The Pennsylvania State University

The Graduate School

College of Health and Human Development

STABILITY AND VARIABLITY IN A RHYTHMIC TASK: BEHAVIORAL DATA AND DYNAMIC MODELS

A Thesis in

Kinesiology

by

Kunlin Wei

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

May 2007

The thesis of Kunlin Wei was reviewed and approved* by the following:

_____

Dagmar Sternad
Professor of Kinesiology
Thesis Adviser
Chair of Committee

_____

Vladimir Zatsiorsky
Professor of Kinesiology and Biobehavioral Health

_____

Stephen Piazza
Associate Professor of Kinesiology

_____

Joseph Cusumano
Professor of Engineering Science and Mechanics

_____

Hermann Müller
Associate Professor of University of Saarbrücken
Special Member

_____

Phillp E. Martin
Professor of Kinesiology
Department Head

*Signatures are on file in the Graduate School.

ABSTRACT

Using the perceptual-motor skill of rhythmically bouncing a ball as an experimental vehicle the present dissertation examined questions about control strategies and their acquisition, adaptation and transfer. Previous studies had already documented that the actor is sensitive to the stability properties of the task dynamics and performs the rhythmic actions with a strategy where effects of perturbations converge back to steady state without requiring error-correcting racket movements. This behavior is consistent with the predictions from stability analyses of a dynamic model of the task. Experiment 1 continued to scrutinize this prediction by applying a range of perturbation magnitudes, designed to be within and outside of the model's basin of attraction. Results showed that even small perturbations that were predicted to equilibrate passively were blended with active control flexibly responding to perceived errors. However, the time course of return to steady state performance was qualitatively consistent with predictions from passive stability. Experiment 2 investigated how the actor combined passive stability with active control when the dynamic stability of the task was manipulated by varying the coefficient of restitution at the racket-ball contacts. To quantify the degree of control the covariance structure of the state variables was compared with model predictions. These predictions were obtained from a model that was extended by stochastic components to yield predictions about the structure of fluctuations at steady state. Results revealed that, paradoxically, variability of performance decreased with decreasing stability, contrary to common expectations in motor control. This was explained by increasing compensatory variability in execution, a signature of control. Hence, actors rely on passive stability when the stability of the system is high and employ more active control when stability is reduced. Applying the same variability and stability analysis Experiments 3 and 4 revisited issues of acquisition, adaptation and transfer in the same skill. Both experiments clearly demonstrate that the performance improvement is correlated with increased sensitivity to passive stability. Variability was evaluated in a space spanned by execution and result variables by applying the *TNC*-decomposition of variability (*T*olerance, *C*ovariation, *N*oise). Results highlighted how learning and adaptation is a migration through the execution space combined with the fine-tuning of covariation between relevant variables and a reduction of noise. Sensitivity to passive stability forms some abstract knowledge that is easily transferable across effectors. Compared to the mere examination of outcome measures, the decomposition method provided finer-grained insights about learning, adaptation, and transfer. Taken together, the four studies extend our understanding about how coordinated performance is achieved by exploiting task stability and active control accompanied by changes in the structure of variability.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 3:

Chapter 4:

Chapter 5:

Chapter 6:

LIST OF TABLES

Chapter 3:

Chapter 4

ACKNOWLEDGEMENTS

I would like to express my deepest appreciation to the following people who helped me through these years in graduate school leading to the completion of my Ph.D degree.

My advisor Dagmar Sternad has helped me in all aspects, from the training of scientific thinking to the details in writing and oral presentation, from the design of experiments to the planning of my academic career, not to mention the care she gave me during my intensive days after my surgery and before my dissertation defense. She has taken the role not only as a mentor to me but also, in some sense, as a mother waiting patiently for the child to grow and to mature. Through hardships and happiness, she has helped me throughout these 6 years to become who I am now. I will always remember everything she did for me.

I would also like to thank my collaborator Tjeerd Dijkstra for the always constructive scientific discussions he provided, his unselfish sharing of his knowledge and skills, and his eternally optimistic attitude to everything. Every single time working with him has been a happy and fruitful experience.

I would also like to express my sincere appreciation to my committee members, Vladimir Zatsiorsky, Hermann Müller, Joseph Cusumano, and Stephen Piazza, for giving me valuable comments and suggestions during the comprehensive exam, the dissertation proposal and finally the defense. They helped me to understand what it takes to be a qualified Ph.D student and encouraged and inspired me to conduct research of high quality.

I would also like to thank all the lab members and alumni, especially Hiromu Katsumata, William Dean, Hong Yu, Tjitske Boonstra, Mohamed Tlili and Xiaogang Hu, for helping me in conducting experiments and data analysis, and for the valuable discussions I had with them and for their friendship.

I also want to give my sincere thanks to the faculty members Mark Latash and John Challis and the departmental secretaries Amy Mitchell and Dori Sunday for their continuous considerate support.

Lastly, I want to thank my family, my dad Wei, Dingqiu, my mom Wang, Puping and my sister Wei, Kewei, for their endless encouragement and support. Their love has provided essential momentum for me to pursue my career and my life.

CHAPTER 1

Introduction and Literature Review

1.1. Redundancy and Task Constraints in Movement Coordination

In everyday life humans are continuously engaged in achieving specific task goals in ever changing environments. How biological systems coordinate and regulate their many degrees of freedom in such varying contexts is arguably the central issue in motor control. One of the well-known example tasks to explain the degree of freedom problem is to reach to a target with the tip of one's finger. The task involves the rotation of several joints, each of which has more than one degree of freedom in orientation and amplitude. Therefore, different joint angle combinations can achieve the same endpoint position. This redundancy similarly exists at the muscular level, where multiple muscles span one or several joints with different lines of action rotate each joint such that different combinations of muscular contractions can lead to the same joint rotation. The same phenomenon repeats itself at more microscopic levels. The fact that the entities in the human body that need to be controlled outnumber the variables describing the motor task is commonly referred as Bernstein's (1967) degrees of freedom problem or redundancy problem.

What is less often emphasized is that not only the human action system poses this challenge but also the task offers redundant degrees of freedom. Many tasks have different solutions that lead to an equivalent level of success. For instance, in a dart throwing task, a large release angle combined with a small release velocity can hit the exact same spot at a bull's eye as a small release angle with a high release velocity. These two variables can be combined in a compensatory manner to offer an infinite– redundant – number of different solutions. Hence, the redundancy problem is not only caused by the complexity of the human action system, but also by the specific task that the actor is engaged in.

There is general agreement in motor control research that it is highly unlikely that

there is an omnipotent controller in the brain that specifies every degree of freedom all the time. It appears to be more reasonable to assume that simpler control strategies have evolved for coordinating the behavioral patterns in human and other biological systems. The question, however, is what are the principles from which these simpler strategies arise? For example the dynamical systems perspective has cast the problem in the language of dynamical systems and complexity theory: behavioral patterns are simple structures that emerge from the complexity of the high-dimensional dynamical system. Such coordinative structures are defined as a group of muscles across a certain number of joints that act as a functional unit, also referred to as synergies in later texts (Turvey, Fitch, & Tuller, 1978). Such functional units are hypothesized to operate as a dynamical system that is subject to constraints which, following Newell's proposition, arise from three sources: the organism, the environment, and the task (Newell, 1986). The question, of course, is *how* these dynamical systems emerge from the constraints.

1.1.1. Organismic, Environmental and Task Constraints

The importance of *organismic* constraints has for example been highlighted in studies where subjects oscillate two pendulums, one held in each hand, with different masses and lengths, emulating limbs with different mass length properties. When subjects are asked to swing them in phase at the same frequency, the pendulum with larger moments of inertia will lag behind the pendulum with less moment of inertia. This phase relationship is systematically dependent on the difference in eigenfrequencies between the two pendular limbs and the observation can be modeled by coupled oscillator system. If the two pendulums are regarded as modeling the actor's limbs, this example shows how physical properties of the actor determine the coordination pattern he/she displays. For the more ecological situation of locomotion it was shown that the eigenfrequencies of the legs successfully predicted the preferred walking speed of humans and quadrupeds (Holt, Hamill, & Andres 1990, 1991; Kugler & Turvey, 1987). Organismic constraints not only include the biomechanical factors, but also include the neural and psychological factors such as handedness or

personal preferences developed from past experience.

The actor interacts with the *environment* both physically via force fields and informationally via perceptual coupling. Examples for the former are that human movements are always embedded in the gravitational field but also all manipulations and collisions with objects produce forces. The informational influence from the environment can also be demonstrated with numerous examples. In static postural control it has been shown that if the task is to stand still, visual and haptic information can significantly influence the migration of the center of pressure. Haptic information from lightly touching an external surface significantly reduces or biases the direction of postural sway and oscillations in the visual surrounds impose synchronized fluctuation patterns with systematic phase relationships (Dijkstra et al., 1994a, b; Jeka et al., 1997; van Asten, Gielen & Denier van der Gon, 1988). A recent study on a reaching task demonstrated that the uncertainty of sensory feedback is combined with prior knowledge about the task to plan the action such the information from the environment largely determines the strategy that the actor selects (Körding & Wolpert, 2004).

The third kind of constraints, those exerted by the *task*, have been investigated less frequently. When the actor is involved in a motor task in interaction with the environment, the specific task solution is restricted by the physics of the task. For instance, Dingwell et al. (2004) studied a reaching task where the actor was required to transport a mass, connected to the hand with a spring, quickly to a target. A physical model of this hand-mass-spring system was developed to derive predictions for the hand movement such that the mass moves in an optimally smooth way. By comparing the predictions from this model with actual human trajectories, the authors found that indeed the trajectory of the transported object was indeed smooth which implied that the hand's trajectory was not smooth and sometimes even biphasic. These results showed task induced unusual trajectories of the hand that no longer displayed minimum jerk (Flash & Hogan, 1985). A similar approach the control problem was taken in a series of studies on cascade juggling (Beek, 1988; Beek & Turvey, 1992;

Beek, 1995). Rather than starting with an analysis of the hand movements or perceptual information about the ball trajectories, Beek and colleagues followed an insight by Claude Shannon who identified that the hand loop times and ball dwell times stood in a very simple relation to the number of balls and the number of hands. This relation presented essential constraints on the duration that the balls had to be held that every juggler had to satisfy. Further analysis of these timing relations showed how jugglers chose the time windows and how they maneuvered their timing to remain maximally flexible and adaptive. These studies exemplify how specific details of the task set windows on how the actor can assemble his/her movement systems to obtain successful performance. The question is how does the actor discover the task dynamics and assemble and tune his/her coordinative structure accordingly?

1.1.2. Bouncing a Ball: A Simple Task with Interesting Dynamics

The present thesis picks up this challenge and examines movement coordination from this perspective. The first step to answer this question is to analyze the task and its constraints. To this end a specific model task must be chosen that is simple enough to permit such analysis but is also representative for a class of actions to allow generalization. The present research has chosen the task of bouncing a ball as its model system in continuation of previous research by our group. The task requires the actor to move a racket to rhythmically propel the ball in the air. The task goal is to keep the ball amplitude consistent, which is consistent with performing a stable rhythmic bouncing pattern. The rationale for choosing this model system is that bouncing a ball presents a tractable dynamical system that, despite its simplicity, shows interesting dynamics. The bouncing ball system has been used in text books on nonlinear dynamics as a example for a nonlinear map that display stability with a period-doubling route to chaos (Guckenheimer & Holmes, 1983; Tufillaro, Abbott, & Reilly, 1992). If an actor performs this rhythmic task, he has to face and solve all the essential “problems” of perception and action. Further, the task is representative for all those movements that involve a contact with an external object where its control is

focused onto one moment. A comparable focus is found in walking and running where control of the movement pattern is confined to the phase of foot contact with the ground. Of course many other hitting tasks are found in sport activities.

1.1.3. Stability as a Soft Task Constraint

Having chosen a dynamical model for the task the next step in our task-based approach is to analyze the model. Specifically, the goal of the analysis is to examine the stability as it originates from the dynamics of the task and how it serves as a constraint for performance. The principal concepts are stability and attractor. An attractor is a stable state of a dynamical system such that any trajectory close to it remains close even if slightly perturbed. In mathematics and nonlinear dynamics stability has several different definitions, including bounded-in bounded-out stability, neutral stability, to Lyapunov stability, and asymptotic stability. For the present thesis we adopted the definition by Hasan (2005) adapted to issues concerning movement: "small effects of a perturbation on a trajectory are posited to remain forever small, whether they are in states of maintained position or movement". This definition is similar to Lyapunov stability as defined in mathematics texts which states that the dynamic system is stable in the Lyapunov sense if infinitesimally small variations in the state of the system remain infinitesimally small forever (Hasan, 2005; Jordan & Smith, 1987).

In a given task stability can arise from the interaction between the actor and the environment. The actor perceives behaviorally relevant *information* from the environment and performs adaptive movements according to this information and the task goal. The action, in turn, modifies the environment *physically* and thereby creates new information for the ensuing action (Warren, 2006). How perceptual information is used in stabilizing coordination patterns has been studied in the context of bimanual rhythmic coordination. For rhythmic movement coordination between two index fingers Kelso and colleagues found that the inphase and antiphase relationship between two oscillations were the two most stable modes and the transition from antiphase mode to inphase mode would happen as a result of stability change upon

increasing oscillation frequency (Kelso, 1984). These phenomena were explained using the concept of stability in the HKB model (Haken, Kelso, & Bunz, 1985). Several follow-up studies showed that the coupling between two oscillations was largely information-based, since the same phenomena happened for different effectors (arm and leg coordination for example, Kelso & Jeka, 1992), and even for inter-person coordination (Richardson, Marsh, & Schmidt, 2005; Schmidt, Carello, & Turvey, 1990; Schmidt et al., 1998). The latter finding highlights that perceptual coupling alone can give rise to the same attractor landscape.

In contrast, the stability arising from the *physical* interaction between the actor and the environment has been rarely addressed in the motor control literature. The physics of some tasks can offer stable fixed point or limit cycle solutions. If this stability originates only from the physical dynamics of the task without including any biological control mechanisms it will be termed *passive stability*.

1.1.4. Dynamic Walking in Robotics: Passive Stability in Locomotion

One area of research where passive stability has been studied at great length is locomotion – or rather passive dynamic walking. Dynamical stability analyses of legged locomotion have revealed that a completely passive two-legged robot without any actuators and control can produce stable locomotion (McGeer, 1990). Such a biped walker with well-tuned physical parameters can walk down a gentle slope with a surprisingly human-like gait. The small fluctuations in the stepping motion are dissipated due to the stability properties of this system. This original passive dynamic walker has been further developed by including relatively simple actuators such that it can now walk on level ground (for a review see Collins et al., 2005). The important implications from these studies for biological movement control is that the dynamics as it arises from the interaction between the actor and its environment can offer inherent stability in the context of a specific task. The exhibited coordination pattern can be significantly influenced by this stability if the actor attunes to the stable regime and exploits it. One core attraction of this understanding of movement is that stability can be obtained without any control. Small fluctuations converge back to the attractor

without requiring explicit error compensation. Hence, performing with passive stability is very efficient in terms of control effort.

However, this exploitation of passive stability can be complemented by additional active *control*. For example, passive dynamic walking models only offer a small basin of attraction and the stability resisting external disturbances is limited (Schwab & Wisse, 2001; Garcia et al., 1998). Active control can enlarge the spectrum of behaviors to come closer to human walking which seems fairly resistant to perturbations. With a view to locomotion in biological systems, which integrates rapid short-loop reflexes and long-loop feedback mechanisms, Kuo (2002) combined a feedforward central pattern generator with a feedback loop in a model of rhythmic locomotory movements. The proposed hybrid feedforward and feedback model not only yielded fictive locomotion, but also compensated for both external disturbances and noise from sensors. In sum, this research in robotics showed that by a subtle combination of active control with passive stability both stability and adaptability is enhanced and ultimately richer behavior is obtained.

1.1.5. Passive Stability and Active Control in Ball Bouncing

Our research on the ball bouncing task has taken a similar strategy as the research in robotics to understanding movement coordination. In previous work a minimal model was developed that describes the physics of the task (Dijkstra, Katsumata, de Rugy, & Sternad, 2004; Schaal, Sternad, & Atkeson, 1996; Sternad, Duarte, Katsumata, & Schaal, 2001). Three elementary physical facts are the foundation for this model: gravity governs the ballistic flight of the ball, the racket moves periodically, the impact between the ball and the racket is instantaneous with a coefficient of restitution capturing the energy loss. In a second step, predictions about stability were derived using linear stability analysis and a non-local stability analysis. Despite its relative simplicity the model has dynamically stable solutions, with period-1 solutions and a period-doubling route to chaos. The stable period-1 solutions are characterized by ball-racket contacts that occur during the upwardly decelerating phase of the racket trajectory. If the actor performs the task with a strategy that is

consistent with this criterion for stability, then small deviations in the movement variables introduced either from motor noise or external perturbations can disappear without explicit corrective adjustments in the racket movements. This stable solution has the advantage in control effort as only open-loop control is needed.

Previous empirical results showed that human subjects indeed exploit passive stability. A series of experiments was run using three different experimental set-ups that consisted of a real racket and a ball suspended to a long boom such that the ball was confined to an approximately vertical trajectory consistent with the vertical motions of the model. In all experiments human actors opted for a strategy where the racket hit the ball at its decelerating phase. Further, some individual subjects' data on learning showed that performance variability decreased across practice and this time course was accompanied by an increasing attunement to the stable solutions (Dijkstra et al., 2004; Sternad et al., 2001). It is important to emphasize that these findings are not trivial as rhythmic bouncing can also be achieved by other control strategies. For example, a robotic implementation of the same task by Bühler and Koditschek (1994) showed how the so-called mirror algorithm also achieves stable behavior. This algorithm scales the racket movements to the ball movements such that the racket velocity continuously "mirrors" the ball velocity to achieve periodic bouncing. This feedback-based strategy has been implemented successfully in a robot juggler operating both in 2D and 3D. Importantly for the present argument is that the implementation of the mirror algorithm yields different performance characteristics than the ones predicted by our stability analyses. Experimental data revealed that the mirror strategy was not adopted by human subjects (Dijkstra et al., 2004).

