFC models coordinated behavior as a tradeoff between maximizing the task reward (e.g., reaching a specific point in space) and minimizing the task cost (e.g., producing energy-expensive movements). Central to this modeling is that perturbations or errors that do not directly affect the task goal are not corrected, therefore making the effector trajectory a priori unknown (Todorov, 2004; Todorov & Jordan, 2002). Although most researchers have focused on voluntary reaching, a few researchers have shown its applicability to rhythmic movements such as locomotion or swimming (e.g., Pham & Hicheur, 2009; Srinivasan & Ruina, 2006; Tassa, Erez, & Smart, 2008). Further support for this framework was provided by linking some of the computational components to particular brain areas that are involved in such functions (Scott, 2004, 2008; Shadmehr & Krakauer, 2008).

Putting the OFC-based approach to test for rhythmic ball bouncing required considering each racket cycle as a separate reach to the next impact. Although pointing to a target position implies zero velocity of the hand at the end of the movement, in the ball-bouncing task the racket moves to a target position with a specific nonzero velocity that bounces the ball to the target height. This is consistent with the recent results by Siegler et al. (2010), who demonstrated that the racket cycle period is actively regulated on every cycle based on perceived variables, but that the racket amplitude is not regulated. It is concluded that the control of racket movements focuses on local variables (e.g., the velocity at the subsequent impact, concluded from the V/H gain analyses) and ignores some more global variables (e.g., the racket amplitude).

To begin, the racket dynamics was modeled as a second-order mechanical system:

$$M\ddot{r}(t) + \gamma \dot{r}(t) = f(t), \quad (4)$$

in which M is the equivalent mass of racket and arm and γ represents the damping in the elbow joint. The net muscular force produced by the elbow flexor and extensor $f(t)$ is low-pass filtered as:

$$\tau \dot{f}(t) + f(t) = u(t) + c u(t) \dot{\sigma}(t), \quad (5)$$

in which τ represents a time constant for the first-order filter. The right-hand side represents the control input $u(t)$ and the second term involving signal-dependent noise acting on the input ($\sigma(t)$ is standard Brownian motion; Harris & Wolpert, 1998; Todorov & Jordan, 2002). We then define the cost function Q comprising the task's costs and rewards:

$$Q = \|r_{k+1} - r_{des,k+1}\|^2 + \|\dot{r}_{k+1} - \dot{r}_{des,k+1}\|^2 + w_{energy} \int_{t_k}^{t_{k+1}} \|u(t)\|^2 dt + w_{rest} \int_{t_k}^{t_{k+1}} \|r(t) - r_0\|^2 dt \quad (6)$$

This cost function Q balances the need to complete the task objective with minimizing control input. The first two terms turn zero when the racket approaches the desired position

$r_{des,k+1}$ and desired velocity at impact $\dot{r}_{des,k+1}$ (i.e., the task is accomplished). The third term penalizes the control input $u(t)$ integrated over the entire cycle duration $\Delta t = t_{k+1} - t_k$ and weighted by w_{energy} . An additional fourth term was added that penalizes large excursions of the racket away from the rest position r_0 . This term prevents the optimization to converge to the trivial solution where the racket is raised until it contacts the ball at the target apex and stays there (Ronsse et al., 2010). The two parameters w_{rest} and w_{energy} are relative weights of the respective cost components.

Using an algorithm derived by Todorov (2005) the time-varying optimal feedback gains L_i were computed that determine the control input $u(t)$ by minimizing the cost function (Equation 6). After minor simplification, these gains can be reduced to the following four terms:

$$u(t) = -L_{rest}(t)(\hat{r}(t) - r_0) - L_{position}(t)(\hat{r}(t) - \hat{r}_{des,k+1}) - L_{velocity}(t)(\hat{\dot{r}}(t) - \hat{\dot{r}}_{des,k+1}) - L_{force}(t)\hat{f}(t), \quad (7)$$

in which $\hat{r}(t)$, $\hat{\dot{r}}(t)$, $\hat{r}_{des,k+1}$, and $\hat{\dot{r}}_{des,k+1}$ are internal estimates of the corresponding variables. These gains can be interpreted as: L_{rest} penalizes large racket excursions away from the rest position r_0 ; $L_{position}$ and $L_{velocity}$ penalize large differences between the actual racket state and the desired one to reach at impact; and L_{force} intends to keep the input force reasonably low. The optimal gains L_i minimizing the cost function (Equation 6) are shaped as shown in Figure 5. Initially, L_{rest} dominates creating a movement to reinitialize the racket position. Later, $L_{position}$ then $L_{velocity}$ dominate to attract the racket toward the desired impact state.

