J Neurophysiol 103: 2482–2493, 2010.
First published February 3, 2010; doi:10.1152/jn.00600.2009.

Optimal Control of a Hybrid Rhythmic-Discrete Task: The Bouncing Ball Revisited

Renaud Ronsse,[1,2] Kunlin Wei,[3,4] and Dagmar Sternad[5]

[1]*Department of Electrical Engineering and Computer Science, Montefiore Institute, Université de Liège, Liège, Belgium;* [2]*Department of Biomedical Kinesiology, Katholieke Universiteit Leuven, Heverlee, Belgium;* [3]*Rehabilitation Institute of Chicago, Northwestern University, Chicago, Illinois;* [4]*Center for Brain and Cognitive Sciences and Department of Psychology, Peking University, Beijing, China; and* [5]*Departments of Biology, Electrical and Computer Engineering, and Physics, Northeastern University, Boston, Massachusetts*

Submitted 14 July 2009; accepted in final form 30 January 2010

Ronsse R, Wei K, Sternad D. Optimal control of a hybrid rhythmic-discrete task: the bouncing ball revisited. *J Neurophysiol* 103: 2482–2493, 2010. First published February 3, 2010; doi:10.1152/jn.00600.2009. Rhythmically bouncing a ball with a racket is a hybrid task that combines continuous rhythmic actuation of the racket with the control of discrete impact events between racket and ball. This study presents experimental data and a two-layered modeling framework that explicitly addresses the hybrid nature of control: a first discrete layer calculates the state to reach at impact and the second continuous layer smoothly drives the racket to this desired state, based on optimality principles. The testbed for this hybrid model is task performance at a range of increasingly slower tempos. When slowing the rhythm of the bouncing actions, the continuous cycles become separated into a sequence of discrete movements interspersed by dwell times and directed to achieve the desired impact. Analyses of human performance show increasing variability of performance measures with slower tempi, associated with a change in racket trajectories from approximately sinusoidal to less symmetrical velocity profiles. Matching results of model simulations give support to a hybrid control model based on optimality, and therefore suggest that optimality principles are applicable to the sensorimotor control of complex movements such as ball bouncing.

INTRODUCTION

Research in motor control to date has mainly focused on two broad classes of behaviors: discrete movements and rhythmic movements. Rhythmic movements like all forms of locomotion are continuous and repetitive, whereas discrete movements like stepping over an obstacle, are bounded by well-defined start and end positions (Hogan and Sternad 2007). The research communities studying rhythmic and discrete movements have been largely distinct and used different experimental paradigms and different theoretical frameworks. Rhythmic movements, particularly animal locomotion, have been shown to result from spinal mechanisms that create periodic patterns in an autonomous fashion, largely independent of higher levels of the CNS. Thus nonlinear oscillators have been predominant in modeling these central pattern generators identified in the spinal cord of many species (Ijspeert 2008). In contrast, for discrete actions imaging work provided evidence that even simple point-to-point movements involve higher cortical areas and therefore may require fundamentally more complex control mechanisms (Schaal et al. 2004). Thus for discrete movements computational frameworks such as optimal feedback control have been invoked (Shadmehr and Wise 2005).

Address for reprint requests and other correspondence: D. Sternad, Dept. of Biology, 134 Mugar Life Science Bldg., Northeastern Univ., Boston, MA 02115 (E-mail: dagmar@neu.edu).

This separation between studies on rhythmic and discrete movements defies the fact that many, if not most, daily activities are a combination of rhythmic and discrete elements. Although walking on even terrain can be executed without much attention and probably without extensive higher-level control, as soon as there are obstacles that perturb the rhythmic pattern, steps have to be placed in a discrete reaching-like fashion. Similarly, piano playing is regarded as an archetypal rhythmic action, yet simultaneously integrates discrete and very accurate reaches across the keyboard. Coordination and control of such hybrid actions may therefore pose more complex problems than each action by itself and probably engage overlapping circuitries and mechanisms (Ronsse et al. 2009; Sternad 2008; Sternad et al. 2000).

The present study examines such a hybrid task with a corresponding two-level approach: using a rhythmic hand–object interaction, bouncing a ball with a racket, we focus on the discrete-like elements of the task that recur in the continuous periodic action. To enhance the discrete component we manipulate the task such that continuous performance breaks down to a sequence of discrete-like actions. Such a transition can be induced by slowing down the tempo of the task such that the continuous rhythmic movement falls apart into a sequence of discrete movements separated by dwell times. This change is expected based on previous studies on bimanual bouncing of a puck, which demonstrated that the racket trajectory changed from a sustained rhythmic actuation to a train of discrete movements as the overall tempo was decreased (Ronsse et al. 2008b). Even though the ball–racket contacts continue to recur in an approximately periodic fashion, each single contact event now requires initiation and termination similar to a target-oriented reach, the hallmarks of discrete performance (Hogan and Sternad 2007).

The same ball bouncing task has been the subject of a series of studies by Sternad and colleagues (Dijkstra et al. 2004; Schaal et al. 1996; Sternad et al. 2001a,b). Results of these studies provided evidence that skilled subjects adopt a performance strategy that established dynamic stability, a regime in which deviations from the intended ball amplitude self-stabilize without involving corrective motor actions. These conclusions were derived from a mathematical model of the task where stability analyses of the periodic behavior provided support that subjects exploited stability of the rhythmic task. Establishing dynamic stability in the ball–racket system makes extensive error corrections unnecessary and makes control less costly (Ronsse et al. 2008a). This is consistent with the sub-

 0022-3077/10 Copyright © 2010 The American Physiological Society

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

jective observation that in skilled performance less attention is needed to maintain a stable pattern. This same approach is used to generate robust execution of locomotion in legged robots. The first and probably most intriguing example is the passive dynamic walker by McGeer (1990) that can travel down a gentle slope and walk stably in two dimensions without any other external source of energy.

Scrutinizing the results on ball bouncing further, more recent investigations revealed that dynamic stability alone is not sufficient to account for human performance (Wei et al. 2007, 2008). When perturbations were applied, subjects' behavior showed that active error-correction processes were recruited to stabilize the pattern. Specifically, the experimental data showed that variations in the ball trajectory (not only induced by explicit perturbations but also resulting from natural variability) can lead to additional error-correcting mechanisms in the racket movement, resulting in compensation of errors faster than expected from the intrinsic dynamic stability. In sum, humans use this "passive" stability offered by the dynamics of the rhythmic movement but also flexibly integrate feedback-based "active control" into the passive stability: the two control modes are blended together (Wei et al. 2007, 2008). However, so far no explicit modeling of the active component has been offered.

The present study aims to develop a control structure to account for such active control. In extension of studies on discrete reaching we adopt an optimal feedback control framework (Todorov and Jordan 2002). However, as pointed out earlier, the task has a hybrid character in that it combines discrete and rhythmic elements. Therefore we extend the current models that have been developed for discrete reaching movements (Diedrichsen 2007; Diedrichsen and Gush 2009; Liu and Todorov 2007; Todorov 2004, 2005; Todorov and Jordan 2002). This study proposes a control architecture that is divided into two layers. A discrete layer is used to determine the desired state for the next impact based on dynamic requirements tuned by available perceptual information. The second layer of control generates the racket trajectory toward this desired impact state at the appropriate impact time (Ronsse et al. 2007), to minimize a given cost function. This continuous-time layer is embedded into the discrete-time layer and thereby influences the discrete dynamics at impact.

With the goal of generating test cases that highlight the interplay between the rhythmic and discrete elements of the task we examine unimanual bouncing of a ball under different rhythmic tempi. Based on the study by Ronsse et al. (2008b) we expect that subjects—as the task's tempo decreased—will no longer show continuously smooth trajectories but rather intermittent dwelling between two contact events. To elicit such slow rhythmic actions, we used a virtual environment: instead of the obvious manipulation of increasing the target height (which poses difficulties in the lab), the virtual setup afforded manipulation of the gravity acting on the ball and thereby significantly slowing the movement. By manipulating gravity, the bounce height could be kept constant.

In sum, the present study examines the task of rhythmically bouncing a ball with the goal of eliciting significant departures from continuous rhythmic movements. We thereby create performance that is a hybrid of rhythmic and discrete movements. We propose an optimal control model to account for the data and demonstrate that this framework successfully accounts for performance of such a hybrid task.

METHODS

Experiment

PARTICIPANTS. Eight volunteers (three females, five males) participated, with ages ranging from 25 to 48 yr. All participants self-reported to be right-handed and used their preferred right hand to bounce the ball with the racket. Before the experiment, all participants were informed about the procedure and signed the consent form approved by the ethics committee of the Pennsylvania State University.

EXPERIMENTAL APPARATUS. For this study, we used the same virtual setup that was already described elsewhere (Wei et al. 2007). This is a two-dimensional (2D) virtual reality setup, in which participants manipulated a real table tennis racket to bounce a virtual ball that is projected onto a screen in front of them (Fig. 1). Participants stood 0.50 m behind the back-projection screen. A PC (2.4-GHz Pentium CPU, Windows XP) controlled the experiment and generated the visual stimuli with a graphics card (Radeon 9700, ATI). The same PC also acquired the data using a 16-bit A/D card (DT322, DataTranslation). The images were projected by a Toshiba TLP 680 TFT-LCD projector and consisted of 1,024 $\times$ 768 pixels with a 60-Hz refresh rate. To measure the vertical displacements a light rigid rod with three hinge joints was attached to the bottom of the racket surface and ran through a wheel whose rotations were registered by an optical encoder. Its accuracy was one pulse for 0.27 mm of racket movement, counted by an onboard counter (DT322). The racket could move and tilt with minimal friction in three dimensions, but only vertical displacement was measured. Given that the ball movements were confined to the vertical direction, racket displacements in other than vertical directions were negligible. Images of racket and ball position were shown on-line on the back-projection screen using custom-made software. The delay between real and virtual racket movement was measured in a separate experiment and was on average 22 $\pm$ 0.5 (SD) ms.