The common assumption is that preferred movements are maximally efficient in some energy measure. For instance, the preferred walking speed was determined to be "optimal" as it is consistent with minimum energy expenditure (Holt, Hamill, & Andres, 1990). Interestingly, the preferred and stable solutions in the ball bouncing task are not the most biomechanically efficient way to satisfy the task demands. To minimize the racket movement for an intended ball bouncing height, the ball-racket

collisions should be happen at the peak velocity of the racket, which corresponds to an impact acceleration of zero m/s^2. As previous studies repeatedly found that especially experienced bouncers opted for negative impact acceleration this suggests that when a less energy efficient solution affords performance stability, actors give priority to stability.

Despite the evidence for the attunement to passive stability, active control can also be expected, not too dissimilar from dynamic walking. This issue was addressed in an earlier experiment by de Rugy et al. (2003). In this study perturbations were applied to the ball trajectory to examine whether subjects relied on the passive properties of the ball-racket system or whether subjects modulated their racket movement to compensate for the applied perturbations. Based on the identified modulations of the racket trajectory, and the fast return to steady state behavior, active control was inferred. The results suggested that passive stability as offered from open-loop dynamics coexists with active control based on perceptual information about the error. This mixture of strategies is functionally advantageous for human movement control as it provides the actor both with stability in performance and flexibility in the face of large perturbations and varying movement contexts.

1.1.6. Motivation of Experiment 1: Passive Stability and Active Control in Response to Perturbations

While this study gave a first indication that actors "actively tracked passive stability", the perturbations were designed without knowledge about the basin of attraction. In the interpretation of the results it could not be distinguished whether a perturbation could have converged back to steady state or whether it was outside the basin of attraction because no theoretical analyses of the basin of attraction were available. Posthoc analyses also revealed that the perturbations in this study were relatively large and very small perturbations were excluded by design. To extend these first investigations Experiment 1 of the present thesis presents a derivation of the basin of attraction for the period-1 attractor of the ball bouncing model. The experimental perturbations were designed such that they covered a wide range both

inside and outside the basin. Furthermore, the coefficient of restitution of the racket, a critical variable influencing the shape of the basin of attraction, was also systematically varied. Based on the locations of perturbations in the basin, predictions about the response of a purely passive strategy could be made. These predictions are then compared with the actors' performance to elucidate the relationship between passive dynamics and active control.

1.1.7. Motivation of Experiment 2: Passive Stability and Active Control During Steady State Performance

Given these results and the findings in Experiment 1, it can be safely inferred that active control is utilized when noticeable perturbations are applied (de Rugy et al., 2003; Wei, Dijkstra, & Sternad. 2007, submitted). However, these results do not speak to the control regime that is used during steady state performance. Could active control also be present during unperturbed bouncing? One route into addressing this question is the analysis of small fluctuations in actor's performance. If no other control were present, the random fluctuations would die out according to the stability properties of the ball-racket system. A statistical way to summarize the return behavior of stochastic fluctuations is to analyze the covariance structure of the state variables. The state variables are velocity and time between impacts. To make predictions about the covariance structure as provided by the open-loop model, the deterministic ball bouncing model had to be extended by adding stochastic components. Following this further model development, analytical and numerical predictions could be derived and compared with similar data extracted from subjects' performance. Using the stochastic version of the bouncing ball map by Dijkstra et al. (2003) Experiment 2 aimed to compare human performance with the predictions of this model. As the model has different stability characteristics for different coefficients of restitutions, the experimental trials were performed with four different coefficients of restitution. Further, attempts to quantify the contribution of active control in this rhythmic goal-directed task will be made. To this end, the additional measure of covariation as developed by Müller and Sternad (2003) was applied. This

measure captures whether compensatory relations between relevant variables existed. Compensatory relation has been argued to be a signature of control (Scholz & Schöner, 1999; Scholz, Schöner & Latash, 2000).

Taken together the results of Experiment 1 and 2 will present a solid foundation to evaluate the blend of active and passive components in the performance of the rhythmic ball bouncing task.

1. 2. Learning, Adaptation and Transfer

Learning new motor skills is essential in human life as humans constantly face new challenges. This begins in early life when an infant has to master all the seemingly simple but essential skills, such as sitting upright, reaching to a toy, and walking. Evidently, such challenges continue throughout the life span when more complex tasks are acquired such as hand writing, driving a car, or riding a bicycle. Interestingly, humans seem to enjoy these challenges as they deliberately create them in their leisure time activities, including such "unnecessary" skills as ice skating or skateboarding. Therefore it is not surprising that skill acquisition, and more generally motor learning, has been studied for more than a century in experimental psychology, neurophysiology, physical education, and in the more recently emerging interdisciplinary research area of motor control. Numerous textbooks speak to the centrality of this topic: frequently adopted texts in current undergraduate teaching include the ones by Schmidt and Lee (2005), Magill (2001), Rose (1997), Rosenbaum (1992) and Shadmehr and Wise (2005).

Side-stepping the century-long history and focusing on the more recent 30 years of research studies in motor learning have been conducted under two theoretically and philosophically different frameworks. One line of investigation has the motor program concept as its central notion and is often referred to as the representational approach (Keele, 1968; Schmidt, 1975; Shadmehr & Wise, 2005). Another line of research has adopted the dynamical systems perspective where the emergence of

patterns is a central notion and tools of nonlinear dynamics figure prominently (Fowler & Turvey, 1978; Newell, 1986; Saltzman & Kelso, 1987; Kelso, 1995; Kugler & Turvey, 1987; Turvey, 1990; Warren, 2006). The present thesis is rooted in the latter approach.

1.2.1. The Representational Approach to Motor Learning

The motor program approach, and more specifically schema theory, is grounded in the computer metaphor where information processing provides the core framework to describe the functioning of the brain and human performance. Perceptual-motor skills are viewed as controlled by abstract representations of the task (Keele, 1968; Schmidt, 1975). Acquisition and adaptation of a motor skill is regarded as first obtaining a motor program followed by the tuning of its parameters with practice. The generalization of the original motor program concept in Schmidt's schema theory was proposed to account for two fundamental problems in motor learning: One is the novelty problem, i.e., how the actor learns a motor task that he/she has not known before. The second problem concerns storage, i.e., how could humans possibly store the details for every single movement acquired over the life span, assuming that the memory capacity of the human nervous system is not large enough. Numerous studies aimed to detail the elements of the generalized motor program or schema, the sensory feedback and its integration into the program (for an overview see Schmidt, 1982).

More recent advances in computational neuroscience similarly have abstract representations as their core concept. The concept of the internal model, central to numerous current studies, can be viewed as a further development of the generalized motor programs (Jordan, 1992; Warren, 2006; Wolpert & Kawato, 1998). In brief, the ability to perform an action is governed by an internal model that includes both forward and inverse models of the human body performing a movement. For a specific task, appropriate motor commands are generated by the forward model as inputs to the dynamic motor system to achieve some desired movement output. The reverse transformation from sensory variables to control variables is captured by the inverse model. Different to the schema theory that relates outcome measure like

reaching or throwing distance to performance measures like force, the mapping in the internal model includes more detailed knowledge of the dynamics of the biomechanical plant.

Acquiring a skill is regarded as the establishing and tuning the internal model for a given task. For example, research in computational neuroscience has been focused on how feedback calibrates the forward and inverse modules in the internal model (Jordan & Rumelhart, 1992; Kawato, 1990; Wolpert & Kawato, 1998). Using imaging tools such as functional fMRI and transcranial magnetic stimulation, the changes in the neuronal activities in motor cortex and the cerebellum have been associated with different functional components in the internal model (Imamizu et al., 2000; Wolpert, Miall & Kawato, 1998). This line of investigation is currently very active and more research will detail more of the postulated relationships and their anatomical and functional correlates.

1.2.2. The Dynamical Systems Approach to Motor Learning

A different approach to skill acquisition is presented by research grounded in the perspective of dynamic systems. Learning a motor skill is conceptualized as the process organizing the complex high-dimensional movement system into a task-specific device by exploring the constraints of the actor, the task, and the environment (Fowler & Turvey, 1978; Newell, 1986; Saltzman & Kelso, 1987). This task-specific organization has been referred to as a coordinative structure (Tuller, Turvey, & Fitch, 1982; Turvey, 1986) or, more recently, coordination dynamic (Mitra, Amazeen, & Turvey, 1998). Turvey, Mitra and colleagues rephrased the process of learning in the language of dynamical systems and dimensionality. In their proposition, the early phase of learning consists of discovering and establishing the relevant collective variable that captures the essence of the coordination pattern (Gelfand & Tsetlin, 1962). The most prominent example for a collective variable has been relative phase as calculated by the phase difference between two rhythmically moving limbs. At this stage the dimensionality of the dynamics is still high due to the independence of many subsystems. The intermediate phase proceeds with standardizing and

stabilizing this variable. In this second phase the dimensionality of the learned coordination dynamics is reduced as expressed by fewer active degrees of freedom. Active degrees of freedom are defined as the minimum number of first-order autonomous differential equation necessary to express the system behavior. The third phase involves a final step of increasing the determinism of the acquired coordination by reducing and organizing the "noise" processes originating from the more microscopic levels. This decrease in active degrees of freedom during skill improvement echoes the general idea of Bernstein (1967) that learning is reflected in the mastery of redundant degrees of freedom.

1.2.3. Learning Dynamics at the Example of Rhythmic Bimanual Coordination

Compared to the wealth of studies from the representational approach, there are relatively few studies that have taken a dynamical systems perspective to motor learning. One series of experiments, however, that directly followed the seminal work on the two-finger coordination paradigm, exemplifies the approach and demonstrates how a different perspective can uncover aspects of this complex problem that were previously less attended to.

On the basis of the Haken-Kelso-Bunz model (1985) for the coordination between two rhythmically moving fingers Kelso and Zanone (1992) examined the acquisition of a new phase relationship between two fingers. As established in previous studies, rhythmic interlimb coordination with 0 deg (inphase) and 180 deg (antiphase) are the two stable modes that are intrinsic to the human movement system. Extending from this fact Zanone and Kelso investigated the acquisition of the novel 90 deg phase lag between the two fingers. Hence, the 90 deg pattern is initially a relatively unstable pattern and had to be specified to the subject by a visual stimulus. The experiment showed that if a subject was unstable in the 90 deg pattern before practice, an abrupt change in the stability with a phase transition to the 90 deg pattern was observed with practice. On the other hand, if the to-be-learned 90 deg pattern was already stable in an individual, a gradual consolidation of the 90 deg pattern was observed. In addition, results also showed that further practice temporarily

destabilized the initially stable pattern of 180 deg. In a subsequent study Zanone and Kelso (1997) found that learning the 90 deg pattern also made the 270 deg phase pattern, the symmetrical counterpart of pattern of 90 deg pattern, simultaneously stable. These findings highlight how existent behavior is modified by the new to-be-learnt pattern and that practice not only creates new stable patterns, but also influences the existing dynamics and potentially changes the entire layout of coordination dynamics. Similar conclusions have been reached in a study of learning a postural pattern with novel relative phase between the hip and ankle joints (Faugloire, Bardy, & Stoffregen, 2006).

1.2.4. Choice of Variables and the Coordinate System

Assessing the change of a certain outcome variable over time has been central to much inquiry in the motor learning literature. Over a century of studies on motor learning revealed many different types of learning curves, i.e., changes of a dependent measure with practice. These curves include exponential, hyperbolic, S-shaped, discontinuous, and power law functions, where the latter have been the most prevailing ones (Newell & Rosenbloom, 1981; Liu, Mayer-Kress, & Newell, 2003, 2006). Newell and colleagues (2001) proposed a dynamical account of changes in learning recognizing that the different learning functions may be an expression of changes on multiple time scales (see also Kugler, 1986; Port & van Gelder, 1995; Thelen & Smith, 1994). For instance, while an exponential function is a reflection of change at a single time scale, the power law function characterizes a change that is organized over all time scales, i.e., scale-free (Schröder, 1991). Newell and colleagues suggested that all learning functions can be realized by viewing the learning process as a migration of the system's trajectory on an evolving attractor landscape, as demonstrated by their simulations. This proposition attempts a unified framework to accommodate many different observations under one umbrella perspective.

What is not explicitly addressed but implicit in their arguments is that the time course of a learning curve is not only dependent on the level of analysis but also highly dependent on the chosen variable. It is noteworthy that in most studies on

motor learning outcome variables have served as the predominant index to assess motor learning (Schmidt & Lee, 2005). Typically, an error score or its variance suggested itself for study as the instruction emphasized this variable. This is understandable as the experimental work on motor learning was largely driven by practical questions such as how to train skills more effectively, how to best deliver knowledge of results to facilitate learning, and how to enable transfer across effectors and task. The focus of attention was on the effect of manipulations on relevant measures of success and little heed was paid to the issue of whether the measured variable was the best reflection of the change in the underlying processes of control. Smeets (2000) has raised this issue demonstrating that the time course of learning is dependent on the dependent measure, pointing to the example of movement time or velocity as frequently used dependent measures. Müller, Frank, and Sternad (2007) similarly demonstrated for a fictive data set how the quantification of the dependent measure can significantly alter the time course. These exemplary demonstrations highlight the importance of the chosen measure, especially when control is at issue.

1.2.5. Relation between Movement Execution and Outcome

In previous work from a dynamical systems perspective this question was regarded as mute or "solved". As the relatively few studies on learning directly built on the seminal work on phase transitions in interlimb coordination which identified relative phase as the collective variable, it was unquestioned that this same variable was the best one to assess the learning process. However, this issue is an interesting problem itself and needs to be addressed when examining new tasks where relative phase is not applicable. The ball bouncing task of the present investigation raises such a problem. However, based on a long series of previous work we suggest that the variable acceleration of the racket at impact constitutes a good candidate measure as it characterized the control strategy of exploiting passive stability. In fact, Sternad et al. (2001) argued this variable to have the status of an essential variable in the sense of early theorizing of Gelfand and Tsetlin (1971). Hence, learning will be assessed by changes in this variable reflecting the change in sensitivity to passive stability, i.e.

characteristics of control.

A second and related point is that most studies in motor learning have evaluated performance improvement by analyzing one or few outcome measures. Typically, an error score or its variability suggested itself for study as it directly reflected the instruction or the task goal. However, in the wake of Bernstein it is important to highlight that execution variables describing the behavior of the action system may be similarly or even more important to reveal the underlying processes of control. The frequently cited example of an expert hammering on an anvil is informative: while the endpoint trajectory is relatively invariant for experienced performers, the joint angle trajectories of the arm may be significantly more variable (Bernstein, 1967; Turvey & Carello, 1995). This at first sight paradoxical observation is explained by the compensatory relationship between the variables: a deviation in one angle is compensated for by a variation in the other angle that offsets the error seen in a third result variable. The lesson to learn from this example is that execution variables may provide the suitable coordinate system in which changes can be analyzed.

The present thesis will investigate performance changes with practice in both its movement outcome and the space of execution variables. In the hammering example by Bernstein movement outcome is the accuracy of the hitting action or, when viewed over time, the trajectory of the tip of the hammer; movement execution refers to the combination of joint angles at the moment of hit or, over time, the trajectories of the joint angles. Essential to this decomposition into two causally related sets of variables is that there is a redundancy, a many-to-one mapping or a many-to-many mapping. Hence, there is a set of different solutions that lead to the same performance and that satisfy the task criteria. The subset in the space of the execution variables that "solves" the task perfectly will be referred to as the solution manifold, a geometric representation of the multiple combinations of execution variables that lead to the same outcome variable. Important for this analysis is that the functional relationship between execution and result is deterministic and known. For the example of multi-joint limb movements the joint angles and the execution of the endeffector

position in external space are connected by a nonlinear function.

The ball bouncing task is a different kind of behavior that also lends itself for analysis from this point of view. Focusing on a single bounce - the racket hitting the ball to achieve a specific target height -, and performing a basic mechanical analysis it can be easily seen that there are three variables that completely determine the subsequent amplitude and error of the ball with respect to a target: racket velocity at ball impact, ball velocity before impact, and impact position relative to the target height. Different combinations of these three execution variables lead to the same result, ball height. The three execution variables form the execution space where the set of perfect solutions form the solution manifold. Performance with a given error form "iso-error bands" around this manifold. Once, the task is represented like this learning can be viewed as the migration of solutions in execution space with respect to this solution manifold. Experiments 3 and 4 will adopt this perspective and analyze learning, adaptation, and transfer in the ball bouncing task within this coordinate representation. Note, in contrast to how ball bouncing was introduced earlier, this analysis perspective analyses each bounce as a separate event.

1.2.6. The TNC-Method: Decomposition of Variability in a Redundant Task

Müller and Sternad developed the so-called TNC-method to analyze the distribution properties of a set of data in execution space with respect to the solution manifold (Müller & Sternad, 2003, 2004; Müller, Frank, & Sternad, 2007; Smeets & Louw, 2007). In the ball bouncing task every contact is a data point, defined by the triplet of execution variables. This data point has a value in result space, the ball's height error. The series of bounces in one trial constitute the sets of data that are analyzed with respect to the solution manifold. The TNC-method decomposes the distributional properties in a set of trials into three components that quantify changes in the result of these data sets over time: improvement originating from exploration and migration in execution space to error-tolerant locations on the solution manifold (called tolerance *T*), improvement due to covariation between execution variables (called covariation *C*), and reduction of stochastic elements in execution (called noise

N). The latter component noise reduction is, of course, a widely acknowledged phenomenon for performance improvement in skill learning. The second component covariation captures compensatory relations between execution variables that achieve a better movement result, in the spirit of the hammering example discussed above. Tolerance is conceptualized with the recognition that certain combinations of execution variables are more tolerant to error than others. This expresses the notion of sensitivity where small deviations in execution may lead to large errors in result (high sensitivity, low tolerance). Improvements in performance, after some exploration of the task, are then related to the migration of performance to more "error-tolerant" locations on the solution manifold.