The time course of the gains may be parsed into a segment that reinitializes the trajectory after the previous bounce, and a second segment that creates an impulse-like movement toward the ball contact. Given that the first reinitialization portion lasts longer for longer inter-bounce durations (smaller gravity acting on the ball), the model produces the feature mentioned previously (i.e., increasing dwell times for increasing inter-bounce periods). Figure 6A depicts simulated racket trajectories for the seven different gravity conditions. All cycle durations were normalized to allow better comparison of the changing shape of the trajectory. This prediction contrasts with the previous open-loop model and also the closed-loop model by de Rugy et al. (2003), which, by definition, only produced quasi-sinusoidal racket trajectories. Note that dwell times have been reported in a bimanual version of the bouncing ball and also in other rhythmic tasks (Adam & Paas, 1996; Ronsse, Thonnard, et al., 2008; van der Wel, Sternad, & Rosenbaum, 2010).

We tested this prediction in the one-dimensional bouncing ball task using the virtual setup depicted in Figure 1. Because the ball was moving in the virtual reality environment, the inter-bounce period could be increased by decreasing the gravity acting on the ball. Figure 6A shows the model trajectories, and Figure 6B shows the average racket trajectories of eight participants in seven gravity conditions (again normalized over time for better comparison). For the nominal gravity (light gray, $g_7 = 9.81 \text{ m/s}^2$, corresponding to a steady-state period of about 0.7 s), the racket profile was close to sinusoidal, as seen in previous studies. In contrast, in the lowest tested gravity (black, $g_1 = 0.61 \text{ m/s}^2$, yielding a steady-state period of about 2.8 s), the racket profile had a marked dwell interval with near-zero velocity in the middle of the cycle. The results quantitatively matched the simulations by the optimal model.

Interestingly, the model was also able to reproduce features that indicated active control in previous studies. For example, the fast relaxation time at steady-state bouncing was obtained by an optimal tuning of the V/HF feedback gain. This indicator quantifies the time to return

to steady state after perturbations. Figure 7A shows these V/HE gains in both model and data; the value at g_7 (nominal gravity) is in good agreement with the results obtained in a previous experiment shown in Figure 4 for $\alpha = .6$ (the same α used in Ronsse et al., 2010). As previously mentioned, this V/HE gain tuning led to the lag-1 autocorrelation coefficient $AC-1$ of the height error HE to go to zero, both in model and data (Figure 7).

Redundancy and Global Optimization

It is important to point out that the implementation of the OFC model proceeded in two steps. First, assuming the racket state to reach at impact ($r_{des}, \dot{r}_{des}$), the optimal gains (7) were computed to minimize Q , as previously described. However, many different racket states could achieve the desired bounce height. Therefore, a second step compared the cost Q across different ($r_{des}, \dot{r}_{des}$) pairs to find the state that globally optimized the cost. The corresponding control policy was found using a simple gradient-descent technique.

From this double optimization another prediction could be derived: If humans find the optimal racket state ($r_{des}, \dot{r}_{des}$) at each cycle, then significant correlations should exist between the observed racket position and velocity and the ones predicted by the model. To explicitly test this prediction, we computed the desired states at the next impact, $r_{des,k+1}$ and $\dot{r}_{des,k+1}$, that minimized the cost Q based on the states at cycle k observed in the data. These individually predicted states were compared with the actual racket states at impact $k+1$ for each of the eight participants and the seven gravity conditions.

Figure 8 shows the regression slopes obtained when regressing actual racket position and velocity against predicted position and velocity. Figure 8A shows slopes that are close to zero for all gravity conditions (i.e., no systematic relation between the actual racket positions and the ones predicted by the optimal model). Impact velocity in Figure 8B showed small but significant regression slopes for all gravities. These results suggest that participants might have implemented a strategy somewhat comparable to our model, but they did not choose the globally optimal set of racket states. Note that this lack of dependency only ruled out the second step in the optimization, not the first step, as the position and velocity pairs did reach the target height. This redundancy in the solutions can be illustrated in more detail as follows.