The virtual racket was displayed at the same height from the floor as that of the real racket and its displacement was the same as that of

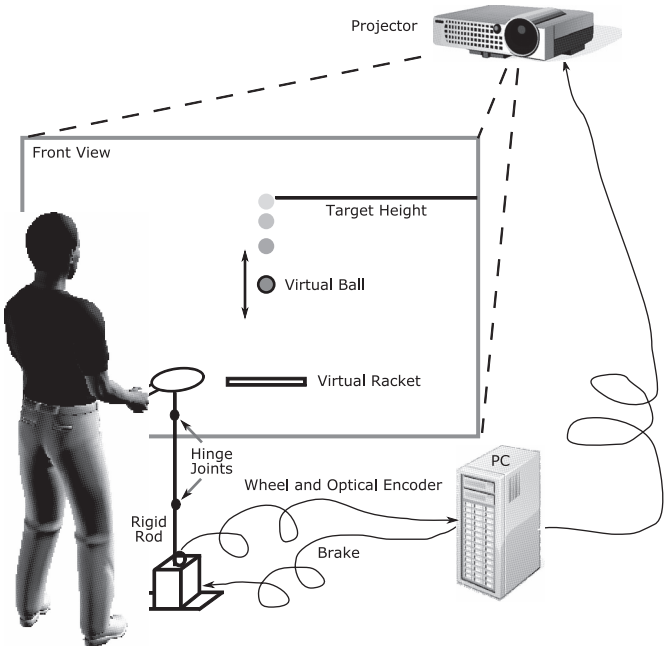

FIG. 1. Sketch of the virtual reality setup for the ball bouncing task together with a front view of the screen display.

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

the real racket (the visuomotor gain was equal to 1). The map for the discrete impacts was based on two assumptions that were implemented in the virtual setup: *1*) between impacts the movement of the ball displayed on the screen follows ballistic flight under the influence of gravity; and *2*) the impacts are instantaneous with the coefficient of restitution capturing the energy loss at impact. The first assumption implies that the ball trajectory between impacts k and $k + 1$ follows a parabola

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

where b_k and $\dot{b}_k$ denote the ball position and velocity immediately after impact k, respectively. *Equation 1* holds for $t_k \leq t \leq t_{k+1}$, where t_k and t_{k+1} denote successive impact times. From *Eq. 1,* the ball velocity obeys $\dot{b}(t) = \dot{b}_k - gt$. The ball apex corresponds to the highest point of the ball trajectory between two impacts. At this point, ball velocity equals zero at the time $t_{apex,k} = \dot{b}_k/g$. The apex position follows from *Eq. 1*

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

The second assumption is that the racket–ball collisions are governed by the impact law

$$(\dot{b}_k - \dot{r}_k) = -\alpha(\dot{b}_k^- - \dot{r}_k)$$
$$\Leftrightarrow \dot{b}_k = -\alpha\dot{b}_k^- + (1 + \alpha)\dot{r}_k \quad (3)$$

where $\dot{r}_k$ is the racket velocity at impact k and $\dot{b}_k^-$ is the ball velocity just before the same impact. The coefficient of restitution α captures the energy dissipation at each impact. The experiments reported here used the single coefficient of restitution $\alpha = 0.6$.

Just before the virtual ball hit the virtual racket, a trigger signal was sent out to a mechanical brake that acted on the rod. The brake was controlled by a solenoid and applied a brief braking force pulse to the rod to create the feeling of a real ball hitting the racket (Magnet-Schultz type R 16 × 16 DC pull, subtype S-07447). The trigger signal was sent 15 ms before the ball–racket contact to overcome the mechanical and electronic delay of the brake. The duration of the force pulse (30 ms) was consistent with the average impact duration observed in a real ball–racket experiment (Katsumata et al. 2003). The brake force was not scaled to the relative velocity of ball and racket but stayed the same for all impacts. Nevertheless, subjects reported the haptic sensation as sufficiently convincing.

To control the ball trajectory and the ball–racket contact the computer program reads the latest racket position from the optical encoder. When the racket was away from the ball, the program updated the ball positions based on the ballistic flight equation (update rate of ~800 Hz). When the ball and racket were close, the computer program would keep a running estimate of time to contact and control the brake accordingly (update rate of ~250 Hz). The update rate was not fixed because Windows XP is not a real-time operating system and thus timing is not deterministic. Thus all data were time-stamped using the high-resolution timer on the Pentium CPU with accuracy <1 ms.

PROCEDURE AND EXPERIMENTAL CONDITIONS. Prior to each experiment, the participant was placed on a support base to adjust for height differences. The support height was adjusted such that at the calibrated nominal position the forearm holding the racket was horizontal and the racket was 10 cm above ground in the virtual workspace. Each trial began with the ball appearing at the right side of the screen and rolling on a horizontal line extending to the center of the screen (see Fig. 1). On reaching the center of the screen, the ball dropped from the horizontal line (height: 0.80 m above ground in the virtual workspace). The task instruction was to rhythmically bounce the ball as accurately as possible to the target line (the same line from which the ball dropped) for the duration of one trial. The experiment consisted of a total of 21 trials per subject, which were collected in a single session lasting about 1 h.

The experiment was conducted in seven blocks, each with a different value of gravity acting on the ball dynamics (Table 1, with g_7 denoting the nominal gravity of 9.81 m/s^2). Each block consisted of three successive trials. For four participants, the gravity conditions were presented in ascending order from g_1 to g_7; the other four participants performed the gravity conditions in descending order. The different gravities were chosen such that the ball flight time produced seven different cycle periods, which are also listed in Table 1. At steady state where $b_k = b^*$ and $\dot{b}_k = \dot{b}^*$, $\forall k$, the interval between two successive impacts, the period T, derives from *Eq. 1*

$$T = 2\frac{\dot{b}^*}{g} \quad (4)$$

Thus T depends only on the ball velocity at release $\dot{b}^*$ and gravity g. By eliminating $\dot{b}^*$ from *Eqs. 2* (at steady-state) and *4,* we find

$$T = \sqrt{\frac{8}{g}(b^*_{apex} - b^*)} \quad (5)$$

which captures how g and the steady-state apex position b^*_{apex} define the impact-to-impact period T for a fixed steady-state impact position b^*. Consequently, assuming $b^*_{apex} - b^* \simeq 0.8$ m, the seven values of g tested in this experiment correspond to the steady-state cycle durations that are given in Table 1.

The gravity values were chosen such that the corresponding steady-state flight times of the ball grew linearly. Each trial stopped after 40 cycles (40 impacts); thus the trial durations depended on the gravity condition: the smaller the gravity value, the longer the trial. None of the subjects reported fatigue as an issue, not even in the longer trials.

DATA REDUCTION AND FILTERING. The first five cycles of each trial were eliminated from the analysis to avoid transients to bias average estimates. The raw data of the racket displacement were resampled at a fixed frequency of 500 Hz and filtered by a fourth-order Savitzky–Golay filter with a window size of 0.01 s on both sides. Compared with conventionally used filters like the Butterworth filter, the Savitzky–Golay filter is superior for smoothing data that have abrupt changes. Abrupt changes occurred at contact when the racket exhibited a sudden drop in acceleration, caused by the brake. The filter order and window size were chosen empirically to remove the measurement noise while not excessively smoothing the signals. The displacements of the ball were generated by the computer and thus contained no measurement noise. Therefore no filtering was applied.

DEPENDENT VARIABLES AND STATISTICAL ANALYSES. Figure 2 shows representative trajectories of ball and racket. Performance was evaluated by the ball height error *HE,* defined as the signed difference between the maximum ball height (the apex) and the target height shown by the line. The cycle period T was defined as the time interval between successive impacts.

To assess changes in the shape of the racket trajectories in the different gravity conditions, the following variables were computed: *1*) the racket amplitude *RA* was defined as the difference between the highest and the lowest racket position within a cycle; and *2*) the activity period *AP* was the percentage of cycle during which the absolute value of the racket velocity was above 5% of its peak value.