This approach is similar in spirit to another line of research, the method of the "Uncontrolled Manifold" (UCM) that has developed a different analysis method to examine variability in redundant movements (Scholz & Schöner, 1999; Scholz, Schöner & Latash, 2000). Central to the UCM approach is that good performance is characterized by variance in execution space that aligns with the tangent of the manifold at mean performance; variance orthogonal to this linearization of the manifold reflects unwanted "motor noise". The alignment of variability in a selected direction was interpreted as the signature of control (Yang & Scholz, 2005; Yang, Scholz, & Latash, 2007).

Analysis of the structure of variability across learning can reveal the relative contributions of conceptually distinct components during the learning process in the coordinate space. Conceptually, this approach to the learning process has a kinship with that of Newell and colleagues (2001) who emphasized that the learning should be interpreted as a migration of the system's trajectory on an evolving attractor landscape. The difference is that the space within which this migration happens is different. For the TNC-method the space is defined by execution variables, where the "potential landscape" is given by the result. Further, it is the distributional properties that are center stage in the TNC-method. Analysis of the structure of variability can highlight new thus far ignored aspects in learning. Hence, the TNC analysis will be

performed in the current thesis for learning and adaptation (Experiment 3) and transfer (Experiment 4).

1.2.7. Motivation of Experiment 3: Variability and Stability during Learning and Adaptation

Having laid the foundation for the analysis of stability and a novel analysis of the distributional properties of execution variables in the redundant task, Experiment 3 will examine the time course of skill improvement in the task of ball bouncing. In addition, having established support for this description of the task in previous experiments, distortions will be introduced in execution space. The goal of this manipulation is to test how actors adapt their strategies to such novel conditions and search for new solutions. Central to the question is how distributional properties tolerance, covariation and noise change. How is stability reestablished after a distortion?

1.2.8. Motivation of Experiment 4: Variability and Stability during Learning and Interlimb Transfer

Acquiring a new skill is only really valuable if such a learnt skill will adapt and generalize to new contexts. Hence, the study of adaptation and transfer is essential and can also retrospectively shed light on the nature of the learning process. An example for such an adaptation was studied in Experiment 3. Experiment 4 extends this question and asks whether the acquired skill of ball bouncing will be transferred to being performed with another limb. Such intermanual transfer will reveal to what degree the skill is peripheral and tied to the specific effector that performs the task or whether the skill implies a more central component that is easily transferred to another limb.

Transfer was a core issue in experimental work on motor learning in the 1960s until 1980s when research was driven by more practical questions (Adams, 1987; Schmidt, 1982).

When tailored more to the conceptualization of the task at hand, the question is

whether learning the system dynamics is similar to obtaining an abstract presentation of the task, the knowledge of exploitation of passive stability learnt from one effector should be transferable to another effector. The Experiment 4 in this thesis will investigate this hypothesis.

CHAPTER 2

Research Questions

Experiment 1: Passive Stability and Active Control in a Rhythmic Task

Previous studies on passive stability in ball bouncing had already documented that the actor is sensitive to the stability properties of the task dynamics and performs the rhythmic actions with a strategy where perturbed performance converges back to steady state without requiring error-correcting racket movements. This behavior is consistent with the predictions from stability analyses of a dynamic model of the task. However, how much stability the task dynamics can offer and perturbations of what magnitudes can be accommodated by using passive stability alone has not been addressed. The first aim of Experiment 1 is to reveal the actor's sensitivity to the global layout of attractor space for a specific task. The second aim is to discover the control strategy when passive stability cannot afford stabilization for large perturbations. The hypothesis is that humans are not only sensitive to task dynamics and deliberately utilize it, but are simultaneously flexible to apply active control when needed. The basin of attraction for the period-1 attractor of the model will be derived to make quantitative predictions about convergence based on the passive dynamics. Human behavior will be compared to these model predictions to reveal the control strategies for accommodating perturbations of different magnitudes.

Experiment 2: Fluctuations and Passive Stability in Ball Bouncing Task

Given the findings of a previous study and Experiment 1, it can be concluded that active control is utilized when noticeable perturbations are applied (de Rugy et al., 2003; Wei, Dijkstra, & Sternad. 2007, submitted). However, when there are no perturbations as in steady-state performance, does the actor solely rely on the automatic correction or blend active control to compensate for small fluctuations? The first aim of Experiment 2 is to demonstrate the presence of feedback-based control blended with passive strategy. It has long been recognized that fluctuations are

inevitable in human performance and often times researchers model them as random noise. To reveal the signature of control from seemingly random fluctuations of the performance, the covariance structure of state variables of a stochastic model, which captures the passive task dynamics, will be compared to actual human performance. The second aim of the study is to investigate whether the actor will resort to more active control as passive stability of the task system decreases. The degree of stability is manipulated by changing the coefficient of restitution. We will use covariation between execution variables of the task as an index to quantify the active control. The result of this experiment will not only elucidate the relative contribution of passive stability and active control for steady-state performance, but also further our understanding about the nature of fluctuations in rhythmic movements.

Experiment 3: Variability and Stability during the Acquisition and Adaptation of a Rhythmic Skill

Previous studies had illustrated that during acquisition of the ball bouncing skill the increasing attunement to passive stable is accompanied by a reduction of variability in performance. While this finding is consistent with the common belief that variability is equal to the inverse of stability, it does not establish the causality relationship between variability and stability in human performance. The first aim of Experiment 3 is to replicate that skill improvement is indeed characterized by utilizing passive stability but additional to show that the structure of variability changes independently. The second aim is to demonstrate learning process of a redundant task can be expressed as systematical changes of distributional properties of the variability of the execution. Applying a novel analysis method the learning process will be evaluated in the execution space, which is spanned by relevant execution variables that lead to the movement outcome variable. A set of individual events will be represented as a distribution in the execution space that contains a manifold of solutions. The third aim of Experiment 3 is to demonstrate that acquisition and adaptation of a learned skill can be characterized by exploration and optimization of

stability properties as analyzed in this space. To induce adaptation the solution manifold will be modified such that actors have to change their strategy and re-establish the passively stable regime.

Experiment 4: Variability and Stability in Learning and Interlimb Transfer

One core issue in motor learning is whether a newly learned skill can be adapted and generalized to other situations or other effectors. The typical indicators to evaluate and quantify adaptation and transfer are movement outcome measures. The aim of Experiment 4 is to demonstrate that utilizing passive stability as originating from task dynamics is subtle and abstract in nature. As a result, the passive strategy is expected to be transferred across two limbs and the transfer is symmetrical. The second aim is to demonstrate the distributional properties of movement execution are also transferable between different effectors. To achieve this, Experiment 4 will also adopt the variability decomposition method to further unravel the structure of variability changes during learning and transfer. This study will provide insights into the nature of transfer by focusing on the changes in the stability properties and of the structure of variability.

CHAPTER 3

Passive Stability and Active Control During in Response to Perturbations

Chapter 3 contains a paper submitted to Journal of Neurophysiology

"Passive Stability and Active Control in a Rhythmic Task"

Kunlin Wei[1], Tjeerd M.H. Dijkstra[2], Dagmar Sternad[1]

[1] Department of Kinesiology and Integrative Biosciences, Pennsylvania State University

[2] University Medical Center, Leiden, The Netherlands

Abstract

Parallel to research on passive dynamic walking rhythmically bouncing a ball with a racket has been demonstrated to be a task that affords a passively stable solution. Results from stability analyses of a mechanical model and accompanying empirical results from human performance support the conclusion that actors exploit this passive dynamics of the task thereby reducing computational demands of the task. Passive stability is characterized as solutions where perturbations converge back to the steady state without active corrective control. The present study investigated the response to systematic perturbations of 14 different magnitudes, in conjunction with predictions derived from the basin of attraction derived from the model. Given that predictions depended on the coefficient of restitution in the model, four different conditions were experimentally tested. Three predictions were tested in a virtual reality set-up with a haptic interface. Relaxation times estimated from the performance errors showed significantly faster returns than predicted from the purely passive model, indicative of active error correction. However, the pattern of relaxation times was qualitatively consistent with predictions based on passive stability: relaxation times were larger for larger perturbations (Prediction 1); the return was shorter for positive and longer for negative perturbations (Prediction 2). For higher coefficients of restitution relaxation time for positive perturbations was longer (Prediction 3). Racket accelerations at contact were negative for all perturbations except the three most negative ones, indicating passive stability. It was concluded that stability properties of the task dynamics are a factor in movement control. During unperturbed performance humans exploit passive stability afforded by the task minimizing the control effort. When facing perturbations, participants show sensitivity to whether perturbations require corrections or whether passive stability is sufficient to maintain the stability.

Introduction

Human actions necessarily express themselves as physical interactions with the environment where the biological system creates and is subject to forces that constrain the execution and the control of action. The stability of these forces is key to their control. Stability in postural and movement control is achieved by many different means: physiologically, it can arise from the stiffness of the muscular and tendonous tissues mediated by stretch reflexes; from a control perspective, stability is attained through feedback-based responses that correct for errors; a third source of stability is of mechanical and dynamical origin: the stability of the actor-environment system as it is set up by the task. This system may also provide stability "for free" that is not brought about by the active compensations of errors, i.e., feedback control, but rather it is brought about by the inherent stability of the task. A prominent example that has demonstrated the importance of this source of stability is the passive dynamic walker. McGeer (1990) showed, by using both an analytic model and a multilink robot, that a purely mechanical system without actuators and control could maintain a stable walking pattern when walking down a gentle slope.

Models of this type achieve dynamic stability without active control, relying solely on the passive dynamics of the physical constructs. Over the past two decades, a lot of attention and effort has been directed to applying the idea of passive dynamics to the design of human-like biped robots walking on level ground (Coleman & Ruina, 1998; Collins, Wisse & Ruina, 2001). Several models demonstrated that walking in two dimensions has inherent stability but to achieve walking in three dimensions additional actuation was needed. However, the fundamental reliance on passive dynamics has offered these 3D walkers efficiency in energy expenditure, reduced demands in control, and elegant mimicry of human motion (Collins et al., 2005; Tedrake et al., 2004). The striking similarity of these minimally controlled walking machines with human walking suggests that passive dynamics may play an important role in shaping coordinated human behavior. Without downplaying the importance of

muscular forces and their tuning by perceptual information in determining behavior, the passive dynamic walkers show that our understanding about control can be deepened by studying the motion that may emerge without control.

To tease apart the contributions of passive dynamics and active control, a motor task is wanted that first affords such a passively stable solution. Sternad and colleagues have shown this to be the case for the task of rhythmically bouncing a ball on a racket (Sternad et al., 2000, 2001; Dijkstra et al., 2004; Schaal, Sternad, & Atkeson, 1996; de Rugy et al., 2003). This task requires an actor to bounce the ball with a racket up in the air to a consistent height over repeated bounces. By viewing the racket as an oscillating planar surface and the ball as a point mass colliding with the racket with an inelastic impact, a simple mechanical model was derived for the ball bouncing task. Stability analyses of this model yielded predictions about criteria for which this passive stability is achieved.

Empirical studies of the ball bouncing task confirmed that humans exploit the passive stability properties as predicted by the model. Experienced actors hit the ball with negative racket accelerations in a variety of experimental conditions: when different ball amplitudes were required (Schaal et al., 1996), when the movement of the ball was confined to the vertical dimension only (Sternad et al., 2000, 2001), or when the ball and the racket moved freely in three dimensions; when the ball was bounced by using a paddle moving downward instead of a hand-held racket moving upward to hit the ball (Schaal et al., 1996), or when the experiment was conducted in a virtual reality setup (de Rugy et al., 2003). In contrast, novice actors performed with positive impact acceleration and only gradually, within approximately 30 minutes of practice, tuned their movements to use negative impact acceleration (Dijkstra et al., 2004). This result was accompanied by decreasing variability supporting the interpretation that performance had improved with this change of strategy. This latter result highlighted that hitting the ball with negative impact acceleration was not an intuitive or trivial solution for actors. The task offers the advantage of stability but it has to be learnt. In sum, the empirical results on the ball bouncing task support the

interpretation that actors exploit the stability properties of the task.

This ball bouncing model shares many features with the passive dynamic walking model. First, both models are formulated over the actor/acting system and the environment that it is in interaction with. Central to the walker are its collisions with the ground, central to the ball bouncing model are the ball-racket collisions. The models are based on mechanics without neuromuscular control mechanisms. Second, the only forces considered are the gravitational and the collision forces. The original 2D passive dynamic walking model is only governed by gravity and inelastic foot-ground contacts. Similarly, in ball bouncing the ball flight is governed by the gravitational force and the inelastic ball-racket impact. Third, both models permit intermittent control at the moment of contact: step-to-step adjustment in walking are only controllable at the moment of foot-ground contact since in the swing phase the leg movement can be viewed as a pendulum (Mochon & McMahon, 1980); equivalently, the ball-racket contact is the only moment when the actor can control the ball flight. Fourth, since both tasks are periodic, the stability analyses used the same Poincare section technique to relate the system at one bounce/step to the next bounce/step. Fifth, both tasks have multiple stable solutions. The passive dynamic walker has a period-1 solution and period-*n* solutions corresponding to a limping gait. Similarly, the ball bouncing model has a period-1 solution but also period-*n* solutions and the so-called sticking solutions where the ball sticks to the racket. This sticking solution corresponds to the situation when the walker falls down.

The presence of multistability immediately raises the question about the boundary between these multiple solutions. How large can a perturbation be before the system ends up in another solution? How large is the basin of attraction? Results of passive dynamic walkers show that the basin of attraction is not very large, even with meticulous tuning of the parameters (Schwab & Wisse, 2001; Garcia et al., 1998). As a result, the bipedal robots are sensitive to initial conditions and demand a careful launch. Even a small perturbation like a small bump on the slope can destabilize the robot and make it fall (Schwab & Wisse, 2001; Wisse & van Frankenhuyzen, 2003).

These observations are in contrast to human walking which is much more robust to perturbations and adaptive to different conditions. This contrast points to the presence and importance of active control in human walking.

Although multistability similarly exists in the ball bouncing map, only period-1 performance has been investigated thus far. Parallel to the line of investigation in passive walking, further understanding about the relationship between passive stability and active control can be obtained by examining the basin of attraction (Schwab & Wisse, 2001; Garcia et al., 1998). Hence, the present study investigates the ball bouncing task with periodic solutions under systematic perturbations that are designed in view of the basin of attraction for the period-1 solution. Will the actor discard the strategy of using passive stability and turn to active control with error feedback on a cycle-to-cycle basis? Alternatively, will the actor rely on passive stability without active error corrections when the ball is only slightly perturbed, with perturbations inside the basin of attraction? Or will the actor adopt a mixture strategy with both the exploitation of passive stability and of perception-guided error correction? In sum, are actors sensitive to the boundaries of the basin of attraction of the period-1 solution?

These questions can be answered by applying perturbations of different magnitudes. In a previous study de Rugy et al. (2003) applied perturbations by randomly changing the coefficient of restitution of the racket upon impacts leading to unexpected under- or overshooting of the ball with respect to the target height. Results showed that actors quickly reestablished negative impact acceleration at impact, indicating the use of passive stability. Yet, for all the applied perturbations modulations of racket movements also indicated signs of active control. While this study gave a first indication that actors “actively tracked passive stability”, the sensitivity to the basin of attraction could not be tested because no theoretical analyses of the basin of attraction were available. Posthoc analyses revealed that the perturbations were relatively large and mostly took the system outside the basin of attraction. In fact, very small perturbations were excluded by design. To extend these

first investigations the current study provides a derivation of the basin of attraction for the period-1 attractor of the ball bouncing map. The experimental perturbations were designed such that they covered a wide range both inside and outside the basin. Furthermore, the coefficient of restitution of the racket, a critical variable influencing the shape of the basin of attraction, was also systematically varied. Based on the locations of perturbations in the basin, predictions about the response of a purely passive strategy can be made. These predictions are then compared with the human performance to elucidate the relationship between passive dynamics and active control.

The Model

The predictions are based on the same model that was derived in earlier work (Dijkstra et al., 2004; Guckenheimer & Holmes, 1983; Holmes, 1982, Sternad et al., 2001; Tuffilaro, Abbott, & Reilly, 1992). As mentioned above, the task of periodically bouncing a ball with a racket to a target height affords a passively stable period-1 solution. This solution is entirely open-loop and, once initiated, requires no control or error correction. The period-1 solution co-exists with other solutions, in particular sticking solutions and period-*n* solutions. The latter solutions give rise to the period-doubling route to chaos and were an important motivation for the initial studies of the map (see Tuffilaro et al., 1992). However, they are of no concern in the current context.

The ball bouncing map is based on the following three assumptions: 1) *Ballistic flight*: between the *k*-th and the *k+1*-th bounce the vertical ball position $x_b(t)$ follows the ballistic flight equation:

$$x_b(t) = x_b(t_k) + v_b^+(t - t_k) - (g/2)(t - t_k)^2 \qquad t_k < t < t_{k+1} \qquad (1)$$

with $x_b(t_k)$ the vertical ball position at the time of the last (*k*-th) impact, v_b^+ the ball velocity immediately after impact, and *g* the acceleration due to gravity (9.81 m/s^2). 2)

Instantaneous impact: the impact is instantaneous such that the ball velocity immediately after impact v_b^+ is determined by:

$$(v_b^+ - v_r) = -\alpha(v_b^- - v_r) \qquad (2)$$

where v_b^- and v_b^+ denote ball velocity just before and after impact, respectively. v_r denotes the velocity of the racket at impact and α denotes the coefficient of restitution which captures the energy loss at the impact. The velocity of the racket does not change during impact because the mass of the racket is much larger than the mass of the ball. 3) *Sinusoidal racket movement*: The racket movement is a pure sinusoid:

$$x(t) = a_r \sin \omega_r t \qquad (3)$$

with amplitude a_r and angular frequency ω_r.