Starting with Equation 1 and its first derivative, it can be shown that the ball height reached after impact k , $b_{H,k}$, is determined by the ball state at this impact ($b_k, \dot{b}_k$) and constrains the ball state that will be reached at impact $k+1$ ($b_{k+1}, \dot{b}_{k+1}$):

$$b_{H,k} = b_k + \frac{(\dot{b}_k)^2}{2g} = b_{k+1} + \frac{(\dot{b}_{k+1})^2}{2g} \quad (8)$$

From Equations 2 and 8, the ball height reached after impact $k+1$ can be written as

$$\begin{aligned} b_{H,k+1} &= b_{k+1} + \frac{(-\alpha \dot{b}_{k+1} + (1+\alpha)\dot{r}_{k+1})^2}{2g} \\ &= b_{k+1} + \frac{(\alpha \sqrt{2g(b_{H,k} - b_{k+1})} + (1+\alpha)\dot{r}_{k+1})^2}{2g} \end{aligned} \quad (9)$$

in which Equation 9 is derived from Equation 8, assuming $\dot{b}_{k+1}^- < 0$. Assuming the ball height $b_{H,k}$ is known to the controller, Equation 9 shows that the racket position ($r_{k+1} = b_{k+1}$) and velocity ($\dot{r}_{k+1}$) at impact $k+1$ are constrained by each other if the objective is to propel the

ball to a target height b_H^* . This redundancy is intuitive: Although for a given impact position only one racket velocity achieves the target height, the same height can be achieved at a higher position but with smaller velocity (taking the modified ball velocity before impact into account). This solution manifold between the two control inputs (the racket position and velocity) presents a redundancy and the global optimization selected one of the infinite set of $(r_{des}, \dot{r}_{des})$ pairs that minimized the cost.

As Figure 8 shows that humans did not use the globally optimal solution, the next question is whether the human participants were sensitive to the solution manifold.

Figure 9 shows the solution manifold (Equation 9) for $b_{H,k+1} = b_H^* = 0.8\text{m}$, $\alpha = .6$, for nominal gravity $g = g_7 = 9.81\text{m/s}^2$. The triplets in white ($b_{H,K}, r_{k+1}, \dot{r}_{k+1}$) are predicted by the model and lie on the manifold as expected (the model tested only $(r_{des}, \dot{r}_{des})$ pairs satisfying (Equation 9)). They are further clustered along a one-dimensional line representing the minimum of the cost function Q for a particular ball height $b_{H,K}$ and initial racket state. In contrast, the corresponding experimental data from a representative subject (black circles) neither lie on the solution manifold nor on the line minimizing the cost Q . However, they do cluster around the solution manifold, as a different perspective in Figure 9B shows.

In sum, the model based on optimal feedback control captured the average control strategy of human participants (i.e., racket trajectories with dwell times and the V/HE gains). However, analyzing the cycle-by-cycle state variables revealed that the model failed to replicate this fine-grained behavior: Humans did not rely on this computationally expensive global optimization when performing the task in real time. Hence, some simplification is required.

We therefore proceeded to hypothesize that the optimal gains $L_i(t)$ are computed only once for all bounces within one trial to minimize the cost Q on average; the control policy (Equation 7) remains stationary across all impacts of a trial. Note that the uncorrelated results for position in Figure 8A implied that the impact position was not controlled across cycles. In contrast, the desired impact velocity was modulated to compensate for the height error HE_k during cycle k (between impacts k and $k+1$) through the following adaptation rule:

$$\dot{r}_{des,k+1} = \dot{r}_{des,k} + \kappa(HE_k), \quad (10)$$

in which κ is the adaptation gain. If κ is equated with the observed V/HE gain, the resulting autocorrelation of HE is close to zero, as observed in the data (Figure 7).

This simplification is also motivated by the fact that impact velocity is a more sensitive input to control the ball height than position: From Equation 9, it can be observed that the ball height depends linearly on the impact position $b_{k+1} = r_{k+1}$, and quadratically on the impact velocity $\dot{r}_{k+1}$. Further, it is computationally less expensive than the former version because the optimization of the gains L_i is performed only once per trial. Nevertheless, we expect this new version of the OFC-based model to match the data at least as well as the previous model because it relies on the same cost function to shape the average racket trajectories among the different conditions.