TABLE 1. *Gravity acting on the ball dynamics and the corresponding steady-state cycle periods T*

Gravity, m/s^2	$g_1 = 0.61$	$g_2 = 0.80$	$g_3 = 1.09$	$g_4 = 1.57$	$g_5 = 2.45$	$g_6 = 4.36$	$g_7 = 9.81$
Cycle periods, ms	$T_1 \simeq 2{,}800$	$T_2 \simeq 2{,}450$	$T_3 \simeq 2{,}100$	$T_4 \simeq 1{,}750$	$T_5 \simeq 1{,}400$	$T_6 \simeq 1{,}050$	$T_7 \simeq 700$

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

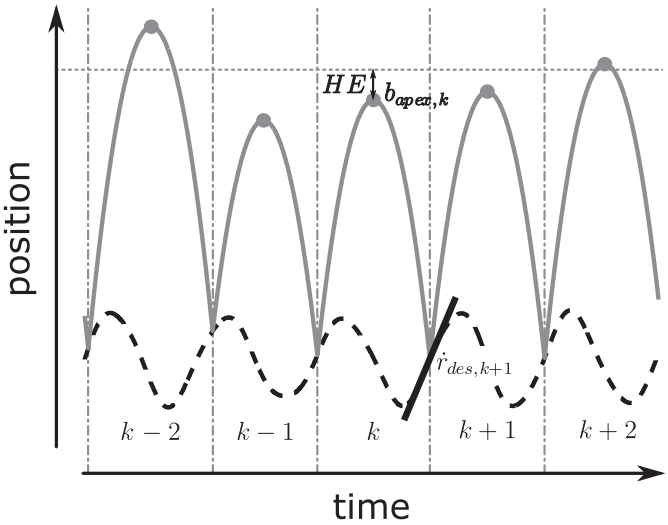

FIG. 2. Representative time series of the racket (dashed black) and ball (solid gray) trajectories over five cycles (from $k - 2$ to $k + 2$). In the model the ball apex during cycle k, $b_{apex,k}$, is evaluated in order to determine the desired racket velocity at the next impact, i.e. $\dot{r}_{des,k+1}$ (black solid line). The height error *HE* is the signed difference between b_{apex} and the target height.

For all four dependent variables, *HE, T, RA,* and *AP,* the means and SDs were calculated over the 35 impacts within a trial. Subsequently, averages across trials and participants were calculated per condition.

One important question is whether a perceived variable is used to control or correct an action variable. Such compensatory adjustments were evaluated by performing linear regressions between variables. In previous work, the most insightful regression proved to be the one of the height error *HE* in cycle *k* (the perceived variable) on the racket velocity at impact $k + 1$ (the controlled action variable): the significant nonzero slope suggested that the larger the *HE* error, the smaller was the velocity, speaking to error corrections (Ronsse et al. 2008b; Wei et al. 2008). This slope will be referred to as the velocity-height *V/HE* gain (see Fig. 2). For a standard proportional controller this slope can be interpreted as a feedback gain: the higher this gain, the faster the perturbation is attenuated (up to the closed-loop stability limit). Note, that Wei et al. (2008) showed that also the stochastic open-loop model produced such correlation due to its self-stabilization. However, the human data showed higher gains than those predicted by the model. Thus only values higher than the predicted values could be interpreted as feedback gains. The present data were analyzed in the same fashion: for each trial all values of racket velocity at contact were linearly regressed against the immediately preceding height error *HE.* The results of all trials in each gravity condition were averaged across subjects.

In addition, time correlations of height errors *HE* were evaluated by computing lag-1 autocorrelations, denoted as *AC*-1. *AC*-1 values were calculated for all trials per condition and then averaged across trials and participants.

ANOVAs were performed on all dependent measures, both on means and SDs. The design was a one-way ANOVA with gravity condition as the within-subject factor. When required, main effects were further analyzed using Tukey's Honestly Significant Difference post hoc tests.

Mathematical model of the hybrid control architecture

The control architecture is divided into two layers: First, because the ball dynamics can be influenced only at impact, the first layer of control is discrete in time and computes the desired racket velocity at the next impact position (Fig. 3*A*). Second, since the racket movements are continuous, the second layer of control aims to drive the racket to the specified target position and velocity at the desired time (Fig. 3*B*).

The discrete layer of control is restricted to a discrete state space of three variables: the racket position and velocity to reach at the next impact to bring the ball to the desired height and the time at which this impact will happen (Ronsse et al. 2007). Impact can occur at any point along the ball trajectory or, stated differently, after any positive flight duration. The aim of the discrete layer is to compute the optimal feedback control policy, minimizing a given cost function for possible cycle durations. From *Eqs. 2* and *3,* it can be shown that the desired racket impact velocity at impact *k* that bounces the ball to the target apex b^*_{apex} during the next cycle is equal to

$$\dot{r}_{des,k} = \frac{\alpha \dot{b}_k^- + \sqrt{2g(b^*_{apex} - b_k)}}{1 + \alpha} \tag{6}$$

depending on the impact position b_k and ball velocity before impact $\dot{b}_k^-$ and, consequently, on the cycle duration since the ball trajectory is deterministic (see *Eq. 1*). The continous layer defines the optimal control policy that minimizes the cost function

$$Q = \|r_{k+1} - r_{des,k+1}\|^2 + \|\dot{r}_{k+1} - \dot{r}_{des,k+1}\|^2 + w_{rest}\int_{t_k}^{t_{k+1}} \|r(t) - r_0\|^2 dt + w_{energy}\int_{t_k}^{t_{k+1}} \|u(t)\|^2 dt \tag{7}$$

over possible cycle durations $\Delta t = t_{k+1} - t_k$. In *Eq. 7* $u(t)$ represents the control input driving the racket. The cost function balances the need to complete the task objectives (reach a specific racket state at impact) and to minimize control. The first two terms are minimized if the task is accomplished, i.e., if the racket approaches the desired position—which is identical to the ball position—and velocity at impact. These two terms are determined from the corresponding cycle duration via *Eqs. 1* and *6.* The last term penalizes the control input $u(t)$ integrated over the entire cycle weighted by w_{energy}. Additionally, the third term penalizes large excursions away from the racket rest position r_0, representing the biomechanical and neurophysiological rest position of the spring/reflex system of forearm and elbow muscles. 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 sticks to it. Instead, it leads to the tendency to keep the elbow at rest as much as possible. This term also reproduces dwell intervals observed at smaller gravity values in the actual data. The two parameters w_{rest} and w_{energy} are relative weights of the corresponding components.

For the continuous-time implementation in the second layer, the racket and forearm were modeled as a rigid mechanical system obeying

$$I\ddot{r}(t) + \gamma\dot{r}(t) = f(t) \tag{8}$$

where *I* denotes the forearm and racket inertia and γ denotes the elbow's damping coefficient (de Rugy et al. 2003). We assume that the elbow's angular displacement is proportional to the racket position $r(t)$ since the displacements are sufficiently small to allow this linearization; $f(t)$ represents the net muscular force produced by the elbow flexor and extensor, which is produced by the first-order low-pass filtered input

$$\tau\dot{f}(t) + f(t) = u(t) + cu(t)\dot{\sigma}(t) \tag{9}$$

where τ represents a time constant. The second term on the right-hand side involving $\dot{\sigma}(t)$ represents signal-dependent noise acting on the input $u(t)$, where $\sigma(t)$ is standard Brownian motion (Harris and Wolpert 1998; Todorov 2005; Todorov and Jordan 2002). The noise strength is scaled by the parameter *c*. The racket variables can thus be assembled into a five-dimensional state vector

$$\mathbf{x}(t) = [r(t), \dot{r}(t), f(t), r_{des,k+1}, \dot{r}_{des,k+1}]$$

The parameter values were taken from other studies: $I = 0.05$ $N{\cdot}m^{-1}{\cdot}s^{-2}$, $\gamma = 0.25$ $N{\cdot}m^{-1}{\cdot}s^{-1}$ (de Rugy et al. 2003), and $\tau = 50$ ms (Liu and Todorov 2007). To simulate the effect of the ball-racket

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

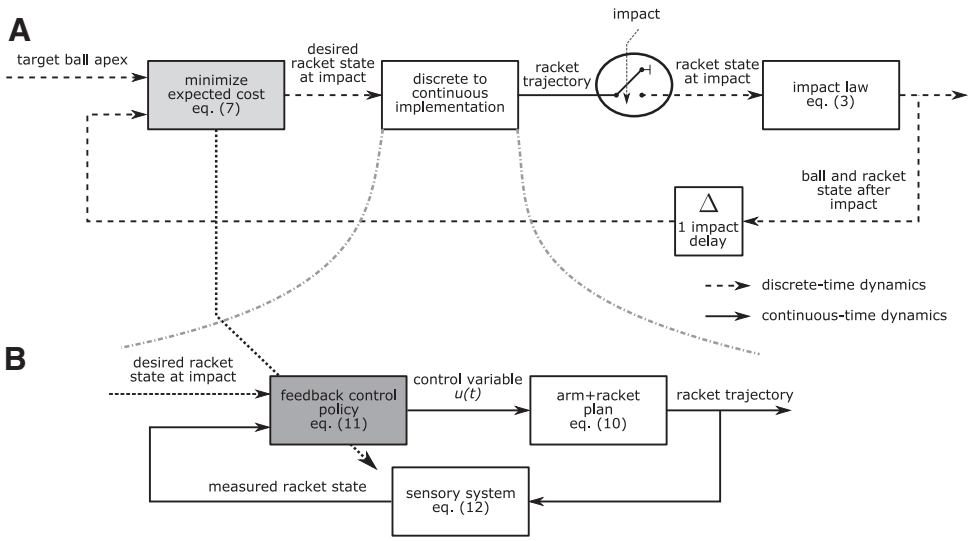

FIG. 3. Two-layer controller for the hybrid task. *A*: discrete dynamics between impacts: the discrete-time controller (light gray box) computes the racket state to reach at the next impact that drives the ball to the desired height. The impact symbol (ellipse) reflects that only the racket state at impact influences the discrete-time dynamics. *B*: the implementation of the racket trajectory is continuous in time and is achieved by the optimal feedback control policy (dark gray box). It uses the measured state evolution of the racket to generate the control variable $u(t)$, which drives the racket to the desired impact state. The corresponding equations of the model are referenced in the figure boxes.

contact, the damping of the racket γ in *Eq. 8* was multiplied by 10 during the first 30 ms after impact, yielding $\gamma = 2.5\ N{\cdot}m^{-1}{\cdot}s^{-1}$.