The validity of these assumptions for bouncing a physical ball is discussed in Dijkstra et al. (2004) and Brody et al. (2002). Since the task is performed in a virtual set-up (see methods below), the ballistic flight and impact assumptions are satisfied by design. The assumption of a pure sinusoid is not obeyed, not even in the virtual set-up: actors pick a periodic waveform with a slope that is slightly steeper than a sine wave in the ascending phase when the racket meets the ball. Unfortunately, there is no simple mathematical description of this waveform. However, we note that only position and velocity with which the racket hits the ball determine the ball trajectory. Thus, for mathematical simplicity we use an equivalent sinusoid that is close to the actual waveform at the impact. The equivalent frequency of this sinusoid is calculated from the period between bounces. The equivalent amplitude of this sinusoid a_r is calculated from the stationary phase of impact θ as (see Dijkstra et al., 2004, eq 8):

$$a_r = \left(\pi \frac{1-\alpha}{1+\alpha}\right) \frac{g}{\omega_r^2 \cos\theta} \qquad (4)$$

From these assumptions, the ball bouncing map can be derived as:

$$v_{k+1} = (1+\alpha)a_r\omega_r \cos\theta_{k+1} - \alpha v_k + (g\alpha/\omega_r)(\theta_{k+1} - \theta_k) \quad (5)$$

$$0 = a_r\omega_r^2(\sin\theta_k - \sin\theta_{k+1}) + v_k\omega_r(\theta_{k+1} - \theta_k) - (g/2)(\theta_{k+1} - \theta_k)^2 \quad (6)$$

This is an implicit map with the two state variables v_k, the ball velocity just after impact, and θ_k, the racket phase of impact. The ball bouncing map has a period-1 attractor which is locally linearly stable when the acceleration at impact, denoted by *AC*, is bounded by:

$$-2\frac{1+\alpha^2}{(1+\alpha)^2} < AC < 0 \quad (7)$$

The period-1 attractor co-exists with other attractors: for *AC* more negative than the lower boundary, there exist attractors where the ball sticks to the racket for part of the cycle; for positive *AC* several attractors exist, among them the period-doubling route to chaos (Tuffilaro et al., 1992).

The domain of attraction of the period-1 attractor depends on the parameters of the map: *g*, the acceleration of gravity, α, the coefficient of restitution, ω_r, the racket frequency, and a_r, the racket amplitude. However, these parameters can be tightly controlled or determined in the experiment to obtain a good quantitative match with the model. The first two parameters, *g* and α, are independent of the actor and can be experimentally manipulated. Since the linear stability in the model and the domain of attraction depend strongly on α (see Figure 1), this parameter was varied as a dependent measure in the current study. The racket period and amplitude are more difficult to control experimentally, since they do depend on the actors' performance. However, the racket frequency ω_r can be fixed by having actors bounce to a visual target. Since the ball amplitude determines the racket period (through the flight equation) and actors hit the ball at an approximately constant height relative to the floor, having a target at a fixed height relative to the floor fixes the racket period. The racket amplitude was not prescribed but the actual movement amplitude was estimated from the impact phase, using eq 4. The model parameter a_r was set accordingly for the calculations.

The domain of attraction of the period-1 solution was calculated by numerically iterating the ball bouncing map (eqs. 5 and 6). The initial conditions and the parameters of the map were chosen to be as close as possible to the experiment. In detail, the initial velocity values were taken from the range 2 to 4.5 m/s (y-axis in Figure 1) and the initial phase values were taken as the stationary phase for that particular coefficient of restitution. Gravity g was 9.81 m/s^2 and the racket frequency ω_r was $2\pi/0.65$ rad/s, which is based on the grand average of the time between bounces over all actors and conditions. The grand average of impact phase varied with the coefficient of restitution α. For the experimentally used values of α, 0.5, 0.6, 0.7, and 0.8, the respective phase values were 6, 7, 8 and 13 deg. From these phase values the equivalent racket amplitudes a_r were calculated using eq. 4. For values of α between the experimentally used ones, spline interpolation and extrapolation was applied to obtain phase values. With these initial conditions and parameters the ball bouncing map was iterated 100 times. First, it was tested whether the state variables were close to the stationary ones; criterion for closeness was the band around the stationary values of 0.1 m/s for velocity and 10 deg for phase. These values for the bandwidth equaled the observed variability during unperturbed bouncing. The results showed that the domain of attraction does not strongly depend on these values.

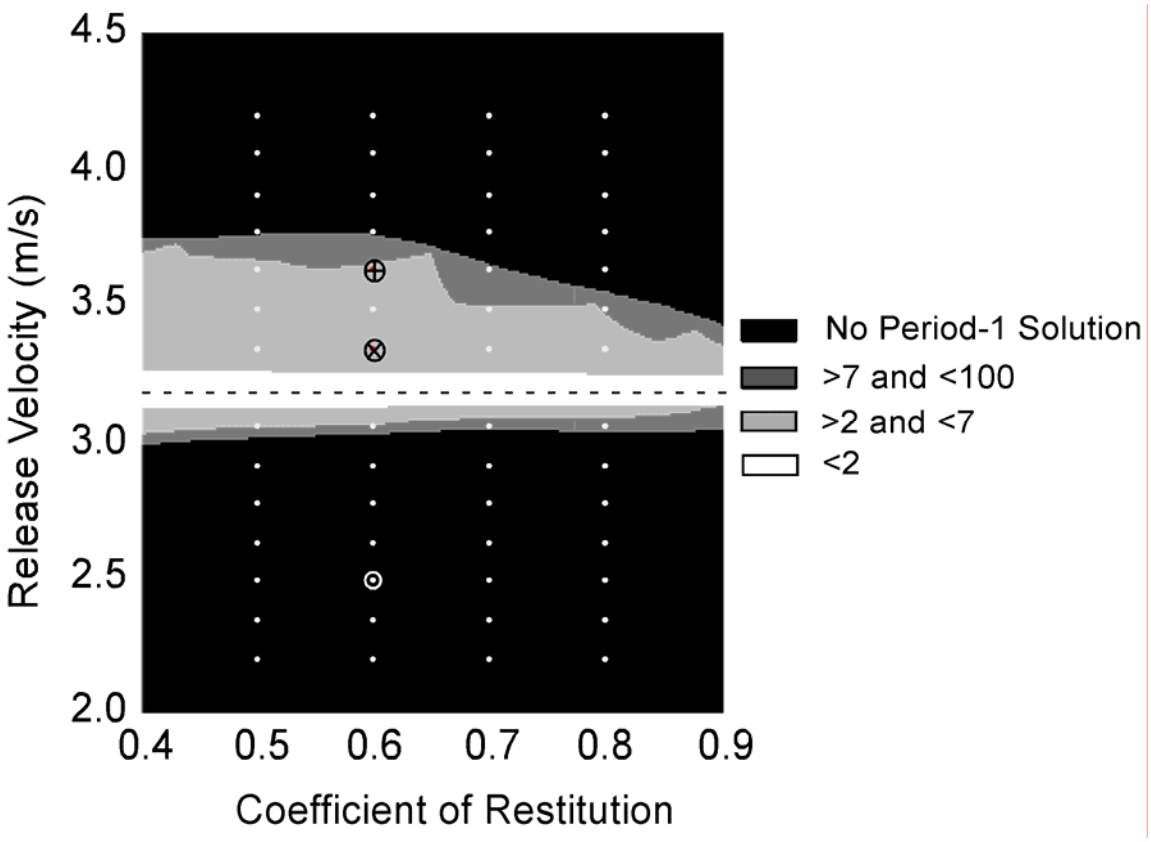

Figure 1: Basin of attraction of the ball bouncing map. The dashed horizontal line denotes the stationary release velocity. The grey shading indicates how many bounces or cycles are needed for the perturbations to die out. The white dots denote perturbations used in the experiment: 14 different magnitudes for each of the four coefficient of restitution. The three larger dots highlight three selected perturbations for $\alpha = 0.6$ for which the time course is illustrated in Figure 2, following the ordering from the top to the bottom.

The results of these computations are presented in Figure 1, where the white and grey shaded areas denote the domain of attraction with increasing relaxation times (counted in number of cycles). The black area indicates (a) initial conditions that did not converge to the stationary state within 100 iterations, or (b) a sticking solution where the time between impacts was smaller than 1 ms. The four vertical columns of dots in Figure 1 indicate the different perturbation magnitudes that were used for the four different α values in the experiment. Figure 2 shows some simulation results for

the three marked selected perturbation magnitudes in Figure 1. Panels A and B illustrate the time course of the two state variables at impact following a large perturbation that takes 7 bounces to return to steady state. Equivalently, panels C and D show the state variables for a smaller perturbation that relaxes back within 4 bounces. The continuous time series on panels E and F illustrates how a sticking solution occurs.

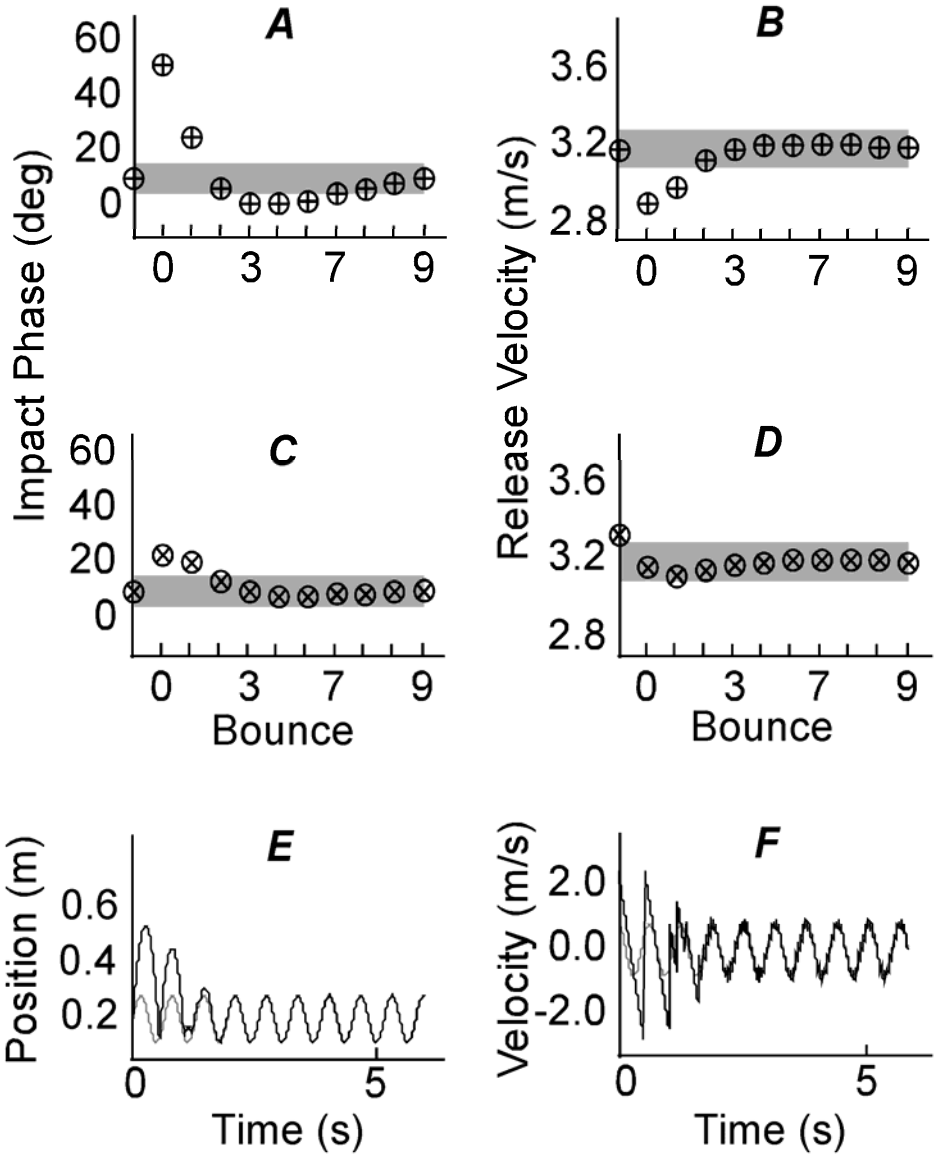

Figure 2: Simulations for three selected perturbations of the ball bouncing map. **A and B**: Time evolution of the state variables impact phase and ball release velocity, respectively. The shaded area denotes the bandwidth when the system is considered to

be back at equilibrium. These two panels correspond to the perturbation marked as a cross in Figure 1. **C and D**: Equivalent plots for the perturbation marked as x in Figure 1. **E and F**: Time series of continuous position and velocity of both the racket and the ball, showing a sticking solution. This was obtained for the perturbation marked as a white circle in Figure 1.

From these analyses the following predictions can be formulated:

Prediction 1: The basin of attraction has a boundary separating stable period-1 solutions from sticking and period-*n* solutions. Actors are sensitive to this boundary and rely on passive stability when perturbations are inside the basin of attraction. For perturbations outside the basin, they adopt an active strategy seen in qualitative changes in the racket kinematics that aim to correct for errors.

Note however that sticking or period-1 solutions have never been observed in the current experimental set-up (de Rugy et al., 2003). Hence, if this discontinuous change in strategy is not seen, we still expect qualitative changes in behavior as a function of perturbation magnitude and coefficient of restitution. Therefore, three more qualitative predictions can be formulated:

Prediction 2a: With increasing perturbation magnitude, the time for returning to steady state performance increases for all coefficients of restitution. Larger perturbations that take the system further out of the basin of attraction lead to longer relaxation times.

Prediction 2b: The relaxation time is longer for negative perturbations than for positive perturbations for all coefficients of restitution (i.e., for release velocities smaller than the average release velocity). This follows from the observation that the lower boundary of the basin of attraction is closer to the stationary state than the upper boundary.

Prediction 3c: For positive perturbations the smaller coefficients of restitution have a wider basin of attraction. Hence, for positive perturbations the smaller coefficients of restitution should show faster returns than the higher coefficients of

restitution. For negative perturbations there should be no difference in relaxation time for different coefficients of restitution.

Method

Participants. Seven volunteers participated, with ages ranging from 23 to 47 years. All participants 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 Regulatory Committee of the Pennsylvania State University.

Experimental Apparatus. In the virtual reality setup, participants manipulated a real table tennis racket in order to bounce a virtual ball that was projected on a screen in front of them (Figure 3). Participants stood about 0.5 m behind a back-projection screen with width 2.5 m and height of 1.8 m. 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 1024 by 768 pixels with a 60 Hz refresh rate. Accelerations of the racket were measured using a solid state piezoresistive accelerometer mounted on top of the racket (T45-10, Coulbourne). The mechanical brake on the rod attached to the racket was controlled by a solenoid (Magnet-Schultz type R 16x16 DC pull, subtype S-07447). A light rigid rod with three hinge joints was attached to the racket surface and ran through a wheel whose rotation was registered by an optical encoder (see Figure 3). Its accuracy was one pulse for 0.27 mm of racket movement. The pulses from the optical encoder were counted by an onboard counter (DT322). The racket could move and tilt with minimal friction in three dimensions but only the vertical displacement was measured. Images of racket and ball position were shown on-line using custom-made software. The delay between real and virtual racket movement was measured in a separate experiment and found to be 22 ms on average with a standard deviation of 0.5 ms.

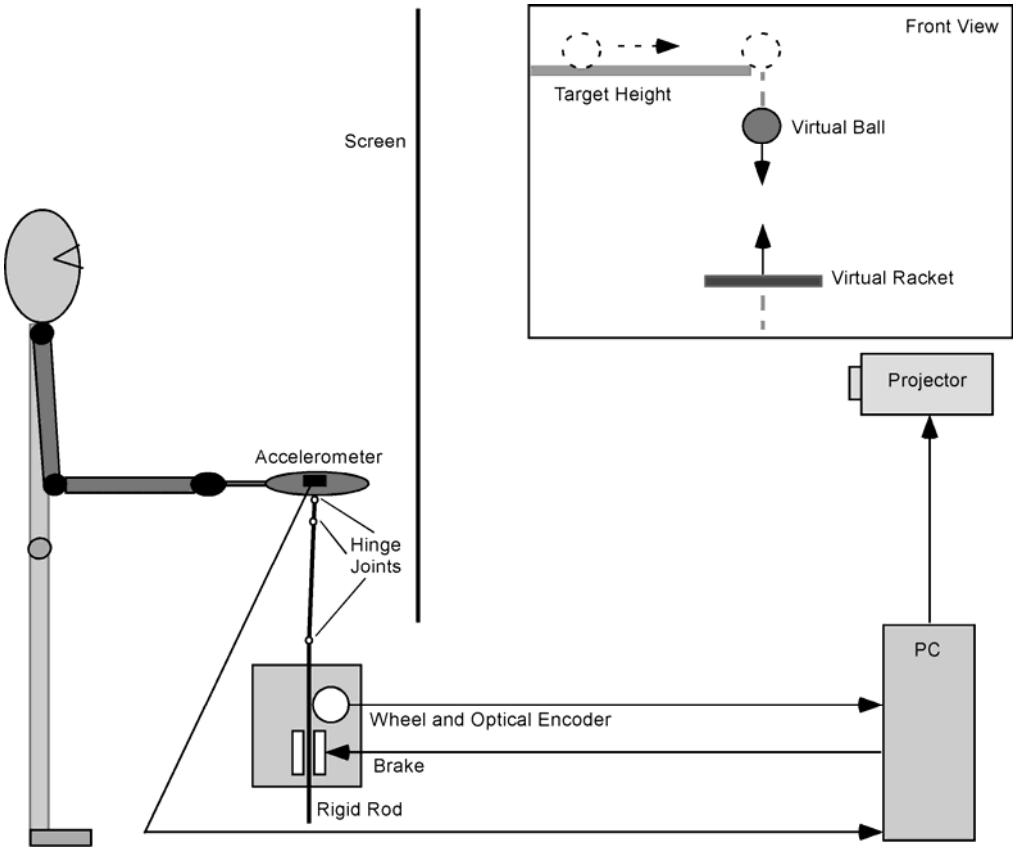

Figure 3: Virtual reality setup for the ball bouncing task, a side view and a front view of the screen display.