Reconciling Passive Stability and Active Control

So far, the link between the feedback-driven optimal model and the reliance on passive stability has not been addressed. Although incorporating elements of open-loop passive stability into an OFC model is the next goal for our research, we may already point toward the compatibility of the two theoretical approaches. As demonstrated by Todorov (2005),

signal-dependent noise $\dot{\sigma}(t)$ in Equation 5 decreases the controller gains to zero at the end of the cycle (see also Figure 5). Consequently, the control input $u(t)$ also goes to zero (see Equation 7), such that signal-dependent noise stays at a moderate level. Because the time constant t is small in Equation 5, the input force $f(t)$ also goes to zero together with $u(t)$ at the end of the cycle. Consequently, in order to impact the ball with a positive racket velocity $\dot{r} > 0$, the racket acceleration must be negative $\ddot{r} < 0$ (see Equation 4). As a result, negative acceleration of the racket at impact becomes a consequence of signal-dependent noise. This rationale is in line with the argument that negative acceleration is also advantageous for robust closed-loop controllers to maximize the stability margin (Ronsse et al., 2007).

Discussion

In this article, we presented our progression to understand the control strategy in the task of bouncing a ball with a racket. These investigations employed a dialogue between model and data, consistent with the strong inference principle (Platt, 1964), each new finding trying to falsify the model hypotheses and thus setting the stage for a new improved model.

Based on the mathematical reasoning that a passive actuation mode exists for the task (i.e., moving the racket sinusoidally without error corrections), our first empirical studies presented evidence for this passive control strategy. The critical result was that participants impacted the ball by decelerating the racket prior to ball contact, a prediction derived from the passive model (Sternad et al., 2001a, 2001b). However, subsequent investigations probed further and showed that active control mechanisms are also in place, both during steady-state bouncing and after perturbations (Wei et al., 2007, 2008). Therefore, an active control model based on optimal feedback control was derived and tested. Although the model captured the qualitative control aspects, finer-grained analyses of cycle-by-cycle data revealed mismatches with the data suggesting that the hypothesized control was more complex than what humans might have applied. Most likely, participants did not rely on a cycle-by-cycle optimization, consistent with other experimental findings on the separation of time scales in computationally demanding tasks (Izawa, Rane, Donchin, & Shadmehr, 2008; Ronsse, Miall, & Swinnen, 2009).

At this stage, we showed a way to accommodate this data into the optimal control model, thereby providing a link between this closed-loop model and the elements of passive stability. This last model awaits to be tested and potentially falsified again by new experimental findings.

Acknowledgments

Renaud Ronsse was funded by the European Union grant Evryon (FP7-ICT-2007.8.5 #231451). Dagmar Sternad was funded by the grants NIH R01-HD045639, NSF PAC-0450218, and NSF DMS-0928587.

References

- Adam JJ, Paas FCGC. Dwell time in reciprocal aiming tasks. *Human Movement Science*. 1996; 15:1–24.
- Åström, KJ.; Murray, RM. *Feedback systems: An introduction for scientists and engineers*. Princeton, NJ: Princeton University Press; 2008.
- Bapat C, Sankar S, Popplewell N. Repeated impacts on a sinusoidally vibrating table. *Journal of Sound Vibration*. 1986; 108:99–115.
- Brown, B.; Zeglin, G. The bow leg hopping robot. *Proceedings of the IEEE International Conference on Robotics and Automation*; 1998. p. 781–786.
- Bühler M, Koditschek D, Kindlmann P. A one degree of freedom juggler in a two degree of freedom environment. *IEEE/RSJ Conference on Intelligent Systems and Robots*. 1988; 1988:91–97.