The mechanical *Eqs. 8* and *9* were discretized with a time step of 2 ms, so that the optimal control input $u(t)$ could be computed (Diedrichsen 2007; Liu and Todorov 2007; Todorov 2005). The discretized equations were written in matrix form

$$\mathbf{x}_{t+1} = \mathrm{A}\mathbf{x}_t + \mathrm{B}(u_t + cu_t\varepsilon_t) \qquad (10)$$

where $\mathbf{x}_t$ is the state vector and u_t the input at time step t. The variable ε_t is the discretized version of Brownian motion, i.e., an independent random variable from a Gaussian distribution with mean 0 and variance 1. It is important to point out that the discretization of the first layer (impact-to-impact, or cycle-to-cycle)—stamped by the variable k—must not be confused with this time discretization t used here.

The Linear-Quadratic-Gaussian (LQG) algorithm found the time-varying feedback control gain vector

$$u_t = -L_t\mathbf{x}_t \qquad (11)$$

where the time-varying control gains L_t were computed such that the cost function (*Eq. 7*) was minimized. Moreover, the control input u_t was determined from the estimate $\hat{\mathbf{x}}_t$ of the true state vector $\mathbf{x}_t$. This estimate was equal to $\hat{\mathbf{x}}_t = [\hat{r}_t, \hat{\dot{r}}_t, \hat{f}_t, \hat{r}_{des,k+1}, \hat{\dot{r}}_{des,k+1}]$. For simplicity, we assumed that the racket state could be measured without noise during the entire movement, such that

$$\hat{r}_t = r_t; \quad \hat{\dot{r}}_t = \dot{r}_t; \quad \hat{f}_t = f_t$$

The desired racket state to reach at impact k could be calculated up to a certain amount of internal noise (Todorov 2005)

$$\begin{pmatrix} \hat{r}_{des,k+1} \\ \hat{\dot{r}}_{des,k+1} \end{pmatrix} = \begin{pmatrix} r_{des,k+1} \\ \dot{r}_{des,k+1} \end{pmatrix} + \eta_t \qquad (12)$$

where η_t is an independent random matrix of a Gaussian distribution with mean 0 and covariance equal to

$$\begin{pmatrix} n_{pos} & 0 \\ 0 & n_{vel} \end{pmatrix}$$

The matrix η_t thus captures the sensory inaccuracy in estimating the ball's state and computing the corresponding racket state to reach at impact.

Since the system dynamics is linear, the terms of the cost function are quadratic, and the noise is Gaussian, the optimization problem can be solved using the LQG framework (see review in Todorov 2005). The first layer contains an inner loop seeking the optimal control policy for a given cycle duration and an outer loop determining the optimal duration associated with the smallest cost Q. This outer loop was implemented using a gradient-descent method. Once the control policy L_t associated with the lowest cost Q—corresponding to a particular state to reach at impact $(r_{des,k+1}, \dot{r}_{des,k+1})$—was found, the second layer determined the racket trajectory to reach this desired state at the optimal impact time (Fig. 3*B*).

In summary, the simulations were performed as follows:

1) The racket state was initialized at impact k

2) For a range of possible cycle durations $T = t_{k+1} - t_k$, the desired racket position and velocity to reach at t_{k+1} (impact) were calculated from *Eqs. 1* and *6* thus forming the state vector $\mathbf{x}$

3) For all possible cycle durations, the control policy L_t and the corresponding cost Q were computed, according to the algorithm developed in Todorov (2005)

4) The optimal control policy was chosen as the one associated with the smallest cost over possible durations

5) The observed state and the corresponding control gains were integrated according to *Eqs. 11* and *12,* respectively, with the gains L_t and the target state corresponding to the optimal control policy

6) When an impact occurred at time t_{k+1}, the ball's state was updated through *Eq. 3* and a new cycle started.

For each gravity condition, 846 impacts were simulated. The first six cycles were removed to obtain 840 impacts for each condition as in the actual data. This procedure was repeated for each of the seven gravity conditions of the experiment. The free parameters of the model are summarized in Table 2, together with their value and the rationale used to tune them.

Note that the proposed model is robust to unexpected perturbations of the racket and ball trajectories, since the first layer can be reinitialized as soon as a perturbation is detected. A systematic test of such perturbations, however, is beyond the scope of this study.

RESULTS

This section reports the human performance data followed by the corresponding model simulations.

Height error HE and movement period T

All subjects performed the task under seven different gravity conditions that resulted in different cycle periods because the target height was the same throughout: the smaller the gravity, the longer the cycle period. Figure 4 displays the ball and

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

TABLE 2. *Model parameters and rationales for setting values*

Parameter	Value	Rationale for Tuning
Coefficient of restitution	$\alpha = 0.6$	
Target apex height	$b^*_{apex} = 0.8$ m	
Zero position of the racket	$r_0 = 0.04$ m	
Arm + racket inertia	$I = 0.05\ \mathrm{N{\cdot}m^{-1}{\cdot}S^{-2}}$	From de Rugy et al. (2003)
Arm + racket damping	$\gamma = 0.25\ \mathrm{N{\cdot}m^{-1}{\cdot}S^{-1}}$	From de Rugy et al. (2003).
	$\gamma = 2.5\ \mathrm{N{\cdot}m^{-1}{\cdot}S^{-1}}$	To match the effect of the break around impact
Muscle time constant	$\tau = 50$ ms	From Liu and Todorov (2007)
Position penalty	$w_{rest} = 0.012\sqrt{g_i}$	To reproduce the actual racket trajectory (see Fig. 7)
Control penalty	$w_{energy} = 0.00002$	To reproduce the actual racket trajectory (see Fig. 7)
Signal-dependent noise strength	$c = 1.3$	To reproduce the actual background racket velocity variability (see Fig. 10)
Desired position uncertainty	$n_{pos} = 0.03$ m	To reproduce the actual racket position variability at impact (not shown)
Desired velocity uncertainty	$n_{vel} = 0.08$ m/s	To reproduce the actual racket velocity variability at impact (see Fig. 10)

racket trajectories for a representative subject. The trajectory in Fig. 4*A* was produced during the condition with normal gravity: 14 impacts occurred in 10 s. The trajectory in Fig. 4*B* is representative for the condition with the smallest gravity: only three impacts occurred in the same time interval.

To assess whether the subjects satisfied the instructions (i.e. repeatedly bounced the ball to the target height), Table 3 reports the average and SDs of the height error *HE*. The one-way ANOVA on *HE* means comparing different gravity conditions did not reach significance ($P > 0.28$). Table 3 reveals that subjects had a bias to slightly overshoot the target,

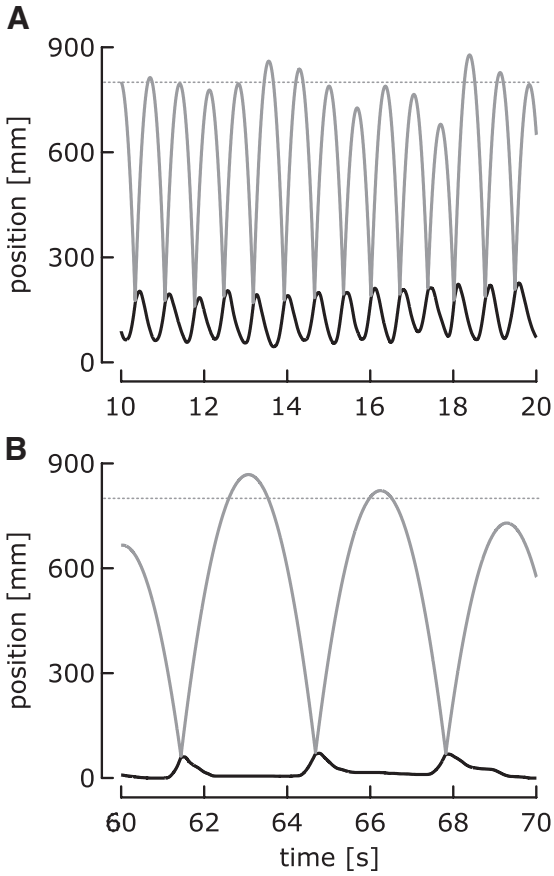

FIG. 4. Evolution of the racket (black) and ball (gray) trajectories over 10 s for a representative subject. The subject had to bounce the ball to the target height shown by a line such that higher gravity induced shorter cycle periods. *A*: trajectories in normal gravity condition g_7. *B*: trajectories under the smallest gravity condition g_1.

as already reported in previous studies (e.g., Wei et al. 2007). However, single-sample *t*-tests between *HE* and the instructed apex height $HE = 0$ were significant for only the one gravity condition g_2 ($P < 0.01$). More interestingly, the same ANOVA on the SDs of *HE* revealed a strong effect of the gravity condition: the smaller the gravity, the larger the SDs (Fig. 5). This result suggests that the bouncing pattern was more difficult to stabilize for smaller gravity values with its longer cycle periods. Post hoc Tukey's tests showed that most pairwise comparisons of two gravity conditions were significantly different ($P < 0.05$), except differences between two adjacent gravity values and between (g_3, g_5), (g_4, g_6), and (g_5, g_7). In sum, subjects reached the target with a slight overshoot and variability increased as gravity decreased.