The virtual racket was displayed at the same height from the floor as the real racket and its displacement was the same as that of the real racket. The movement of the ball displayed on the screen was governed by ballistic flight and an instantaneous impact event when the virtual racket impacted the virtual ball. Just before the virtual ball hit the virtual racket a trigger signal was sent out to the mechanical brake that was attached to the rod. The trigger signal was sent out 15 ms before the ball-racket contact to overcome the mechanical and electronic delay of the brake. The brake applied a brief downward force pulse to the rod to create the feeling of a real ball hitting the racket. The duration of the force pulse (30 ms) was consistent with the impact duration observed in a real ball-racket experiment (Katsumata, Zatsiorsky, & Sternad, 2003). The brake force was not scaled to the relative velocity of the ball and

the racket but stayed the same for all impacts.

The computer program controlling the experiment would read the latest racket position from the optical encoder and racket acceleration from the accelerometer. When the racket was away from the ball, the program would update the ball positions based on the ballistic flight equation. On the 2.4 GHz computer under Windows XP this led to an update rate of around 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. The increased computational load led to a slow-down of the update rate to around 250 Hz in the 30 ms surrounding an impact. The update rate was not fixed because Windows XP is not a real-time operating system and thus timing is not deterministic. Hence, all data were time-stamped using the high-resolution timer on the Pentium CPU, with an accuracy better than 1 microsec.

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 the height of the racket, when held with the forearm horizontally, was 10 cm above its lowest position. This ensured that the base height of the racket position was the same for all participants. Each trial began with a ball appearing at the left side of the screen and rolling on a horizontal line extending to the center of the screen (Figure 3 inset). Upon reaching the center, the ball dropped from the horizontal line (0.7 m high). The task instruction was to rhythmically bounce the ball for the duration of a trial (55 s) as accurately as possible to the target line (the same line that the ball started on). The experiment consisted of a total of 80 trials, which were collected in two sessions. Each session lasted approximately one hour.

The entire experiment was divided into 4 blocks of 20 trials, one block for each value of the coefficient of restitution α: 0.5, 0.6, 0.7, and 0.8. The blocks were presented in either ascending or descending order, counterbalanced among participants. The first two and the last two trials of each block were control trials without any perturbation. In the remaining 16 experimental trials perturbations were applied at random impact times. A perturbation was created by an abrupt change of

the ball release velocity immediately after the ball-racket impact. This led to an unexpected ball amplitude without any other noticeable change in the ball trajectory. This perturbation in velocity was chosen randomly from 14 magnitudes which were added or subtracted from the current release velocity -1.0, -0.86, -0.71, -0.57, -0.43, -0.29, -0.14, 0.14, 0.29, 0.43, 0.57, 0.71, 0.86, 1.0 m/s. Negative values led to smaller ball amplitudes, positive values to larger ball amplitudes. With an average ball amplitude of 0.55 m and a corresponding average bounce period of 650 ms, the effect of these perturbations can be converted to deviations from the target height. The largest positive perturbation caused an overshoot of 0.37 m above the target, and the largest negative perturbation an undershoot of 0.27 m. The smallest perturbation of 0.14 m/s caused an overshoot of 0.047 m and -0.14 m/s an undershoot of 0.045 m. The 14 different magnitudes of perturbations ranging from -1 m/s to +1 m/s were labeled as P-7, P-6, P-5, …. to P+6, P+7.

The complete set of 14 perturbations for each α was delivered on two successive trials, with 7 perturbations within one trial in randomized order. As each trial had approximately 80 to 90 bounces, the perturbations occurred randomly on the 8^{th}, 9^{th}, or 10^{th} bounce relative to the previous perturbation. Across the 16 experimental trials, the set of 14 perturbations could be administered 8 times. With randomization of both time and magnitude of perturbation, participants were unable to anticipate the ball amplitude or the time of the perturbations.

Data Reduction and Analysis. The raw data of the racket displacement and acceleration were resampled at a fixed frequency of 500 Hz and filtered with a 4^{th}-order Savitzky-Golay filter with a window size of 0.01 s on both sides (Gander & Hrebicek, 2004). The filter order and window size were chosen empirically to remove measurement noise while not excessively smoothing the signals. The Savitzky-Golay filter is superior for smoothing data that have abrupt changes as compared to conventional filters like Butterworth filters. These abrupt changes occurred in the data when the racket exhibited a sudden drop in acceleration, caused by the brake. The ball displacement was generated by the computer so it contained no measurement noise.

Therefore, no filtering was necessary. As a verification of our filtering procedure, racket displacement was double-differentiated using a Savitzky-Golay filter and compared with the acceleration data collected by the accelerometer. The comparison showed a good match between the two types of data, supporting the validity of the data acquisition.

Dependent Measures. Figure 4 illustrates the primary dependent measures. Performance was evaluated by the ball height error, *HE*, which was defined as the signed difference between the maximum ball height and the target height. Height error was equivalent to the state variable ball velocity at the moment of release from the racket, as this velocity determined the subsequent ball amplitude in the gravitational field if the impact position was relatively constant. The racket amplitude, A_R, was calculated as half the distance from the minimum to the maximum of the racket trajectory during one cycle. The period between impacts, *T*, was calculated from the time intervals between the impacts of successive bounces. The acceleration of the racket at impact, *AC,* was determined from the accelerometer signal one sample before the time of impact.

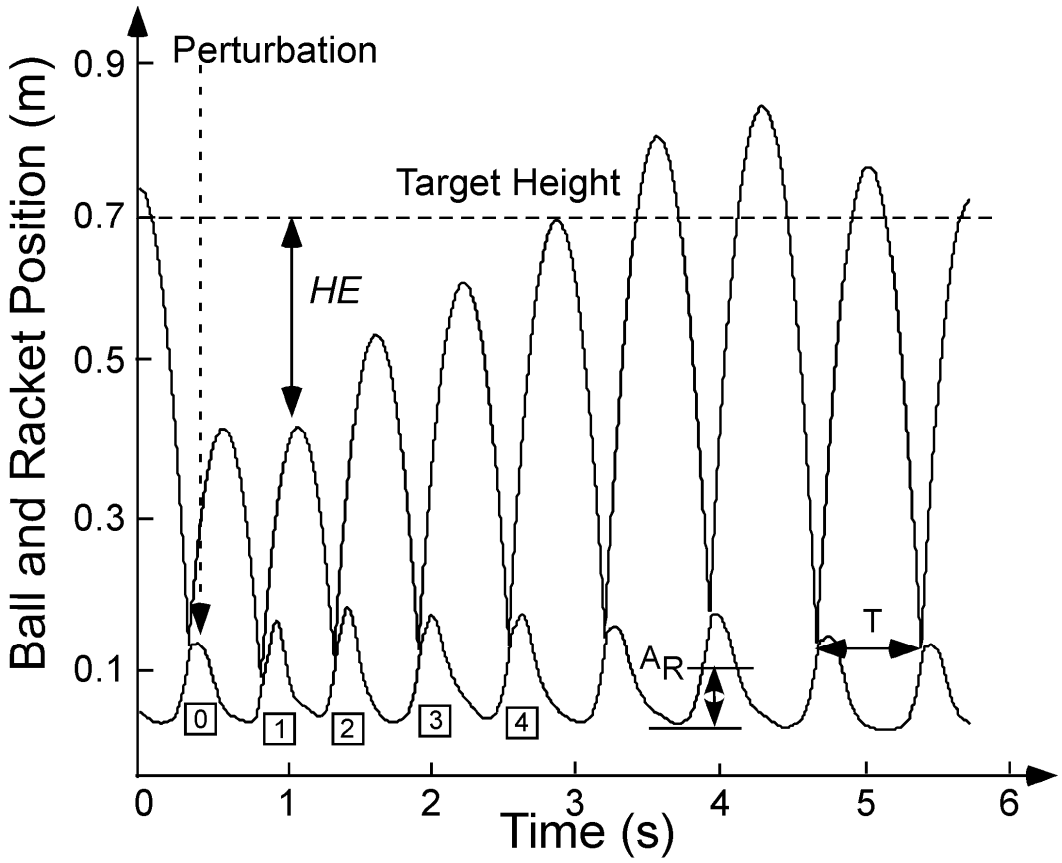

Figure 4: Segment of an exemplary trial with a large perturbation (*P-7*). The impacts are marked by numbers and the perturbed impact is marked as the 0^{th} impact.

Relaxation Time after Perturbation. The bouncing cycles were labeled in sequential order starting from the unperturbed impact directly before the perturbed impact. This last unperturbed cycle was labeled *C-1*, the perturbed cycle was labeled *C0,* and the following cycles were labeled *C1*, *C2,* and so on (Figure 4). Some of the dependent measures, for example *HE* and A_R, showed an abrupt change after a perturbation with a subsequent exponential return to a final level. However, this final level did not always coincide with the pre-perturbation levels. In order to quantify this exponential return with variable final level, least-squares fitting of an exponential function was performed using the following functional form:

$$y_k = y_0 e^{-k/\tau} + y_\infty \qquad (8)$$

where y_k was the dependent measure of interest and k was the cycle number with $k = 0$

denoting the perturbed bounce at *C0*. Fitted parameters were the amplitude of the perturbation y_0, the final level y_∞ and the relaxation time τ. We used the Levenberg-Marquardt least squares fitting algorithm (Matlab 6.5, Mathworks).

Results

I. Performance during Control Trials

Racket Acceleration. Given that the objective of the study examined perturbations away from stable performance the data first had to be examined whether actors indeed performed the task in accord with the criteria of passive stability. Hence, performance was evaluated in the control trials (four trials for each of the four α condition). The primary measure indicating performance at passive stability is the mean impact accelerations *AC* across all bounces of one trial (typically 70-80 bounces during 60s long trials). The mean values across the four control trials, determined separately for each of the four α conditions, are listed for all 7 participants in Table 1. Overall, participants showed negative *AC* values as predicted by the model and seen in previous studies, with only two exceptions: Participant 1 and 2 had small positive values for $\alpha = 0.8$ and 0.5, respectively. Excluding these two cases, these results verified that all participants indeed performed the task consistent with criteria for passive stability. A 4 (α) x 7 (participant) ANOVA was performed on these data with participant treated as a random factor. The results showed no significant differences between different α values, $F(3, 84) = 0.63$, $p = 0.607$. The main effect of participant and the interaction were significant, $F(6, 84) = 9.98$, $p < 0.0001$, and $F(18, 84) = 3.02$, $p < 0.0001$, respectively. The grand mean of these control data were used as the baseline for the design of the perturbation magnitudes and the calculations of the basin of attraction.

Table 1:

Impact accelerations (m/s^2) of individual participants in control trials. The values are means across 6 trials. Standard deviations are shown in parentheses.

	$\alpha = 0.5$	$\alpha = 0.6$	$\alpha = 0.7$	$\alpha = 0.8$	Participant Mean
Participant 1	-1.25 (3.29)	-0.07 (2.11)	-2.26 (2.10)	0.29 (2.04)	-0.82 (2.39)
Participant 2	1.00 (2.70)	-0.11 (2.68)	-1.08 (2.75)	-0.72 (2.09)	-0.23 (2.56)
Participant 3	-2.43 (2.28)	-2.22 (1.98)	-1.43 (1.50)	-1.50 (1.25)	-1.90 (1.75)
Participant 4	-4.21 (2.48)	-3.56 (1.74)	-4.37 (1.67)	-3.64 (1.09)	-3.95 (1.75)
Participant 5	-3.21 (2.39)	-3.46 (1.79)	-2.66 (2.05)	-2.76 (1.56)	-3.02 (1.95)
Participant 6	-0.59 (2.13)	-0.34 (1.94)	-1.60 (1.90)	-2.57 (2.05)	-1.28 (2.01)
Participant 7	-0.96 (2.46)	-0.68 (2.04)	-0.84 (1.13)	-2.00 (0.91)	-1.12 (1.64)
Grand average	-1.66	-1.49	-2.03	-1.84	

Height Error. Performance during control trials was also evaluated in terms of the primary performance measure, height error *HE*. Subjecting the mean values determined across each trial in the four α conditions to a 4 (α) x 7 (participant) ANOVA did not identify significant difference between α conditions. Differences between individuals were significant, $F(6, 84) = 7.90$, $p < 0.0001$. Overall, actors tend to slightly overshoot the target by an average of 0.016, 0.017, 0.017 and 0.023 m for $\alpha = 0.5$, 0.6, 0.7, and 0.8, respectively.

Racket Period and Amplitude. For a better characterization of the task performance the continuous racket trajectories during steady state were assessed by their mean period and amplitudes per trial. The same 4 (α) x 7 (participant) ANOVA performed on period yielded significant differences between different α values, $F(3, 84) = 8.91$, $p < 0.005$, and between different participants, $F(6, 84) = 20.56$, $p < 0.0001$. The interaction between participants and α was also significant, $F(18, 84) = 4.44$, $p < 0.0001$. A small but clear trend can be seen: the racket periods slightly

increased for larger α conditions (634±38, 642±44, 654±45, 679±44ms). The cause for this trend is that in the higher α conditions participants impact the ball at slightly lower positions although performance measured in *HE* was not affected. The equivalent ANOVA on amplitudes revealed a decreasing trend for higher α, $F(6, 84) = 173.22$, $p < 0.0001$, indicating that subjects moved the racket less when the ball-racket contact was bouncier: 0.067±0.007m, 0.049±0.005m, 0.036±0.003m and 0.025±0.005m for α conditions 0.5, 0.6, 0.7 and 0.8, respectively. The main effect for participant and the interaction were significant, $F(6, 84) = 3.87$, $p < 0.05$, $F(18, 64) = 6.76$, $p < 0.0001$.

II. Performance Following Perturbations

Height Error. Figure 5 displays the grand averages of *HE* over all repetitions and all seven participants as a function of cycle number directly before and following the perturbation. Due to space limitations, only 8 perturbation magnitudes are displayed. As to be expected, large effects of the perturbation are observed at the perturbed cycle *C0* and *HE* deviates from the baseline level with a magnitude that scales with the perturbation; larger perturbations lead to larger amplitudes of the ball, hence larger *HE* values. During subsequent cycles, *HE* shows an approximately exponential return back to pre-perturbation values, indicating that the experimental perturbation had a significant effect on the performance. This return is not symmetrical for positive and negative perturbations of the same magnitudes. Comparing the largest negative perturbation, *P-7*, with the largest positive perturbation, *P+7*, it is apparent that *P+7* shows a faster return, even though for *P+7* the ball amplitude deviated more from target (approximately +0.37 m vs. –0.27 m for *P+7* vs. *P-7,* respectively).

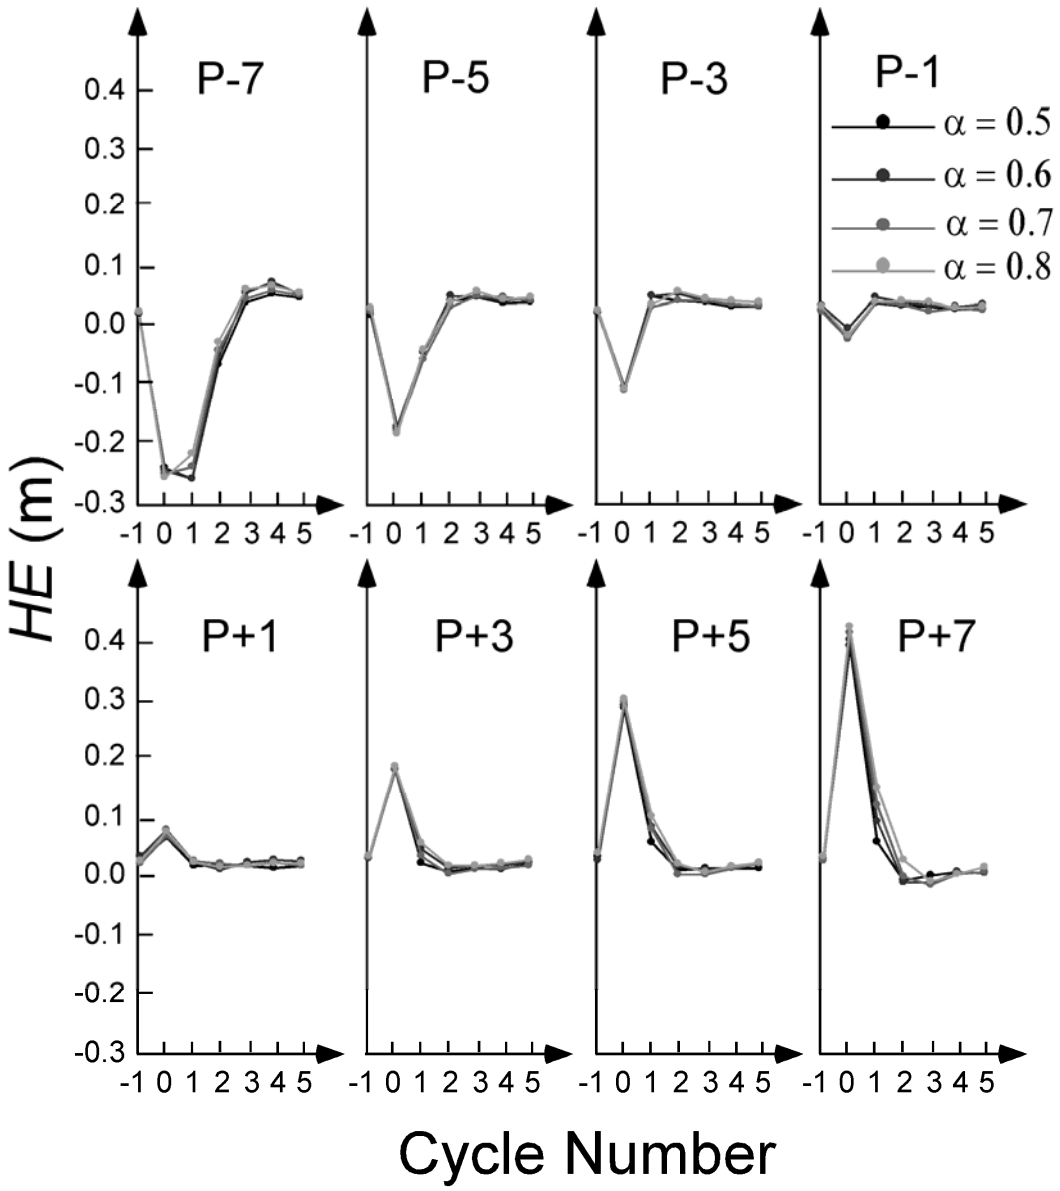

Figure 5: Grand averages over all repetitions and participants of height error *HE* plotted as a function of cycle number. The eight panels show the data for eight perturbation magnitudes. Different α conditions are shown by different grey shades.