- Bühler M, Koditschek D, Kindlmann P. A family of robot control strategies for intermittent dynamical environments. *Control Systems*. 1990; 10(2):16–22.
- Bühler M, Koditschek D, Kindlmann P. Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research*. 1994; 13:101–118.
- Cavagna GA, Franzetti P, Heglund NC, Willems P. The determinants of the step frequency in running, trotting and hopping in man and other vertebrates. *Journal of Physiology*. 1988; 399:81–92. [PubMed: 3404473]
- de Rugy A, Wei K, Muller H, Sternad D. Actively tracking “passive” stability in a ball bouncing task. *Brain Research*. 2003; 982:64–78. [PubMed: 12915241]
- Dégallier S, Ijspeert A. Modeling discrete and rhythmic movements through motor primitives: a review. *Biological Cybernetics*. 2010; 103:319–338.10.1007/s00422-010-0403-9 [PubMed: 20697734]
- Diedrichsen J. Optimal task-dependent changes of bi-manual feedback control and adaptation. *Current Biology*. 2007; 17:1675–1679.10.1016/j.cub.2007.08.051 [PubMed: 17900901]
- Diedrichsen J, Gush S. Reversal of bimanual feedback responses with changes in task goal. *Journal of Neurophysiology*. 2009; 101:283–288.10.1152/jn.90887.2008 [PubMed: 18987120]
- Dijkstra TMH, Katsumata H, de Rugy A, Sternad D. The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies*. 2004; 11:319–344.
- Guckenheimer, J.; Holmes, PJ. *Nonlinear oscillations, dynamical systems and bifurcations of vector fields*. New York: Springer-Verlag; 1986.
- Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature*. 1998; 394:780–784.10.1038/29528. [PubMed: 9723616]
- Hogan N, Sternad D. On rhythmic and discrete movements: Reflections, definitions and implications for motor control. *Experimental Brain Research*. 2007; 181:13–30.
- Holmes PJ. The dynamics of repeated impacts with a sinusoidally vibrating table. *Journal of Sound Vibration*. 1982; 84:173–189.
- Izawa J, Rane T, Donchin O, Shadmehr R. Motor adaptation as a process of reoptimization. *Journal of Neuroscience*. 2008; 28:2883–2891.10.1523/JNEUROSCI.5359-07.2008 [PubMed: 18337419]
- Liu D, Todorov E. Evidence for the flexible sensori-motor strategies predicted by optimal feedback control. *Journal of Neuroscience*. 2007; 27:9354–9368.10.1523/JNEUROSCI.1110-06.2007 [PubMed: 17728449]
- Morice AH, Siegler IA, Bardy BG, Warren WH. Learning new perception-action solutions in virtual ball bouncing. *Experimental Brain Research*. 2007; 181:249–265.
- Pham Q-C, Hicheur H. On the open-loop and feedback processes that underlie the formation of trajectories during visual and nonvisual locomotion in humans. *Journal of Neurophysiology*. 2009; 102:2800–2815.10.1152/jn.00284.2009 [PubMed: 19741106]
- Platt JR. Strong inference: Certain systematic methods of scientific thinking may produce much more rapid progress than others. *Science*. 1964; 146:347–353.10.1126/sci-ence.146.3642.347 [PubMed: 17739513]
- Popper, KR. *The logic of scientific discovery*. New York: Basic Books; 1959.
- Reist, P.; D’Andrea, R. Bouncing an unconstrained ball in three dimensions with a blind juggling robot. *Robotics and Automation. ICRA 2009 IEEE International Conference*; 2009. p. 1774-1781.
- Ronsse R, Lefèvre P, Sepulchre R. Sensorless stabilization of bounce juggling. *IEEE Transactions on Robotics*. 2006; 22:147–159.
- Ronsse R, Lefèvre P, Sepulchre R. Rhythmic feedback control of a blind planar juggler. *IEEE Transactions on Robotics*. 2007; 23:790–802.
- Ronsse R, Lefèvre P, Sepulchre R. Robotics and neuroscience: A rhythmic interaction. *Neural Networks*. 2008; 21:577–583.10.1016/j.neunet.2008.03.005 [PubMed: 18490135]
- Ronsse R, Miall RC, Swinnen SP. Multisensory integration in dynamical behaviors: maximum likelihood estimation across bimanual skill learning. *Journal of Neuroscience*. 2009; 29:8419–8428.10.1523/JNEUROSCI.5734-08.2009 [PubMed: 19571132]