The objective behind testing different gravities was to indirectly elicit different movement periods. Table 3 also summarizes the average and SDs of the movement periods *T* and confirms that this was indeed achieved. First, the one-way ANOVA on *T* established significant changes across gravity conditions [$F(6,49) = 1122.4$, $P < 0.0001$], strongly supported by post hoc Tukey's tests, which identified pairwise differences between any two gravity conditions ($P < 0.01$). Note that the actual periods *T* were not exactly identical to those listed in Table 1, since the average impact position (corresponding to the model steady-state b^*) was not exactly the same across gravity values. Furthermore, the standard deviations of *T* also depended on the gravity acting on the ball: the smaller the gravity, the more variable the period. Post hoc Tukey's tests identified significant differences between any two gravity conditions ($P < 0.05$), except between (g_1, g_2), (g_3, g_4), (g_5, g_6), (g_5, g_7), and (g_6, g_7). Note that Table 3 reports the variability in milliseconds; the effect was still present when comparing the SDs of *T* normalized by its corresponding average (coefficient of variation): for g_7, the SDs of *T* represent about 3% of the average period, whereas for g_1, this coefficient grew to 9%. In sum, the smaller the gravity acting on the ball, the slower and more variable the bouncing pattern.

Regression and correlation analyses

Wei et al. (2008) established that sufficiently large negative correlation between *HE* and racket velocity at the next impact is indicative of active control of the racket based on the perceived height error. Thus slopes of the regression of height error on the subsequent racket velocity at impact were interpreted as the feedback gain. These slopes, or *V*/*HE* gains, are

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

TABLE 3. *Means and SDs of the height error (HE) and cycle period (T) for the seven gravity conditions*

	Gravity							
	g_1	g_2	g_3	g_4	g_5	g_6	g_7	ANOVA
HE, mm	32	31	27	18	21	6	12	$F(6,49) = 1.3$
SD *HE,* mm	141	123	100	84	75	55	43	$F(6,49) = 16.4$*
T, ms	3,094	2,647	2,254	1,856	1,467	1,067	714	$F(6,49) = 1,122.4$*
SD *T,* ms	281	232	164	115	85	49	24	$F(6,49) = 45.1$*

*$P < 0.0001$.

displayed in Fig. 6*A* (solid line). A one-way ANOVA established a significant dependence of *V*/*HE* on *g* ($P < 0.0001$): the larger the value of *g*, the steeper the slope or higher the gain.

Note that the value of −0.84 for the nominal gravity g_7 is in the range reported by Wei et al. (2008): for $\alpha = 0.6$, values were between −0.75 and −0.85. Importantly, these values exceeded those resulting from purely self-stabilizing properties of the dynamic model. Thus the regression slopes arise from both the passive and active properties of the control system.

As another test of the degree of error compensation and, indirectly, of the feedback gain, the autocorrelation *AC*-1 of the sequence of height errors was calculated. Assuming feedback control, positive correlations across successive *HE* would be indicative of insufficient error compensation; high negative lag-1 autocorrelations would signal overcompensation (see the following text how this feature was quantified in model simulations). The *AC*-1 values obtained from the data were invariant across different gravities and close to zero, as shown in Fig. 6*B* (solid line). Single-sample *t*-tests for each of the seven data sets showed that *AC*-1 values were not significantly different from zero ($P < 0.01$), with exception of g_6 and g_7. In addition, the one-way ANOVA did not establish a dependence of *AC*-1 on *g* ($P > 0.56$). In sum, assuming feedback control on height errors, the lag-1 autocorrelations of *HE* suggested that height errors were appropriately compensated in the next cycle.

Racket kinematics

The next focus was on the continuous trajectory of the racket. As expected from previous work, the racket profiles changed markedly with increasing period, as shown in Fig. 7. Figure 7, *A* and *C* display the racket position (*top*) and velocity (*bottom*) averaged across subjects. For better comparison across the different gravity conditions the profiles were time-normalized. The normalized profiles were obtained by resampling the trajectories between two successive impacts picking 500 equally time-spaced samples. Subsequently, for each subject the three trials in one condition were averaged. The velocity traces have additionally been rescaled by the factor g_l/g_7 such that the velocity peaks are the same.

The data illustrate that the increase in cycle period did not lead to a simple scaling of the racket trajectory. With increasing period, the average racket trajectory changed shape from an approximately sinusoidal profile to one with a plateau of low or zero velocity (dwell interval) in the middle of the cycle. This transition in trajectory shape was quantified by two measures to illustrate this dependence on the gravity values: racket amplitude *RA* and activity period *AP*. These measures should have stayed invariant if the change in gravity had simply "stretched" the racket trajectory. Figure 8*A* shows a significant increase of

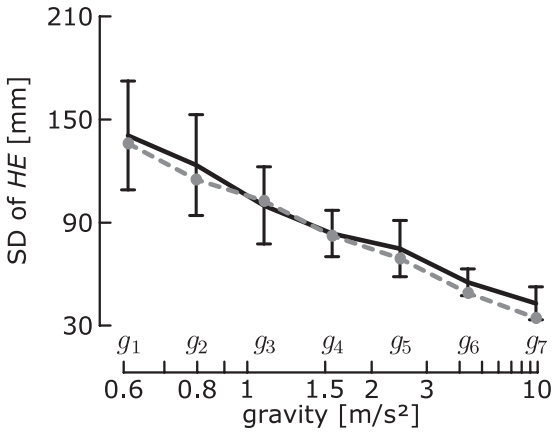

FIG. 5. SDs of the height error *HE* depending on the gravity condition: data (solid line; see also Table 3) and model-based simulations (dashed gray line). The error bars denote the 95% confidence intervals in the human data.

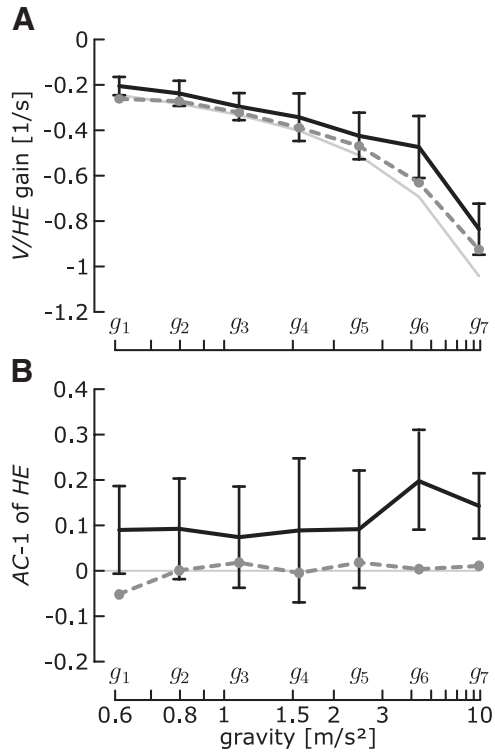

FIG. 6. Regression analyses. *A*: regression slopes between height error *HE* and the racket velocity at the next impact, *V*/*HE,* for the human data (solid line) and results simulated by the model (dashed and solid gray lines; see Model simulations in RESULTS). *B*: lag-1 autocorrelation coefficient *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 SDs of the human data.

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

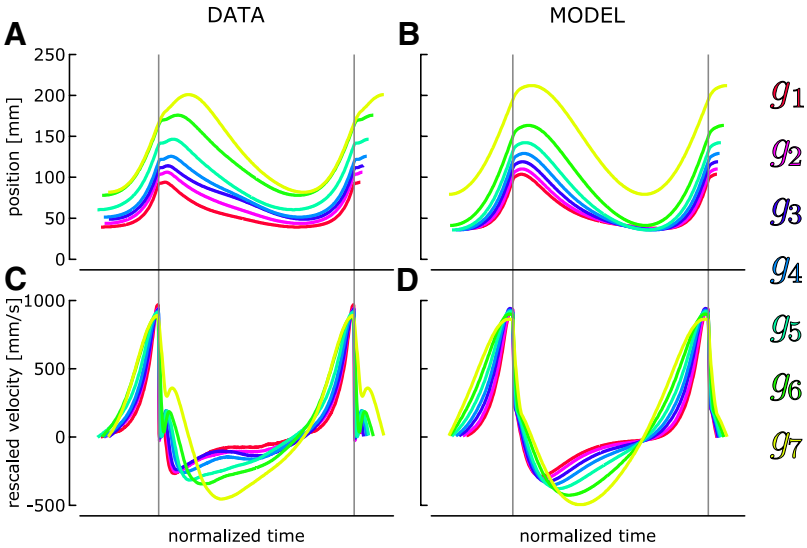

FIG. 7. Averaged trajectories of the racket position and rescaled velocity over the normalized time in the 7 gravity conditions. The racket kinematics have been resampled between 2 consecutive impacts (represented by the 2 vertical lines). *A*: averaged racket trajectories of the human subjects in the 7 gravity conditions. *B*: racket trajectory as generated by the model for all gravity conditions. *C*: rescaled velocity profile for the human subjects (averaged). *D*: rescaled velocity profile as generated by the model.