The second observation is that for small perturbations the time of return to pre-perturbation values is as short as one cycle and only approximately three cycles for large perturbations. This indicates that participants managed to recover from perturbations faster than the model predicted. Even large perturbations that take the system outside of the basin of attraction are compensated within three cycles. There is no visible difference between the different coefficients of restitution for negative

perturbation, but some for large positive perturbations: for example, the data for *P+7* shows that higher values of α have slower relaxation times than lower α. These observations will be quantified by the following curve fitting analysis.

Relaxation Times of Height Error. In order to quantify the rate of return, an exponential function was fitted to the *HE* data (eq. 8). The relaxation time τ expresses how fast the system returns to the baseline level. The exponential curve fits are illustrated for all 14 perturbation magnitudes and for α = 0.5 in Figure 6. The different lines represent the fitted curves for *HE* over eight cycles following the perturbation (from *C0* to *C7*). The curve fits rank order with perturbation magnitude, with larger perturbations producing larger perturbation amplitudes (y_0) and longer relaxation times (τ). The R^2 values for the 14 fits for all participants ranged between 0.84 and 0.99, with an average of 0.96.

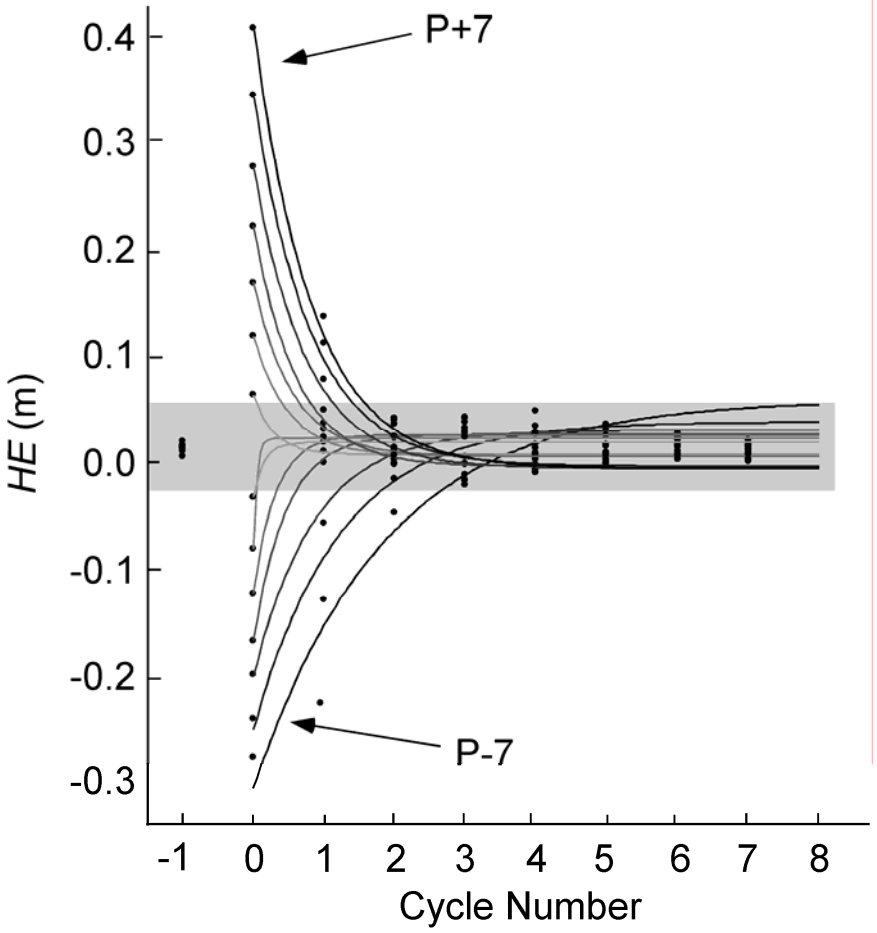

Figure 6: Exponential curve fits of the height error *HE* for 14 different perturbation magnitudes for coefficient of restitution 0.5. The dots are grand averages over repetitions and participants. The shaded area indicates the band of variability measured in the unperturbed control trials.

The relaxation time estimates obtained from these fits are plotted in Figure 7. The filled symbols denote perturbations inside the basin of attraction, the hollow symbols larger perturbations outside the basin of attraction. Note that the boundary is different for negative and positive perturbations and also for different α conditions. This figure provides a first basis to test the predictions. Prediction 1 anticipated a qualitative change in behavior from perturbations inside to outside the basin of attraction. However, the figure does not reveal such discontinuous change in the relaxation constant. Still, the τ values show a distinct pattern that can be evaluated in view of the second set of predictions. Consistent with Prediction 2a, relaxation times were higher for larger perturbations. Further, there was a noticeable asymmetry between negative and positive perturbations; relaxation times were larger for large negative perturbations than for their corresponding positive ones, consistent with Prediction 2b. It can also be seen that different α conditions did not induce differences in relaxation times across all perturbation magnitudes, with the only exception that for positive perturbations lower α conditions appeared to show shorter relaxation times. This finding supports Prediction 2c: the basin of attraction narrows for larger α, but only on the positive side.

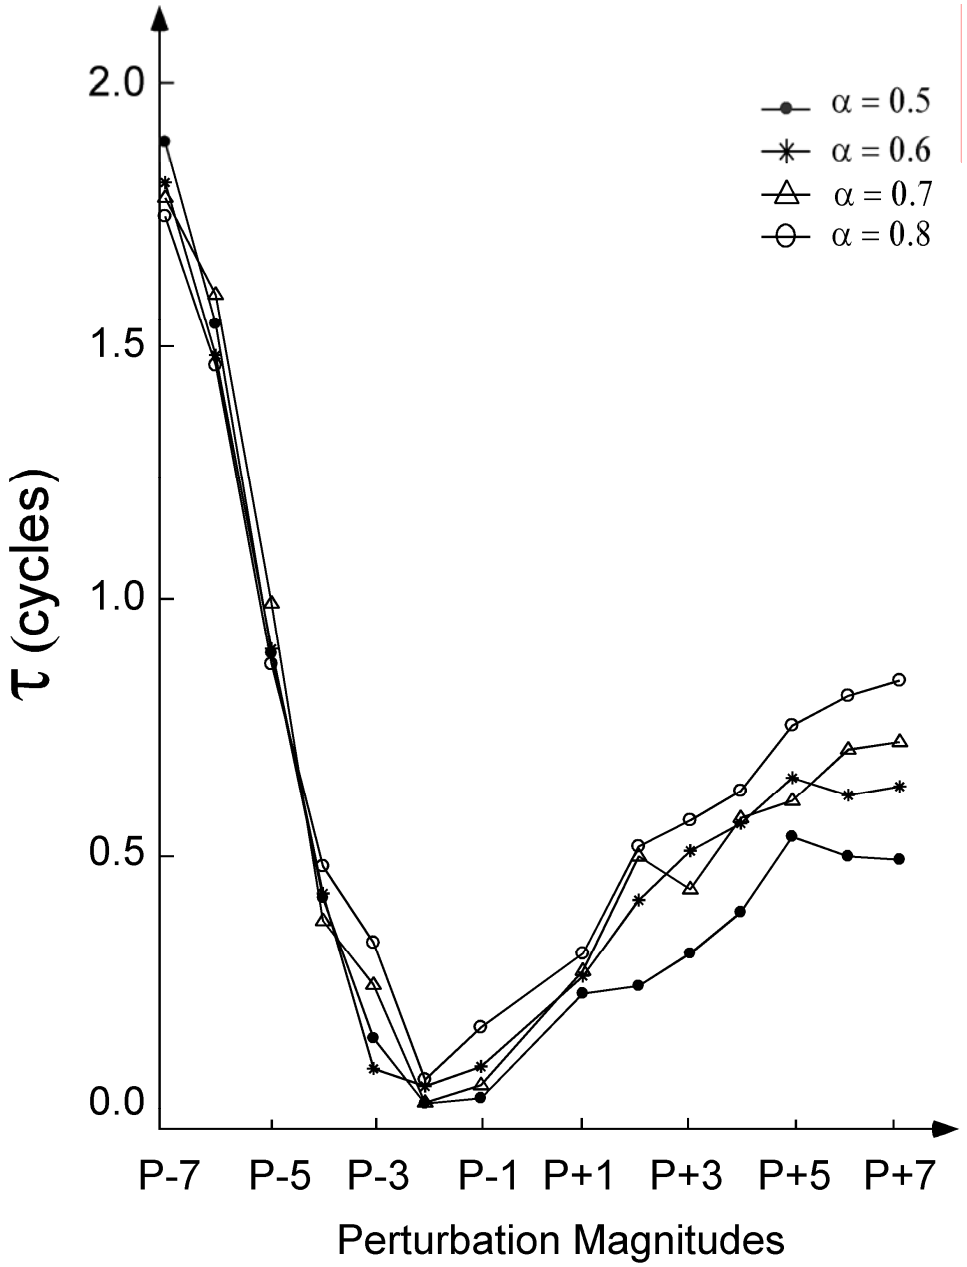

Figure 7: Relaxation times τ obtained from the exponential fits of the height error *HE*. Different α conditions are shown by different lines and symbols.

Racket Acceleration. A first assessment of the strategy that actors applied can be obtained by the analysis of the racket-ball contact characteristics. Figure 8 shows the grand averages over repetitions and participants of *AC* plotted over cycles before and after the perturbations; again, only 8 of the 14 perturbation conditions are shown. Only the three largest negative perturbations (*P-7, P-6, P-5*) introduced significant

deviations in *AC* from the baseline. For all positive perturbations and small negative perturbations *AC* did not change. This observation was supported by calculating the percentage of trials that had a significant change in the cycle directly following a perturbation. In the first cycle after a perturbation (*C1*), the percentages of trials with a change of more than one standard deviation from mean control trial performance (baseline) were 10%, 7%, 4%, 9%, 10%, 11%, 8%, 12% and 11% for perturbation magnitudes *P-2* to *P+7*, respectively. (Note that at cycle *C0* at which the perturbation was applied, 6% of trials had *AC* values different from baseline.) Only the three largest negative perturbations (*P-7, P-6, P-5*) introduced significant deviations in *AC*: 83%, 78% and 59% of trials showed changes in *C1* greater than one standard deviation. The values of *AC* jumped from around -2.0 m/s^2 at the pre-perturbation cycles to around +10 m/s^2 in *C1*. The large positive impact acceleration suggested that the ball contacted the racket earlier in the upward movement due to the undershoot of the ball amplitude caused by these three large negative perturbations.

In sum, the perturbations had minimal effects on the racket-ball contact and the racket continued to contact the ball with negative accelerations even for large perturbations outside the basin of attraction.

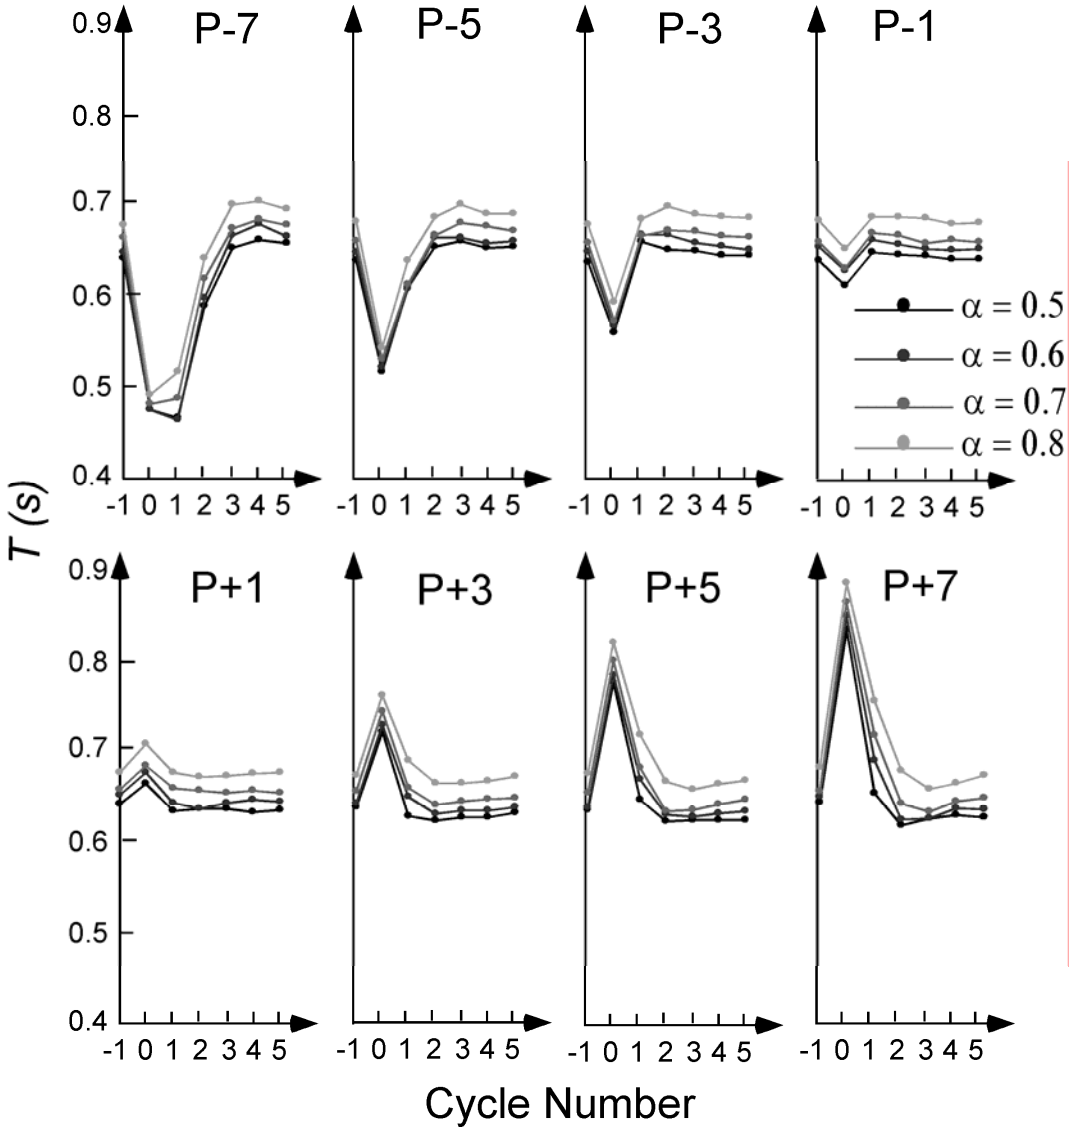

Figure 8: Impact acceleration *AC* (averages over participant and repetition) plotted as a function of cycle number. Each of the eight panels shows the data for one perturbation magnitude. Different α conditions are shown by different lines.

Thus far, the data give some seemingly inconsistent answers to the predictions. 1) The racket contact results (*AC*) indicate that participants overall continue to use passive stability to recover from perturbations. 2) However, the return behavior is quantitatively inconsistent with the predictions resulting from passively stable behavior, as the return is significantly faster. 3) There is no sign of a qualitative change when the perturbations are inside or outside the basin of attraction. 4)

However, there are several qualitative features that speak to the fact that the basin of attraction does play a role in the return behavior. These results point to the fact that actors must have made adjustments in their racket trajectories to achieve such fast recovery from the perturbations. If so, were these adjustments in the racket movements sensitive to the boundary of the basin of attraction, relying on passive stability for small perturbations? These questions were addressed in a detailed analysis of the racket kinematics.

Racket Periods and Amplitudes. Figures 9 and 10 illustrate how the racket periods T and amplitudes A_R changed in the face of perturbations. As for *AC* and *HE*, T and A_R were averaged across trials and participants of each α condition and plotted against cycle number for different perturbation magnitudes. The period results in Figure 9 show a rank ordering of α conditions: smaller α conditions lead to shorter racket periods, similar to the behavior observed in the control trials. Additionally, T also showed systematic deviations from baseline in *C0* and the magnitudes of deviations scaled with the perturbation magnitude similar to *HE*. The changing pattern of T indicates that the racket periods were adjusted according to the perturbed ball trajectory such that the racket periods were scaled with the ball amplitude following the perturbations. This coupling between the racket and the ball was in effect during the very first cycle following perturbation. There was no discernable difference between α conditions.