- Ronsse R, Thonnard J-L, Lefevre P, Sepulchre R. Control of bimanual rhythmic movements: Trading efficiency for robustness depending on the context. *Experimental Brain Research*. 2008; 187:193–205.10.1007/s00221-008-1297-9
- Ronsse R, Wei K, Sternad D. Optimal control of a hybrid rhythmic-discrete task: The bouncing ball revisited. *Journal of Neurophysiology*. 2010; 103:2482–2493.10.1152/jn.00600.2009 [PubMed: 20130042]
- Schaal S, Atkeson CG, Sternad D. One-handed juggling: A dynamical approach to a rhythmic movement task. *Journal of Motor Behavior*. 1996; 28:165–183. [PubMed: 12529218]
- Schaal S, Sternad D, Osu R, Kawato M. Rhythmic arm movement is not discrete. *Nature Neuroscience*. 2004; 7:1136–1143.
- Scott SH. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*. 2004; 5:532–546.
- Scott SH. Inconvenient truths about neural processing in primary motor cortex. *Journal of Physiology*. 2008; 586:1217–1224.10.1113/jphysiol.2007.146068 [PubMed: 18187462]
- Shadmehr R, Krakauer JW. A computational neuroanatomy for motor control. *Experimental Brain Research*. 2008; 185:359–381.10.1007/s00221-008-1280-5
- Siegler IA, Bardy BG, Warren WH. Passive vs. active control of rhythmic ball bouncing: the role of visual information. *Journal of Experimental Psychology: Human Perception and Performance*. 2010; 36:729–750.10.1037/a0016462. [PubMed: 20515200]
- Srinivasan M, Ruina A. Computer optimization of a minimal biped model discovers walking and running. *Nature*. 2006; 439:72–75.10.1038/nature04113 [PubMed: 16155564]
- Sternad D. Juggling and bouncing balls: Parallels and differences in dynamic concepts and tools. *International Journal of Sport Psychology*. 1999; 30:462–489.
- Sternad D, Duarte M, Katsumata H, Schaal S. Bouncing a ball: Tuning into dynamic stability. *Journal of Experimental Psychology: Human Perception and Performance*. 2001a; 27:1163–1184. [PubMed: 11642701]
- Sternad D, Duarte M, Katsumata H, Schaal S. Dynamics of a bouncing ball in human performance. *Physical Review E*. 2001b; 6301:011902.
- Tassa, Y.; Erez, T.; Smart, W. Receding horizon differential dynamic programming. In: Platt, J.; Koller, D.; Singer, Y.; Roweis, S., editors. *Advances in neural information processing systems*. Vol. 20. Cambridge, MA: MIT Press; 2008. p. 1465-1472.
- Todorov E. Optimality principles in sensorimotor control. *Nature Neuroscience*. 2004; 7:907–915.
- Todorov E. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*. 2005; 17:1084–1108.10.1162/0899766053491887. [PubMed: 15829101]
- Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*. 2002; 5:1226–1235.
- Tufillaro, NB.; Abbott, T.; Reilly, J. *An experimental approach to nonlinear dynamics and chaos*. Redwood City, CA: Addison-Wesley; 1992.
- Tufillaro NB, Albano AM. Chaotics dynamics of a bouncing ball. *American Journal of Physics*. 1985; 54:939–944.
- Van Der Wel RPRD, Sternad D, Rosenbaum DA. Moving the arm at different rates: Slow movements are avoided. *Journal of Motor Behavior*. 2010; 42:29–36.10.1080/00222890903267116
- Wei K, Dijkstra TMH, Sternad D. Passive stability and active control in a rhythmic task. *Journal of Neurophysiology*. 2007; 98:2633–2646. [PubMed: 17881482]
- Wei K, Dijkstra TMH, Sternad D. Stability and variability: Indicators for passive stability and active control in a rhythmic task. *Journal of Neurophysiology*. 2008; 99:3027–3041. [PubMed: 18353911]
- Wiesenfeld K, Tufillaro NB. Suppression of period doubling in the dynamics of a bouncing ball. *Physica D*. 1987; 26:321–335.
- Zavala-Rio A, Brogliato B. On the control of a one degree-of-freedom juggling robot. *Dynamics and Control*. 1999; 9:67–90.

Zavala-Rio A, Brogliato B. Direct adaptive control design for one-degree-of-freedom complementary-slackness jugglers. *Automatica*. 2001; 37:1117–1123.