RA with gravity such that *RA* was almost twice as high for g_7 as for g_1. This was confirmed by a significant main effect of the one-way ANOVA [$F(6,49) = 8.6$, $P < 0.0001$]. Similarly, the activity period *AP* showed a significant increase with gravity [$F(6,49) = 8.8$, $P < 0.0001$], suggesting that the subjects interspersed dwell intervals as the cycle period became longer (Fig. 8*B*). Whereas at g_7 the racket moved with a velocity >5% of its peak during >90% of the cycle, at g_1 this activity period *AP* dropped to about 50%.

Model simulations

This section summarizes the simulation results of the model with parameters listed in Table 2. To afford direct comparisons, the simulation results are presented in the same figures together with the human data.

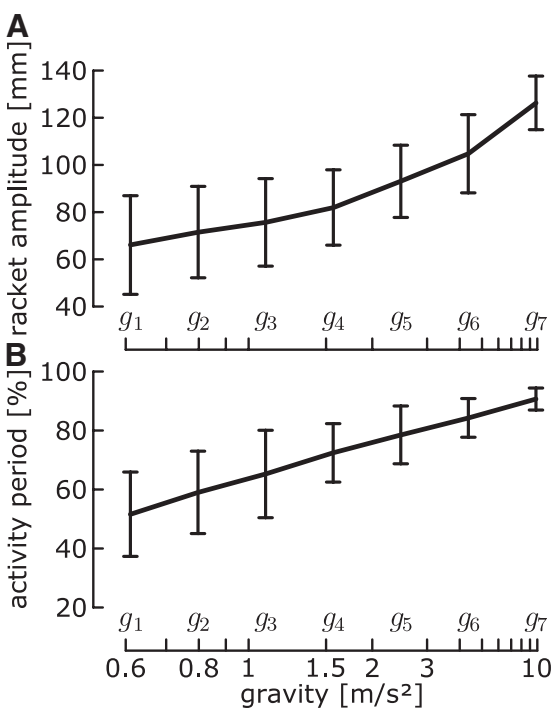

FIG. 8. Parameters of the racket kinematics depending on the gravity acting on the ball. *A*: racket amplitude *RA*. *B*: activity period *AP*. The panels show the participant means ± 95% confidence intervals.

The model was designed to perform the task by canceling the ball height errors *HE* within one impact, i.e., by calculating the necessary desired velocity (*Eq. 6*). Using *Eqs. 4* and *5,* it can be shown that this required a *V*/*HE* gain equal to

$$V/HE_{optimal} = \frac{-\alpha}{2(1+\alpha)}\sqrt{\frac{2g}{\Delta b}} \qquad (13)$$

where Δb is the distance between the target apex and the average impact position. This optimal gain is depicted in Fig. 6*A* (light gray). Its absolute value increased with gravity *g,* in agreement with what was observed in the actual data. The *V*/*HE* gains were also calculated from the simulated data (dashed line) and were close to the theoretical values (*Eq. 13*). As a consequence, *AC*-1 values of the simulated data were very close to zero for all gravity conditions; the correlation coefficient was never significantly different from zero (all $P > 0.1$). In sum, model simulations achieved results comparable to actual data for both the *V*/*HE* gain and *AC*-1. Only for the two largest gravity values, the actual *V*/*HE* gain was slightly smaller than predicted, resulting in slightly positive *AC*-1.

To assess what *AC*-1 values would be obtained in the absence of feedback control we also ran model simulation with the *V*/*HE* gain set to zero. To this end we forced the desired impact position and velocity of the racket to be constant, and computed the racket velocity from the steady-state value of *Eq. 6.* The resulting SDs of *HE* were about 30% higher and ranged between 200 mm (for g_1) and 50 mm (for g_7). As a result, the *AC*-1 values were also significantly higher, ranging between $r = 0.60$ and $r = 0.66$ across gravity conditions (all $P < 0.01$).

Model simulations also reproduced the pattern of change in the racket trajectories (second-layer of the controller). The average controller gains[1] for the seven gravity values computed by the optimization method are displayed in Fig. 9. Initially, all gains were low except $L_{1,t}$, which penalized large racket excursions away from the rest position, and was thus

[1] These were obtained by computing the control gains over a sequence of cycles based on the noiseless racket trajectories. The iteration was considered to have converged when the durations of two consecutive cycles were equal, and the sum of absolute differences between control gains was $<10^{-6}$.

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

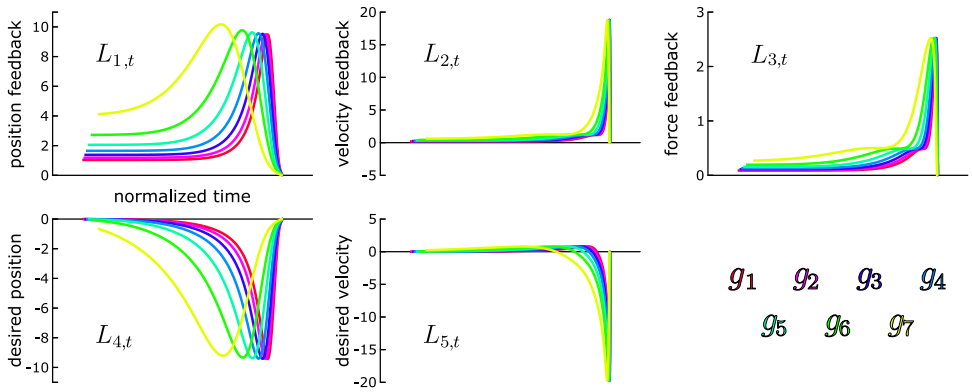

FIG. 9. Average gain vector $L_t = [L_{1,t}, L_{2,t}, L_{3,t}, L_{4,t}, L_{5,t}]$ computed by the iterative method described in *Mathematical model of the hybrid control architecture in* METHODS. The 5 panels refer to the gains multiplying the 5 state variables **x.** At the beginning of the cycle, the controller acts like a spring (high position gain). It then switches to a mode where it minimizes the difference between the actual and the desired racket states *1st* and *2nd columns*.

independent of the state to reach at impact. The control gains driving the racket to the desired impact state, $L_{4,t}$ and $L_{5,t}$, increased only during the second half of the cycle. Returning to Fig. 7, the right panels depict the racket position (Fig. 7*B*) and rescaled velocity (Fig. 7*D*) as generated by the model. The simulation results were resampled with the same procedure as the data. As can be seen, the racket trajectories changed across gravity conditions similar to the data, i.e., from a sinusoidal shape to trajectories with increasingly more pronounced dwell intervals. When the cycle period increased, the first and third terms of the cost function (*Eq. 7*) favored such slowing down in the racket trajectory preceding the impact. To obtain this matching the weights of the cost function were tuned to $w_{rest} = 0.012\sqrt{g_i}$ and $w_{energy} = 0.00002$. Note that w_{rest} was proportional to $\sqrt{g_i}$ or, equivalently, to $1/T$ (see *Eq. 5.*) Since the corresponding term of the cost function was integrated over the cycle with duration T, this term penalized the proportion of the movement deviating from the rest position r_0.

Sources of variability in the racket velocity

The SDs of *HE* showed a marked increase in variability with longer periods (Table 3 and Fig. 5). Are such changes in variability also observable across the cycle? Such on-line changes present another way of assessing the model, specifically the contribution of the signal-dependent noise (*Eqs. 9* and *10*). The SDs of the racket's normalized velocity during the cycle are displayed in Fig. 10*A*. Based on the same normalization and resampling procedures as for Fig. 7 the SDs were determined for each time bin. The variability profiles of all gravity conditions show pronounced peaks at the impacts. Although the magnitude of variability was comparable for the seven gravity conditions in most parts of the cycle, the peaks before impact increased for smaller gravities. A one-way ANOVA across subjects confirmed this impression: the main effect for gravity was significant [$F(6,49) = 6.2$, $P < 0.0001$].

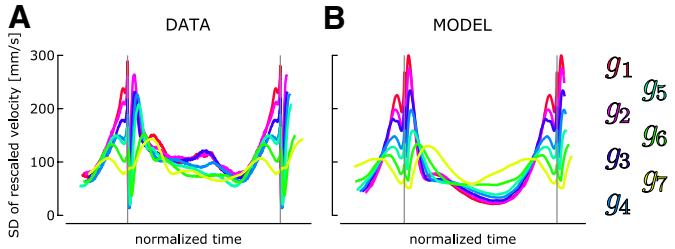

FIG. 10. SDs of the rescaled racket velocity over the normalized time. The 2 vertical lines represent 2 consecutive impacts. *A*: human data for the 7 gravity conditions. *B*: model simulations for the 7 gravity conditions.

To produce variability in the trajectory between impacts, the model had to include noise. Signal-dependent noise in the input $\dot{\sigma}(t)$ generated such variability in the racket trajectory, but would only slightly change the variability of ball height. Indeed, in the presence of signal-dependent noise the controller gains L_t (*Eq. 11*) decreased at the end of the cycle such that the control inputs became very small and, consequently, virtually noise-free around the impact (see Fig. 9). The racket states reached at impact were then close to the desired ones.