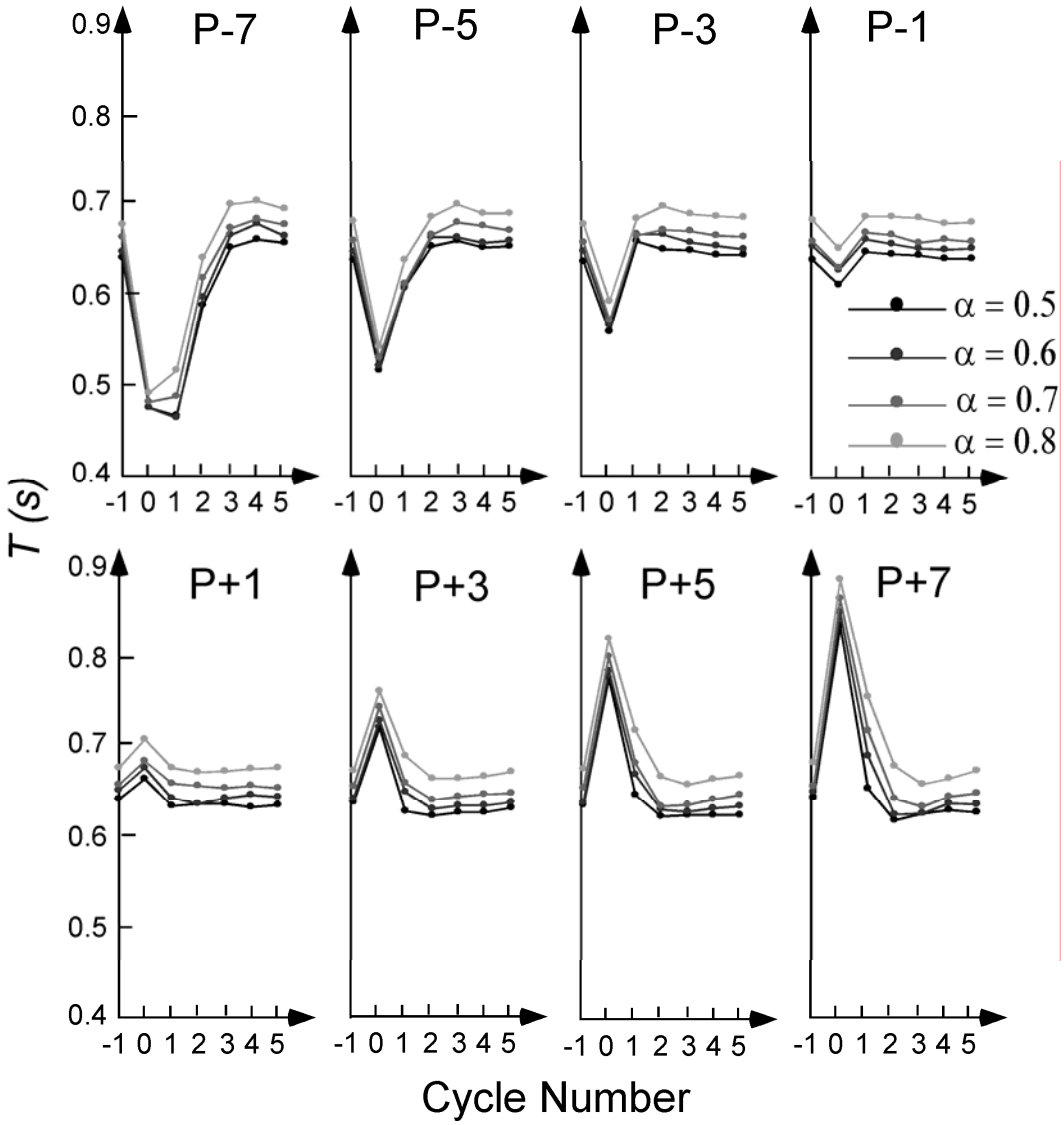

Figure 9: Period averages (over participant and repetition) between impacts *T* as a function of cycle number. Each of the eight panels shows the data for one perturbation magnitude. Different α conditions are shown by different lines.

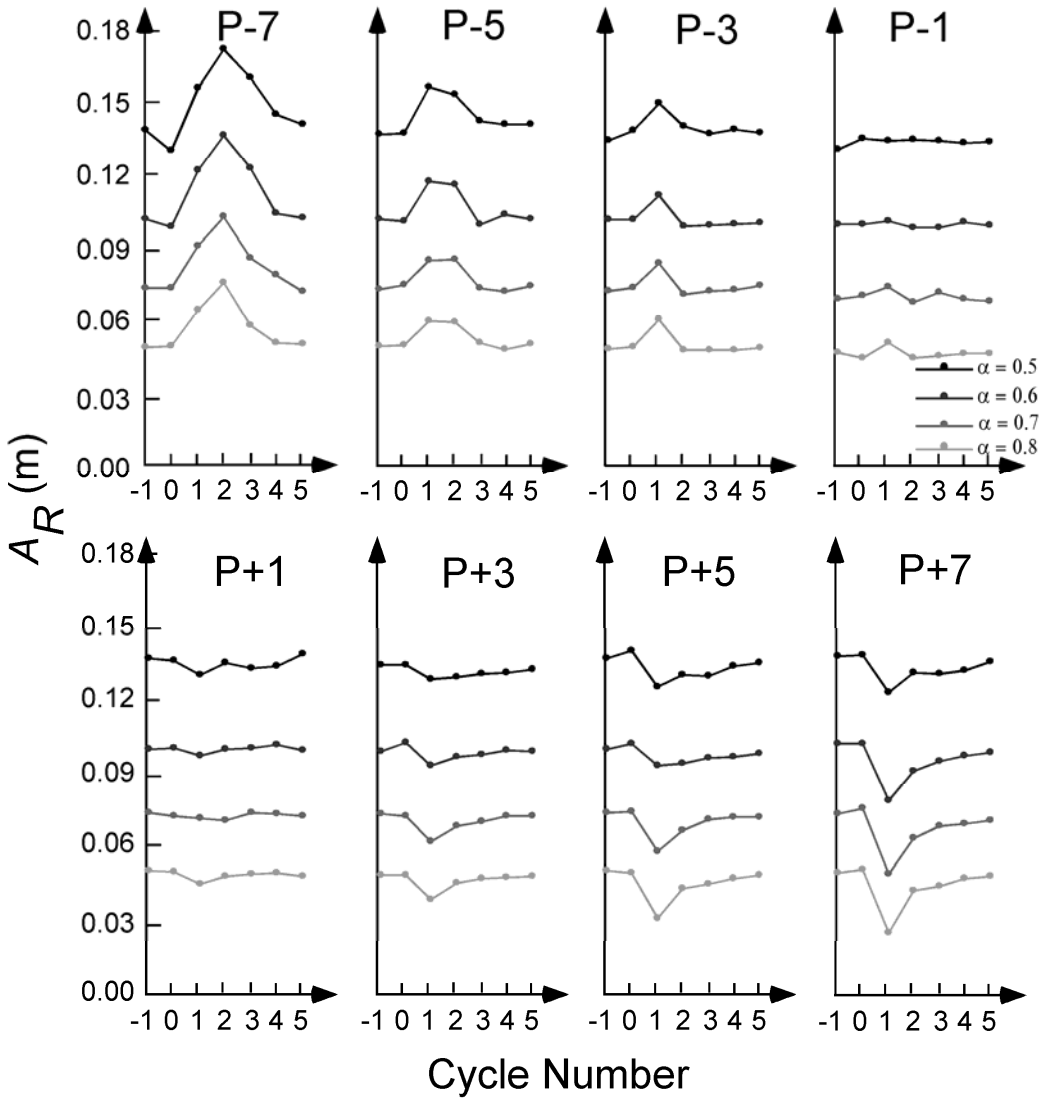

Figure 10: Racket amplitude averages (over participant and repetition) A_R plotted as a function of cycle number. Each of the eight panels shows the data for one perturbation magnitude. Different α conditions are shown by different lines.

Figure 10 summarizes the amplitudes A_R: a strong dependence on α, i.e., the more elastic the ball-racket contact was (higher α value), the smaller was A_R. (Note that no change is seen for *C0* since the perturbation was applied at *C0* and did not yet affect the racket movement.) In addition, A_R systematically changed following the perturbation, increasing or decreasing depending on the sign of the perturbation. The larger the perturbation, the larger were the changes in A_R. However, for the very small perturbations, the changes in A_R were relatively small. There was no discernable

difference between α conditions in terms of changing pattern of $A_{R.}$

This qualitative display indicates changes in the racket kinematics which are consistent with the observations of the fast return following perturbations and the lack of change in the racket-ball contacts. To assess whether these racket adjustments were sensitive to the boundary of the basin of attraction (Prediction 1) exponential fits were applied in the same fashion as for *HE*. Using eq.8 the three parameters of the exponential function were fitted: the relaxation time τ, the gain y_0, and the baseline value y_∞. The parameter values for both period and amplitude data are summarized in Figure 11. The pattern of the τ values shows a qualitatively similar pattern as the ones for *HE*, although slightly less systematic in racket amplitude. The gain parameter y_0 shows a linear change with the perturbation magnitude for both period and amplitude. The baseline parameter shows a gradual decrease for period with perturbation magnitude; the same parameter for the amplitude fits is unmodulated by perturbation magnitude. With a view to Prediction 1, no support can be found that the racket movements show qualitative changes to adapt to large perturbations outside the basin of attraction. It is also noteworthy that the gradual change also covers the very small perturbation magnitudes.

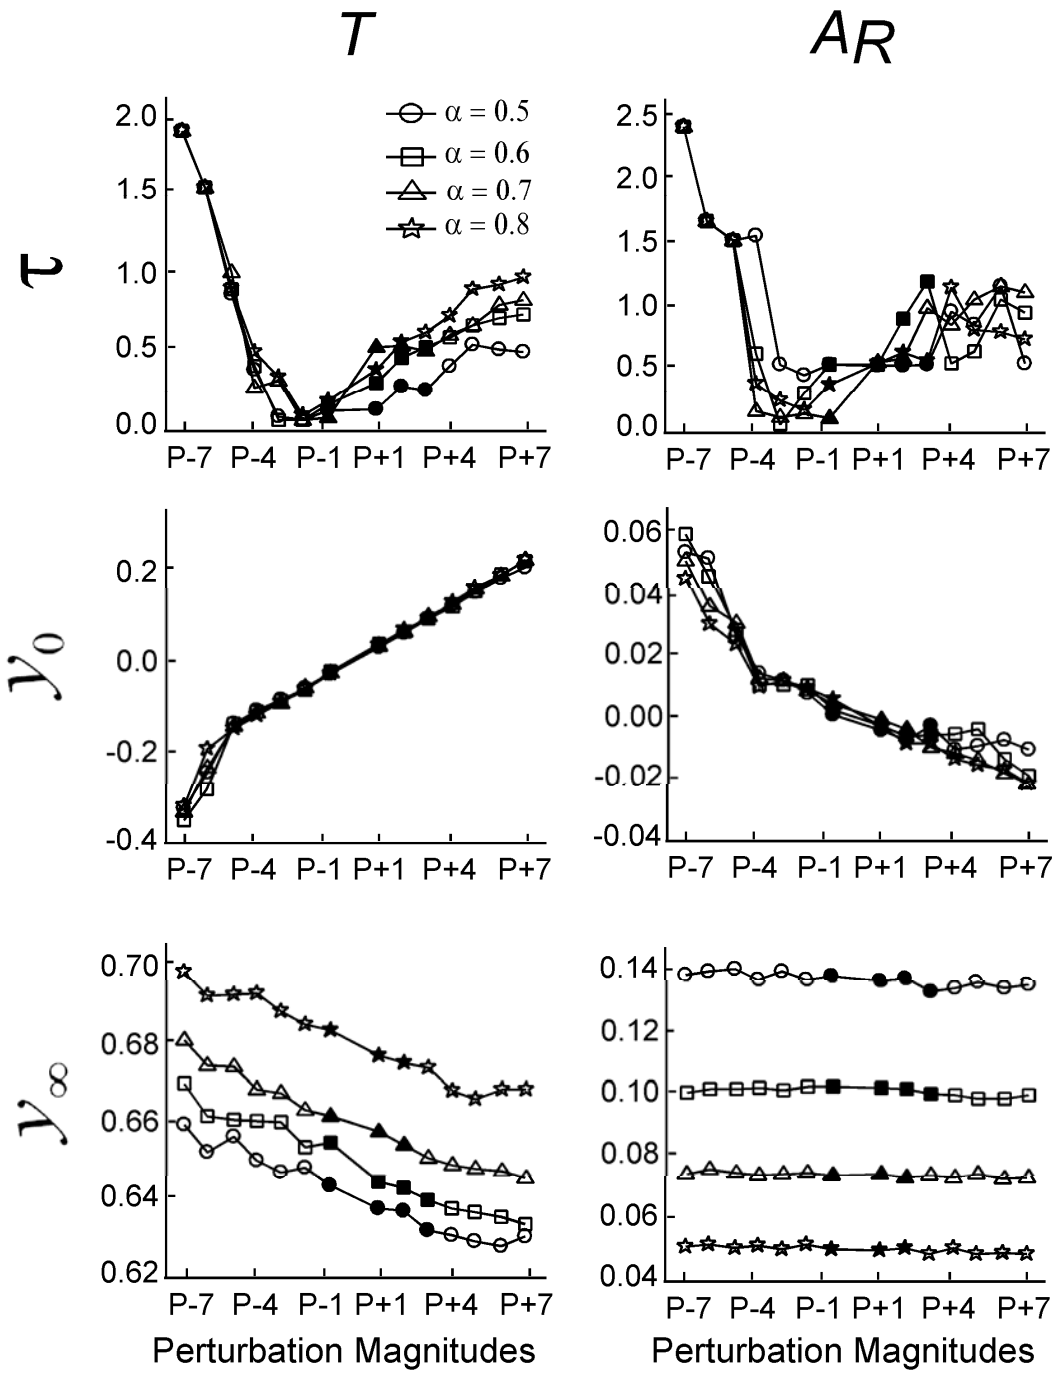

Figure 11: Parameters obtained from the exponential fits of the racket periods T (left column) and racket amplitudes A_R (right column) following the perturbation. The three parameters (the relaxation time τ, the gain y_0, and the baseline value y_∞; see eq. 8) are plotted as a function of the 14 perturbation magnitudes. Results from different α conditions are indicated by different symbols. The filled symbols indicate perturbations inside of the basin of attraction, the hollow symbols denote perturbations outside the basin of attraction.

Discussion

The present experiment examined the role of passive stability and active control in a rhythmic perceptual-motor task where participants aimed to steadily bounce a ball to a given target height. Previous studies on steady state performance of the same experimental task established that actors are sensitive to and exploit the stability properties defined by the dynamics of the racket and ball movements (Schaal, et al., 1996; Sternad, et al., 2000, 2001; Dijkstra et al., 2004). This strategy implies that small errors, such as unintended deviations of the ball from the target height, require no corrections as the ball trajectory passively relaxes back to its steady state. Large perturbations, however, may take the system outside its basin of attraction and stable period-1 solutions are lost if the errors are not actively corrected for. The present study applied perturbations to examine whether actors change their strategy to correct for such induced errors and whether their strategy is dependent on the magnitude of perturbations.

A set of 14 perturbations of different magnitudes was designed for each of the four coefficients of restitution based on the basin of attraction that was determined from the ball bouncing model. As the stability boundaries were different for different coefficients of restitution, the experiment was conducted with four coefficients of restitution. The predictions about relaxation behaviors after perturbations were based solely on the passive dynamics of the ball-racket system as the ball bouncing model does not include any active control of the racket movements. The first prediction was that actors exploit this passive stability and do not change their racket movement according to perceived error information as long as errors are small; in contrast, for larger perturbations, actors either lose stability or apply active corrections to their racket movements to regain stability. Such sensitivity to the boundary of stability is present in purely passive dynamical walking models where disturbances outside the narrow basin of attraction make the bipeds fall (Schwab and Wisse, 2001; Garcia, Chatterjee, Ruina and Coleman, 1998). A second set of predictions, also derived from

the passive model, formulated more qualitative expectations: a) With increasing perturbation magnitude, the time for returning to baseline performance increases for all coefficients of restitution α. b) The relaxation time is longer for negative perturbations than for positive perturbations for all α. c) For positive perturbations the smaller α should show faster returns than the higher α. The alternative hypothesis was that actors perceived all applied errors, even small errors, and actively adjusted their racket movements to regain steady state performance.

The previous study by de Rugy et al. (2003) already examined perturbations in the ball bouncing task and identified active modulations of the racket trajectory. However, several essential aspects were different in this previous study. First, the perturbations were applied in terms of changes in the coefficient of restitution α, a parameter of the model, and not in terms of ball release velocity, a state variable of the model. Therefore, the actual perturbation effects on the height error depended also on the ball and racket velocity. As perturbation magnitudes were not as accurately controlled as in the present study, the resulting effects were not analyzed as a function of perturbation magnitude but rather pooled over all perturbations. Further, small perturbations were excluded to ensure that effects were observable. As no analyses of the basin of attraction were available, the effect of the stability boundary on behavior could not be addressed. Also, the virtual set-up only provided a visual interface and was therefore not as realistic as the current development with the haptic contact. Despite these shortcomings, the findings clearly indicated that actors adjusted their racket movements in order to re-establish the stable pattern, i.e., they "actively tracked passive stability". The present study built on this experiment but significantly developed the theoretical framework and fine-tuned experimental approach to afford model-based quantitative and qualitative predictions about the effects of perturbations on dynamically stable behavior.

Interestingly, the results provided no support for Prediction 1: Participants' behavior showed no signs of sensitivity to the boundary of the basin of attraction. Both the performance measure height error and also the racket trajectories showed gradually more pronounced adaptations to increasing perturbation magnitudes. The

result of these adaptations was that the height error quickly decreased or increased to baseline, i.e., stable performance was quickly resurrected. The gradual change in racket kinematics also reveals that compensatory modulations in the racket trajectory existed for very small perturbations. This indicated that for all applied errors some compensatory behavior is seen, despite the dynamic stability afforded by the task. Yet, the focal dependent measure, acceleration of the racket at contact, exhibited very little changes for all but the three largest negative perturbations.

Analysis of the performance after perturbation showed qualitative agreement with the second set of predictions of the model that deemphasized the discontinuous boundary of stability. First, the return to steady state after a perturbation took longer for larger perturbation magnitudes in all α conditions. This was evidenced by the increasingly longer relaxation times of the ball height errors over successive cycles after the perturbation with larger perturbation magnitudes. Second, relaxation times were shorter for positive compared to negative perturbations of corresponding magnitudes, mirroring the asymmetry in the basin of attraction. Third, lower α conditions exhibited faster returns for positive perturbations than higher α conditions. This is consistent with the topology of the basin of attraction whose width is wider for higher α conditions but only on the positive side.

While these results were *qualitatively* consistent with the model predictions, the results did not match the predictions *quantitatively*. Overall, the relaxation behavior was considerably faster than the model predicted. Based on the assumption that the racket trajectory remains unchanged in the face of perturbations, the model predicted that the return to steady state should take two to several tens of bounces for perturbations, even when the perturbation was still inside the basin of attraction. Solutions other than period-1 or the sticky solutions as predicted for perturbations outside the basin of attraction were never observed. This shows that in all cases participants actively accelerated their returns to the pre-perturbation steady state to as fast as one to three bounces.