\$watermark-text

\$watermark-text

\$watermark-text

FIGURE 1.

(A) Sketch of the ball-bouncing task: The participant has to propel the ball to a target height by impacting it with a racket. (B) Sketch of the virtual reality setup (i.e., there is no physical ball) for the task together with a front view of the screen display.

FIGURE 2.

Racket acceleration at impact $\ddot{r}_*$ across a series of bouncing trials (30 s each). Each data point represents the average of six participants' trial means. The solid line represents the exponential fit with the equation and its R^2 shown at the top. Reproduced from Sternad et al. (2001b).

FIGURE 3.

Response to perturbations of 14 different magnitudes. **(A)** Grand averages of height error (HE) plotted as a function of cycle number (cycle 0 is the perturbed cycle). The 8 panels show the data for 8 of 14 perturbation magnitudes: $P - n$ are negative perturbations (ball post-impact velocity was lower than expected), and $P + n$ are positive perturbations (ball post-impact velocity was higher than expected), the larger n , the larger the perturbation magnitude. Four different α conditions are shown by different gray shades. The error bars denote the average of the seven individuals' standard deviations across trials. **(B)** Relaxation times t obtained from the exponential fits of HE . Different α conditions are shown by different lines and symbols. The gray shaded box indicates that the t values for the smallest perturbations have to be interpreted with caution, since the postperturbation variables did not significantly differ from steady-state behavior. Although the relaxation times t were significantly faster than predicted by the model, the qualitative pattern of t as a function of perturbation magnitude corresponded to the basin of attraction. Adapted with permission from Wei et al. (2007).

FIGURE 4.

The slopes of the regressions of the racket velocity on the ball height error at the previous bounce (V/HE gain), for both the model and the data. **(A)** Average data of the four participants from the unexperienced group (open circles) and the simulated predictions using the parameters from the four individuals (crosses). Error bars denote the standard deviations across trials averaged over the four participants in this group. **(B)** The same data and simulations from the experienced group (four participants). Reprinted with permission from Wei et al. (2008).

FIGURE 5.

Temporal profile of the optimal gains L_i . At the start, the racket reinitializes its position, by returning close to the rest position. Later, the racket performs an impulse to reach the desired impact state. These gains were obtained for g_4 , but look similar across the seven gravity conditions.

FIGURE 6.

Mean trajectories of the racket position over the normalized time in seven gravity conditions, corresponding to seven different cycle periods (g_1 corresponds to the longest period, about 2.8 s; and g_7 —the nominal gravity—corresponds to the shortest period, about 0.7 s). The racket kinematics have been resampled between two consecutive impacts (represented by the two vertical lines). **(A)** Racket trajectories generated by the model for all seven gravity conditions. **(B)** Averaged racket trajectories of the human participants (eight) in the same seven gravity conditions. Adapted with permission from Ronsse et al. (2010).

FIGURE 7.

Regression analyses. **(A)** Regression slopes between height error HE and the racket velocity at the next impact V (V/HE gain) for the human data (solid line) and simulation results from the model (dashed line). **(B)** Lag -1 autocorrelation coefficient $AC-1$ of the height error HE at impact k and $k+1$ in the human data (solid line) and in simulation results (dashed line). The error bars denote the standard deviations of the human data (8 participants). Reprinted with permission from Ronsse et al. (2010).

FIGURE 8.

Regression slopes obtained from regressions between the actual and predicted racket state at impact, for all gravity conditions g_i ([A] position; [B] velocity). The error bars denote the standard deviations for eight participants.

FIGURE 9.

(A) Three-dimensional representation of ball-racket contacts of one trial: racket impact velocity r_{k+1} as a function of the racket impact position r_{k+1} and the ball height during the preceding flight $b_{H,k}$. Real data are shown as black circles, model data are shown as white squares. The two-dimensional solution manifold (with $b_{H,k+1} = b_H^*$) is represented in gray. (B) Section of space with the solution manifold: the lines represent the intersection of the solution manifold with different impact positions, $r_{k+1} = 0.17\text{m}$ (solid line) and $r_{k+1} = 0.12\text{m}$ and $r_{k+1} = 0.22\text{m}$ (dotted lines). Data were taken from a representative participant, and for g_7 (nominal gravity).