Thus another source of noise was needed to understand the data. Therefore, some uncertainty in computing the actual velocity to reach at impact $\dot{r}_{des,k+1}$ was included. This noise was quantified by the parameter matrix η_t in *Eq. 12*. By tuning $n_{vel} = 0.08$ m/s^2, the model generated variability of the ball apex of magnitude similar to that of the data. Note that the value of this parameter did not influence the computation of the optimal gains L_t.

Since the internal noise n_{vel} introduced uncertainty in the desired velocity that was not signal-dependent, the signal-to-noise ratio became smaller for smaller gravity conditions since they necessitated smaller desired racket velocities at impact. This phenomenon can be directly understood from *Eq. 2;* adding some uncertainty (noise) to $\dot{b}(t)$ would result in uncertainty of b_{apex}, with higher gains for smaller g. The correspondence of the variability in actual racket velocity and simulation results over the seven gravity conditions was obtained by tuning only the one parameter n_{vel}. Note that the parameter n_{pos} also reproduced the variability of the racket impact position (data not shown). To reproduce the noise across the cycle, the signal-dependent noise gain was set to $c = 1.3$ (see *Eq. 10*).

DISCUSSION

Bouncing a ball rhythmically to a target height constitutes an interesting task that combines continuous rhythmic movements of a racket with controlled discrete contacts between the racket and the ball. On the one hand, the task involves continuous rhythmic movements and it is conceivable that control and coordination are established by assembling the action system into an oscillator (Saltzman and Kelso 1987). On the other hand, the requirement to bounce and propel the ball to a specified height implies that the racket has to impact the ball with a contact velocity, depending on the impact position and the ball's state. For this latter discrete type of tasks various feedback control models including optimal feedback control

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

have been proposed. Note that the hybrid nature of the present task is not unusual and, for example, all locomotory activities have the same hybrid dynamics: the leg movements are continuous, whereas the control of propulsion of the body is confined to intermittent impacts with the ground.

How are such hybrid movements controlled? Previous work of our group has examined the ball bouncing task with a focus on its rhythmicity. Using a mathematical model of the task we showed how dynamic stability of the periodic dynamic system played a significant role in achieving the task (Dijkstra et al. 2004). Using a discretized version of the continuous mechanical system, we showed that fixed-point attractor stability existed. Such dynamic stability implies that small errors self-stabilize and need not be compensated for explicitly. However, more recent research also revealed that active error corrections are intertwined with the self-stabilizing properties of the task (Wei et al. 2007, 2008). Especially in the presence of larger perturbations, the ball–racket impacts become controlled events and the racket impact velocity of the next impact is actively adjusted based on the perceived height errors of the ball. These conclusions were partly supported by changes in the racket trajectory that deviated from their usual sinusoidal profile to prepare and modify the upcoming impact. Such deviations from sinusoidal movements were no longer consistent with the assumptions of the model. However, no alternative model has been suggested that addressed these active corrections.

The present study aimed to examine these explicitly controlled features of task execution. To provide an appropriate test set of data, the experiment presented conditions that scaled the tempo of performance such that the rhythmic stable performance, predominant at preferred movement rates, would give way to a more discrete type of performance at slower rates. A number of previous studies reported that increasing the cycle duration of a rhythmic movement ultimately induced dwell intervals into each cycle because people tend to avoid very slow movements; thereby, the continuity of the movement gives way to a string of discrete segments (Adam and Paas 1996; Nagasaki 1991; Ronsse et al. 2008b; van der Wel et al. 2010). The present experiment constrained the cycle duration of the task by changing the gravity. The alternative straightforward manipulation would have been to increase the target height. However, this was not practical because similar tempo changes would have required target heights that exceeded the ceiling height of the lab. Although changing gravity led to the desired slowing of the periods, this might have introduced a potential concern: previous studies demonstrated that humans anticipate the dynamics of falling bodies guided by the nominal gravitational field and would thus misestimate the ball trajectory in unusual gravity conditions (e.g., Zago and Lacquaniti 2005; Zago et al. 2004). However, these studies invariably compared performance between nominal and zero gravity. In contrast, in our experiment the ball was always in a gravitational field.

To begin, at normal gravity our results replicated performance characteristics reported in earlier studies. The target height was achieved with racket movements that were relatively smooth and sinusoidal. The variability of ball apex errors was within the same range. However, when gravity was decreased and the movement tempo became slower, the smooth trajectories between the individual contact events started to separate. This was concluded from the racket trajectories that had interspersed intervals of rest. Each contact became like a single target event where the racket trajectory accelerated toward the ball for contact. Interestingly, variability in this latter type of performance was significantly higher, indicating that noise became more prominent and control was more difficult for slower velocities.

To shed light on what processes gave rise to these data we developed a model that was based on optimal feedback control. We implemented a hierarchical controller that included two layers responsible for the continuous and discrete aspects of the task (Todorov et al. 2005). The first layer was discrete in time: after contact, it computed the state that the racket had to reach at the next impact. The intermittent targets for movement execution were specified based on state information about the ball trajectory. The second layer generated the continuous racket trajectories from one impact to the next, as specified by the first layer. The racket trajectory was generated by optimal feedback control: the controller determined a control policy that minimized a given cost function, penalizing both the discrepancy between the target and the actual state of the effector (maximizing the task's reward), the effort associated with the movement production (the control cost), and the racket's movements away from its rest position. This framework extends previous studies developed on the basis of stochastic optimal feedback control that used only a single-layer model (e.g., Davis and Vinter 1985). Unlike our present approach, most studies that combined experimental data with modeling focused on unimanual and bimanual reaching (Diedrichsen 2007; Diedrichsen and Gush 2009; Liu and Todorov 2007; Todorov 2004, 2005; Todorov and Jordan 2002). Only few studies investigated rhythmic movements such as locomotion or swimming from a similar optimality perspective, although they reported only simulation results (e.g., Srinivasan and Ruina 2006; Tassa et al. 2008).

One focus of our study was to identify how subjects adjusted their racket actions to accurately bounce the ball to the target height as instructed. Data showed that subjects consistently tried to cancel errors in ball height from one cycle to the next, by adapting the racket velocity at the next impact. This led to negative correlations between height error and racket velocity at contact in all gravity conditions. A consequence was that successive height errors were not, or only weakly, correlated during steady-state bouncing (lag-1 autocorrelation close to zero). The optimal model similarly revealed significant negative correlations between the perceived height error and the following impact velocity of the racket and near-zero lag-1 autocorrelation of height errors in all gravity conditions. An additional model simulation that eliminated the optimization of racket velocity at impact and forced the impact racket velocity to be constant across cycles showed consistently less autocorrelation and higher variability in the height error. Although other modeling scenarios might also lead to less variability in the height error than the one obtained with *Eq. 6,* this would have been accompanied by nonzero lag-1 autocorrelation of height errors, inconsistent with the data. Interestingly, the results on the velocity profiles were consistent with this tuning of impact velocity: the data showed that the variability of racket velocity was low at the end of the cycle, suggesting that accurate control of the racket velocity in that portion of the cycle had some priority. This last result also suggests that the

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

fine control of the impact velocity maximized the responsiveness to perturbations, i.e. within one cycle and much faster than predicted by the stochastic model of dynamic stability (Wei et al. 2008). Additionally, Wei et al. (2007) studied how experimentally applied perturbations of the ball trajectory (induced by modifying the post-impact velocity of the ball) were compensated by the actor. They showed that humans were also able to compensate within one cycle directly following the perturbation, which is much faster than predicted by the deterministic model of dynamic stability. The present study illustrated that these active compensations could result from an adaptation of the racket velocity at impact due to feedback from the ball height during the preceding cycle.

The model simulations also replicated the observed variability in the continuous racket trajectory and at impact by introducing two separate sources of noise (Faisal et al. 2008). The continuous control layer included signal-dependent noise in the trade-off between maximizing task reward and minimizing control cost (Harris and Wolpert 1998; Todorov 2005). In addition, the discrete layer included a constant noise parameter that compounded the impact with noise. The effect of the latter noise term was that the variability of the height error became larger with slower execution periods. Stated differently, the longer the interbounce interval, the more uncertainty in the next impact, as frequently documented in experiments on production and perception of rhythmic temporal intervals (e.g., Ivry and Hazeltine 1995). Our model therefore implies that slow velocities of the limbs are relatively more difficult to estimate. This might be related to the fact that people avoid moving slowly and prefer to introduce dwell times (Adam and Paas 1996; Nagasaki 1991; Ronsse et al. 2008b; van der Wel et al. 2010).

Given the good match between simulation with experimental results in previous work and our current modeling, the question is how the model relates to the functioning of the nervous system. Recently, Shadmehr and Krakauer (2008) reviewed neurological studies to identify correspondence between computational stages to dedicated neural substrate, including dysfunctions associated with specific brain lesions (Scott 2004, 2008). Justification for optimality principles is seen in the fact that reward-based learning with cost estimation is a key function of the basal ganglia. In this spirit, the proposed stratification of control may be associated with different foci of brain activation. For example, in a task involving visually cued finger movements the supplementary motor area and the lateral premotor cortex were mostly activated during presentation of the visual cue, whereas the primary motor cortex was activated during execution of the movement. During the phase of movement execution, the supplementary motor areas experienced a gradual shift in activation from the anterior to the posterior portion (Lee et al. 1999).

The two-layered structure of the proposed model was motivated by the hybrid nature of the task, intertwining continuous dynamics (control of the racket) with discrete events (impacts with the ball). This particular control structure also provided mathematical convenience, for example it afforded to use the LQG framework. However, we believe that any alternative model would have to distinguish between two possibly intertwined components to generate the same results: *1*) planning the next impact based on the sensory feedback of the current cycle, represented by the discrete layer in our model; and *2*) executing the continuous racket movement, similar to the extensively modeled reaching movements where continuous control appears most appropriate. Note, though, that in the brain even the discrete layer is likely to be a continuous process, similar to a discrete judgment that is achieved by a neural process that accumulates evidence to reach a final decision threshold (for a review see e.g., Gold and Shadlen 2007). These two elements would have to be combined in a stratified structure: the desired impact velocity could be (re-) computed based on perceived perturbations, leading to possible recomputation of the feedback gains driving the racket movement. The two-layer distinction is also consistent with the proposition that discrete and continuous elements of a movement may be under separate control, as suggested by several behavioral and imaging studies (Buchanan et al. 2006; Huys et al. 2008; Schaal et al. 2004; Smits-Engelsman et al. 2006; Sternad et al. 2000; Sternad 2008).

In conclusion, the present paper suggests that the control of hybrid movements, consisting of continuous rhythmic movements with intermittent discrete interactions between effector and object, has two nested layers of control. Given the significantly more extensive involvement of cortical areas in discrete movements, a next step to test the viability of this interpretation may be to examine whether the balance between both control strategies depends on the level of expertise and on the cognitive resources available during task execution, affected by a dual-task or stress (Ehrlenspiel et al. 2010).

ACKNOWLEDGMENTS

We thank the reviewers for their help in improving the manuscript.

GRANTS

This work was supported by a Belgian Network in Dynamical Systems, Control, and Optimization grant to R. Ronsse, funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State Science Policy Office; the Francqui Foundation; and the Fonds de la Recherche Scientifique (Belgium) for a scholarship grant for R. Ronsse to visit D. Sternad at the Pennsylvania State University. D. Sternad was funded by a National Science Foundation Grant PAC-0450218, National Institute of Child Health and Human Development Grant R01-HD-045639, and Office of Naval Research Grant N00014-05-1-0844.

REFERENCES

Adam JJ, Paas FGWC. Dwell time in reciprocal aiming tasks. *Hum Mov Sci* 15: 1–24, 1996.

Buchanan JJ, Park JH, Shea CH. Target width scaling in a repetitive aiming task: switching between cyclical and discrete units of action. *Exp Brain Res* 175: 710–725, 2006.

Davis MHA, Vinter RB. *Stochastic Modelling and Control, Volume XII* (Monographs on Statistics and Applied Probability). London: Chapman & Hall, 1985.

de Rugy A, Wei K, Muller H, Sternad D. Actively tracking "passive" stability in a ball bouncing task. *Brain Res* 982: 64–78, 2003.

Diedrichsen J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* 17: 1675–1679, 2007.

Diedrichsen J, Gush S. Reversal of bimanual feedback responses with changes in task goal. *J Neurophysiol* 101: 283–288, 2009.

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* 11: 319–344, 2004.

Ehrlenspiel F, Wei K, Sternad D. Open-loop, closed-loop and compensatory control: performance improvement under pressure in a rhythmic task. *Exp Brain Res* 201: 729–741, 2010.

Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.

Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007.

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026

Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.

Hogan N, Sternad D. On rhythmic and discrete movements: reflections, definitions and implications for motor control. *Exp Brain Res* 181: 13–30, 2007.

Huys R, Studenka BE, Rheaume NL, Zelaznik HN, Jirsa VK. Distinct timing mechanisms produce discrete and continuous movements. *PLoS Comput Biol,* 4: e1000061, 2008.

Ijspeert AJ. Central pattern generators for locomotion control in animals and robots: a review. *Neural Netw* 21: 642–653, 2008.

Ivry RB, Hazeltine RE. Perception and production of temporal intervals across a range of durations: evidence for a common timing mechanism. *J Exp Psychol Hum Percept Perform* 21: 3–18, 1995.

Katsumata H, Zatsiorsky V, Sternad D. Control of ball–racket interactions in rhythmic propulsion of elastic and non-elastic balls. *Exp Brain Res* 149: 17–29, 2003.

Lee KM, Chang KH, Roh JK. Subregions within the supplementary motor area activated at different stages of movement preparation and execution. *NeuroImage* 9: 117–123, 1999.

Liu D, Todorov E. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27: 9354–9368, 2007.

McGeer T. Passive dynamic walking. *Int J Robot Res* 9: 62–82, 1990.

Nagasaki H. Asymmetrical trajectory formation in cyclic forearm movements in man. *Exp Brain Res* 87: 653–661, 1991.

Ronsse R, Lefèvre P, Sepulchre R. Rhythmic feedback control of a blind planar juggler. *IEEE Trans Robotics* 23: 790–802, 2007.

Ronsse R, Lefèvre P, Sepulchre R. Robotics and neuroscience: à rhythmic interaction. *Neural Netw* 21: 577–583, 2008a.

Ronsse R, Sternad D, Lefèvre P. A computational model for rhythmic and discrete movements in uni- and bimanual coordination. *Neural Comput* 21: 1335–1370, 2009.

Ronsse R, Thonnard J-L, Lefèvre P, Sepulchre R. Control of bimanual rhythmic movements: trading efficiency for robustness depending on the context. *Exp Brain Res* 187: 193–205, 2008b.

Saltzman E , Kelso JA. Skilled actions: a task-dynamic approach. *Psychol Rev* 94: 84–106, 1987.

Schaal S, Atkeson CG, Sternad D. One-handed juggling: à dynamical approach to a rhythmic movement task. *J Mot Behav* 28: 165–183, 1996.

Schaal S, Sternad D, Osu R, Kawato M. Rhythmic arm movement is not discrete. *Nat Neurosci* 7: 1136–1143, 2004.

Scott SH. Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* 5: 532–546, 2004.

Scott SH. Inconvenient truths about neural processing in primary motor cortex. *J Physiol* 586: 1217–1224, 2008.

Shadmehr R, Krakauer JW. A computational neuroanatomy for motor control. *Exp Brain Res* 185: 359–381, 2008.

Shadmehr R, Wise SP. *The Computational Neurobiology of Reaching and Pointing. A Foundation for Motor Learning.* Cambridge, MA; MIT Press, 2005.

Smits-Engelsman BC, Swinnen SP, Duysens J. The advantage of cyclic over discrete movements remains evident following changes in load and amplitude. *Neurosci Lett* 396: 28–32, 2006.

Srinivasan M , Ruina A. Computer optimization of a minimal biped model discovers walking and running. *Nature* 439: 72–75, 2006.

Sternad D. Towards a unified framework for rhythmic and discrete movements: behavioral, modeling and imaging results. In: *Coordination: Neural, Behavioral and Social Dynamics,* edited by Fuchs A, Jirsa V. New York: Springer, 2008, p. 105–136.

Sternad D, Dean WJ, Schaal S. Interaction of rhythmic and discrete pattern generators in single-joint movements. *Hum Mov Sci* 19: 627–664, 2000.

Sternad D, Duarte M, Katsumata H, Schaal S. Bouncing a ball: tuning into dynamic stability. *J Exp Psychol Hum Percept Perform* 27: 1163–1184, 2001a.

Sternad D, Duarte M, Katsumata H, Schaal S. Dynamics of a bouncing ball in human performance. *Phys Rev E* 6301: 011902, 2001b.

Tassa Y, Erez T, Smart W. Receding horizon differential dynamic programming. In: *Advances in Neural Information Processing Systems,* edited by Platt J, Roller D, Singer Y, Roweis S. Cambridge, MA: MIT Press, 2008, vol. 20, p. 1465–1472.

Todorov E. Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.

Todorov E. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17: 1084–1108, 2005.

Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.

Todorov E, Li W, Pan X. From task parameters to motor synergies: a hierarchical framework for approximately-optimal control of redundant manipulators. *J Robot Syst* 22: 691–710, 2005.

van der Wel RPRD, Sternad D, Rosenbaum DA. Moving the arm at different rates: slow movements are avoided. *J Mot Behav* 42: 29–36, 2010.

van Mourik AM, Beek PJ. Discrete and cyclical movements: unified dynamics or separate control? *Acta Psychol* 117: 121–138, 2004.

Wei K, Dijkstra TMH, Sternad D. Passive stability and active control in a rhythmic task. *J Neurophysiol* 98: 2633–2646, 2007.

Wei K, Dijkstra TMH, Sternad D. Stability and variability: indicators for passive stability and active control in a rhythmic task. *J Neurophysiol* 99: 3027–3041, 2008.

Zago M, Bosco G, Maffei V, Iosa M, Ivanenko YP, Lacquaniti F. Internal models of target motion: expected dynamics overrides measured kinematics in timing manual interceptions. *J Neurophysiol* 91: 1620–1634, 2004.

Zago M, Lacquaniti F. Internal model of gravity for hand interception: parametric adaptation to zero-gravity visual targets on earth. *J Neurophysiol* 94: 1346–1357, 2005.

Downloaded from journals.physiology.org/journal/jn (188.252.120.050) on September 22, 2026