This fast return behavior was accomplished by active modulation of the racket trajectory in both amplitude and period. Both of these racket measures showed

systematic deviations from baseline levels in the cycles after perturbations with an exponential return to their respective baselines. Interestingly, these modulations were seen from very small to very large perturbations, showing that even very small perturbations induced compensatory adaptations in the racket movements. In contrast to these significant modulations, impact accelerations of the racket remained negative during all perturbations, except for the three largest negative perturbations (reasons discussed below). Thus, while racket trajectories were actively adjusted to compensate for perturbations, conditions for passive stability, negative acceleration at impact, were maintained or immediately reestablished. One needs to keep in mind that this is possible with small deviations from purely sinusoidal trajectories. Viewing these changes in racket movements together permits the conclusion that actors may set up conditions for passive stability to assist return to steady state by active correction.

For the three most negative perturbations (*P-7, P-6, P-5*) high positive impact accelerations were observed in the two cycles following the perturbation, coincident with relatively large changes in the racket amplitudes. Two effects might contribute to this observation. First, assuming that actors perceived the smaller ball amplitude in sufficient time, they may have made a faster upswing of the racket to propel the ball to the desired amplitude in the following bounce. If the impact happened during this accelerating phase, the impact acceleration would be positive. Support for this interpretation comes from the observation that the values of *AC* on bounce *C1* scaled with α values (see first and second panel of Figure 8). The higher the α, the less forceful an impact was required, thus the *AC* values were lower. A second and similarly likely explanation is that for large negative perturbations the ball amplitudes were significantly lower and, concomitantly, the impact-to-impact periods significantly shorter. Assuming approximately sinusoidal racket movement, an earlier impact during the upward movement of the racket will result in a more positive racket acceleration. Thus, the large positive values of *AC* following the three largest negative perturbations might be, at least partially, a direct consequence of the unexpected

undershooting of the ball. This explanation is supported by the fact that the period following these perturbation in *C1* was merely about 500 ms. Assuming the actor can only observe the perturbation when the ball reaches the maximum height, it will only leave him about 250 ms, i.e., half of the period, to prepare a more appropriate impact.

The findings of the present study support the conclusion that in the ball bouncing task human actors adopt a strategy of blending active control with exploiting the passive dynamics of the task. Active control and passive stability have been studied extensively in human walking and walking robots. Mathematical models of purely passive dynamic walkers (without any actuators) revealed that stability is attainable from few mechanical assumptions, such as ballistic motion of the swing leg and an instantaneous double support phase. These passive dynamic walkers do not have actuators similar to the neuromuscular apparatus and the only energy input was the potential energy derived from descending a gentle slope (McGeer, 1990). However, the basin of attraction of the passive walker is narrow and several constraints have to be satisfied to utilize this passive stability: the walker has to be started off correctly, the energy input has to be high enough to overcome the dead point where the potential energy of the system is the highest, and the swing foot has to reach a forward position in time. During walking, the passive dynamic walker is also not very resistant to perturbations and small unevenness of the floor can cause the walker to fall (Schwab & Wisse, 2001; Wisse & van Frankenhuizen, 2003).

Stability becomes even more critical if the walking robot can also move in the lateral direction, i.e., when the passive walking model is extended from 2D to 3D: lateral balance requires additional control or constraints. One solution was to mechanically couple the lateral motion of the leg to the hip rotation such that the side-to-side shift of the center of pressure caused by the change of the supporting leg was counteracted by a rotation of the hip (Wisse, Schwab & Linde, 2001). Although this solution brought lateral stability, humans much rather show adjustments in their foot placement between successive steps, i.e., they actively modify their behavior instead of using a simple mechanical coupling between hip and leg movements.

Additionally, very fast adjustments were observed in lateral ankle movements during every step (Hof, van Bockel, Schoppen, & Postema, 2006). These modifications are obviously based on perceived deviations from stability. In other words, an active control mechanism based on sensory information and feedback is necessary to extend the passive model to mimic biological behavior.

These considerations lead to the intriguing question about the role of perceptual information for achieving or tuning stable behavior. For walking, Bauby and Kuo (2000) suggested that humans harness passive dynamic properties of the limb in the sagittal plane, but they also include perceptually guided active control to stabilize lateral motion. In ball bouncing, little work has examined the role of perceptual information. In one comparisons of performance with visual and haptic information the results showed that both types of information provide sufficient information to maintain the bouncing action, even though haptic information ensure more stable behavior than visual information alone (Sternad et al., 2000). In a two-dimensional version of the task, where the task is extended to a ball-in-a-wedge billiard Ronsse and colleagues showed that visual information alters the behavioral strategies that humans adopt compared to performance with haptic information only (Ronsse, Levebre, & Sepulchre, 2006). Given that the dimensionality of this task is higher, it is also likely that perceptual information plays a more important role. Yet, more work is needed to examine the online modulation and the required perceptual information.

An active form of control has been implemented by Bühler, Koditschek and colleagues (Bühler & Koditschek, 1990, Bühler, Koditschek, & Kindlmann, 1994). In their studies on ball bouncing in 2D and 3D, using a planar robot arm that bounces a puck, they implemented the so-called mirror law. With this algorithm the racket movements “mirror” the down coming ball movements such that the velocity of the racket is tightly coupled to the visually perceived ball velocity. One interesting outcome of this tight perception-action coupling is that the ball-racket contact then occurs at positive impact accelerations, which, according to our model, produce unstable solutions. Nevertheless, using this mirror algorithm the robot juggler

obtained successful bouncing actions in 2D and 3D.

In summary, the present study on perturbed ball bouncing revealed that actors apply a blend of passive stabilization and active perceptual control. On the one hand, the actor utilizes the passive stability afforded by the task and the performance qualitatively bears the signature of the basin of attraction derived from a mechanical model of the task. On the other hand, there were significant modulations of the actions upon changes brought about by perturbations of all sizes. The theoretically predicted boundary of stability did not seem to induce significant qualitative changes in performance. This is consistent with the idea that actor stabilize and flexibly control their behavior in the context task, environmental, and informational constraints (e.g., Newell, 1986; Warren, 2006). To account for this blend of control, perceptually-based tuning of the passively stable dynamics should be included in the purely passive ball bouncing model. This will be left for future work.

Acknowledgements

This research was supported by grants from the National Science Foundation BCS-0450218, the National Institutes of Health R01 HD045639, and the Office of Naval Research N00014-05-1-0844.

Literature

Bauby, C.E. & Kuo, A.D. (2000). Active control of lateral balance in human walking. *Journal of Biomechanics*, *33*, 1433-1440.

Brody, H., Cross, R., & Lindsey, C. (2002). *The physics and technology of tennis.* RacquetTech Publishing.

Bühler, M., Koditschek, D.E. (1990). From stable to chaotic juggling: Theory, simulation, and experiments. Paper presented at the IEEE International Conference on Robotics and Automation. Cincinnati, OH.

Bühler, M., Koditschek, D.E., & Kindlmann, P.J. (1994). Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research, 13*, 101-118.

Coleman, M. J., & Ruina, A. (1998). An uncontrolled walking toy that cannot stand still. *Physical Review Letters, 80*, 3658-3661.

Collins, S.H., Wisse, M., & Ruina, A. (2001). A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research, 20(7)*, 607-615.

Collins, S.H., Ruina, A., Wisse, M., & Tedrake, R. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science*, *307*, 1082-1085.

de Rugy, A., Wei, K., Müller, H. & Sternad, D. (2003). Actively tracking 'passive' stability in a ball bouncing task. *Brain Research, 982,* 64-78.

Dijkstra, T. M. H., Katsumata, H., de Rugy, A., & Sternad, D. (2004). The dialogue

between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies, 11(3)*, 319-345.

Gander, W., & Hrebicek, J. (2004). *Solving problems in scientific computing using Maple and MATLAB*. 4th Ed. New York: Springer.

Garcia, M., Chatterjee, A., Ruina, A., & Coleman, M. (1998). The simplest walking model: stability, complexity, and scaling. *Journal of Biomechanical Engineering*, *120(2)*, 281-288.

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

Hof, A.L., van Bockel, R.M., Schoppen, T., & Postema, K. (2006). Control of lateral balance in walking: Experimental findings in normal subjects and above-knee amputees. *Gait & Posture*. [Epub ahead of print].

Holmes, P.J. (1982). The dynamics of repeated impacts with a sinusoidally vibrating table. *Journal of Sound and Vibration*, *84(2)*, 173–189.

Jin, H.L., & Zacksenhouse, M. (2003). Oscillatory neural networks for robotic yo-yo control. *IEEE Transactions in Neural Networks, 14(2)*, 317–325.

McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research 9(2*), 62–82.

Mena, D., Mansour, J. M. & Simon, S. R. (1981). Analysis and synthesis of human swing leg motion during gait and its clinical applications. *Journal of Biomechanics, 14*, 823–832.

Mochon, S. & McMahon, T.A. (1980). Ballistic walking: An improved model. *Mathematical Biosciences*, *52(3-4),* 241-260.

Newell, K.M. (1986). Constraints on the development of coordination. In M.G. Wade and H.T.A. Whiting (Eds.), *Motor development in children: Aspects of coordination and control* (pp. 341–360). Dordrecht, The Netherlands: Nijhoff.

Ronsse, R., Lefevre, P., & Sepulchre, R. (2006). Sensorless stabilization of bounce juggling. *IEEE Transactions of Robotics, 22(1)*, 147-159.

Schaal, S., Sternad, D., & Atkeson, C.G. (1996). One-handed juggling: Dynamical approaches to a rhythmic task. *Journal of Motor Behavior*, *28(2)*, 165-183.

Schwab, A. & Wisse, M. (2001). Basin of attraction of the simplest walking model. *Proceedings of DETC'01, ASME 2001*, Pittsburgh, Pennsylvania, September 9-12.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2000). Dynamics of a bouncing ball in human performance. *Physical Review E, 63*, 011902-1 –011902-8.

Sternad, D., Duarte, M., Katsumata, H. & Schaal, S. (2001). Bouncing a ball: Tuning into dynamical stability. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 1163-1184.

Tedrake, R., Zhang, T.W., Fong, M., Seung, H.S. (2004). Actuating a simple 3D passive dynamic walker. *IEEE International Conference on Robotics and Automation*, *5*, 4656-4661.

Tufillaro, N.B., Abbott T. & Reilly, J. (1992). *An experimental approach to nonlinear dynamics and chaos*. Redwood City, CA: Addison-Wesley.

Warren, W.H. (2006). The dynamics of perception and action. *Psychological Review*, *113(2),* 358-389.

Wisse M., Schwab, A.L., & van Linda, R.Q. (2001). A 3D passive dynamic biped with yaw and roll compensation. *Robotica*, *19*, 275-284.

Wisse, M. & van Frankenhuyzen, J. (2003). Design and construction of MIKE: a 2D autonomous biped based on passive dynamic walking. *Proceedings of the 2nd International Symposium on Adaptive Motion of Animals and Machines*, Kyoto, March 4-8.

CHAPTER 4

Passive Stability and Active Control During Steady-state Performance

Chapter 4 contains an original manuscript prepared for submission to a journal.

"Kunlin Wei[1], Tjeerd M.H. Dijkstra[2], Dagmar Sternad[1]"

[1]Department of Kinesiology and Integrative Biosciences, Pennsylvania State University

[2]University Medical Center, Leiden, The Netherlands

Abstract

We investigated movement control in a task where subjects rhythmically bounce a ball with a hand-held racket to a prescribed amplitude. In previous studies on the same task, stability analysis of an open-loop mechanical model revealed that passively stable solutions exist if the racket is decelerating to hit the ball. Several studies reported that experienced subjects indeed bounce the ball with negative racket acceleration, consistent with the notion of passive stability. Perturbation studies also identified active adjustment of the racket movement with respect to the perturbed ball movement. However, whether active control exists at steady-state performance has been left unchecked, due to the difficulty to quantify the active adjustment when the racket movement is stereotypical with little fluctuations. The present study extended the open-loop mechanical model with stochastic components allowing us to derive predictions for the correlation functions of the state variables, ball velocity after impact and time between impacts. Predictions from this model are that fluctuations take longer to die out with increasing $\tilde{\alpha}$. In particular, the lag-1 correlation increases with α.

A virtual reality set-up was used where participants held a real racket and manipulated a virtual ball. Eight subjects performed in this virtual reality set-up for values of the simulated coefficient of restitution (α) of the racket ranging from 0.3 (sticky racket) to 0.9 (bouncy racket) in steps of 0.1. Results indicated that correlation functions of state variables matched well with the predictions for small α conditions. In trials with large α, performance decreased as measured by the lag-1 correlation but was more stable than the model predicted. Paradoxically, variability of performance decreased with decreasing stability. These observations can be explained by increasing compensatory variability in execution, a signature of control. Hence, actors rely on passive stability when the stability of the system is high and employ more active control when stability is reduced.

Introduction

It is widely agreed upon that stability is an indispensable goal for many motor tasks. Despite the influence of external perturbations and intrinsic noise in the action system itself, static postures must be maintained and movements need to be performed consistently and reliably with little variability. The human perceptual-motor system has a variety of solutions to meet the challenge of stability, ranging from stiffness of muscular and tendonous tissues to reflexes of different latencies and to feedback-based voluntary intervention. However, besides these purely internal contributions to stability, there is another type of stability originating from the interaction between the actor and the environment. Since humans interact with objects that pose physical constraints, an analysis of the task performance, including the evaluation of stability, must take the environment, the action system and their interaction into account.

Some tasks have inherent stability that requires little or no additional control. Small perturbations die out without active error correction if the task is performed to exploit this "passive" stability. One example is the passive dynamic walking model from robotics. McGeer (1990) first demonstrated that a passive mechanical biped could walk along a gentle slope without actuation and control. The small fluctuations in the stepping motion could be dissipated, solely relying on the stability properties of this "life-less" system. The principles of passive dynamic walking have been further developed and implemented for robots walking on the level ground (for a review see Collins et al., 2005). The important implication from these studies for human movement control is that stability analysis of movements should not be limited to the neuromuscular substrate of the actor alone. A comprehensive analysis of the whole system, including the physics of both the actor and the physical environment, must be conducted to understand the task dynamics that constrains the available coordination strategies.

The idea of passive stability has been investigated in a series of studies on the task of rhythmically bouncing a ball on a racket (Schaal et al., 1996; Sternad et al.,

2000, 2001; de Rugy et al., 2003; Dijkstra et al., 2004). The task requires the actor to use a hand-held racket to bounce a ball vertically into the air to a consistent height. Similar to the passive dynamic walking studies, the research started with a model of the mechanics of the task ignoring neuromuscular aspects or control mechanisms. Stability analysis of this model revealed that the task affords a passively stable solution if the ball is hit by a decelerating racket. Small fluctuations or perturbations do not have to be explicitly corrected and stability of the system can be achieved with open-loop control. Empirical studies confirmed that experienced actors indeed adopted this strategy and organized the racket movement to hit the ball with negative racket acceleration (Dijkstra et al., 2004; Sternad et al., 2001).

Such reliance on passive stability does not necessarily exclude active control of racket movement using perceptual information about the ball movement. A learning study revealed that novice actors started the bouncing task with positive impact acceleration and gradually tune to the passive regime with negative impact acceleration after approximately 30 minutes of practice (Dijkstra et al., 2004). This finding suggests that utilizing negative impact acceleration is not an intuitive solution for the actor and exploitation of passive stability is learnt through perception-guided practice. It was also found that actors adjusted their racket movements according to target errors induced by experimentally controlled perturbations (de Rugy et al., 2003; Wei, Dijkstra, & Sternad. 2007, submitted). At the same time, the racket impact acceleration in the subsequent bounces following a perturbation remained unchanged in the negative range despite changes in other racket kinematical variables. As a result, the error introduced by the perturbation was corrected much faster than the passive model would predict. It was suggested that actors shortened the relaxation process by using perceptually guided active control. However, actors simultaneously tune their performance to be dynamically stable to facilitate the correction of errors. Taken together, these results suggest that a mixture strategy, exploitation of passive stability and active control, was adopted by the actor when faced with perturbations.

The question remains open whether the actor, when performing at steady state, solely relies on the passive strategy identified in previous studies. Negative impact

accelerations were found during steady state performance (Schaal et al., 1996; Sternad et al., 2001; Dijkstra et al., 2004) but the possibility of additional active components cannot be eliminated. It is possible that examination of racket kinematic variables has been too coarse-grained such that active control of racket movement with a more subtle presence has not been revealed. One route to evaluate performance in a finer-grained fashion is to examine the correlation structure of the state variables in the system. To make fine-grained predictions about the correlation structure, the current deterministic ball bouncing model must be extended with stochastic components. The present study adopted the stochastic model for the ball bouncing task proposed by Dijkstra and his colleagues (Dijkstra et al., 2004, see *Model* session).

Extending a deterministic model with a stochastic component is one approach to further test a dynamic model. Fluctuations in movement execution are always present, due to intrinsic noise in the biological system. For example, the HKB model (Haken et al., 1985) was developed as a deterministic model to account for differential stability and the observed transition between the in-phase and anti-phase coordination modes. To account for the enhanced fluctuations preceding a phase transition, Schöner, Haken, and Kelso (1986) extended the HKB model by adding stochastic terms to the state dynamics. This stochastic model not only quantitatively accounted for fluctuations during the transition, it also made new predictions. For example, the so-called critical slowing down phenomenon before transition was one of the predicted features from the stochastic model. These predictions were confirmed in later experiments (Scholz, Kelso & Schöner, 1987).

The current study followed the same strategy and added stochastic components to the original deterministic ball bouncing model representing the fluctuations during steady-state performance. Assuming that they are small, these fluctuations propagate across bounces with a specific relaxation behavior that can be captured by the covariance structure of the state variables and quantitative comparisons between the model and the data can be performed. Does the actor solely rely on passive stability for correcting small fluctuations in performance?

Model

The model is based on the same bouncing map previously derived in (Tuffilaro, Abbott, & Reilly, 1992; Sternad et al., 2001; Dijkstra et al., 2004). Since we study stationary performance, we compare theory and experiment using the correlation functions of the state variables as yardstick. In order for the model to make predictions about the correlation functions, we briefly review the deterministic ball bouncing map and its stochastic extension (Dijkstra et al., 2004). The map is based on three assumptions: (1) between bounces the ball follows ballistic flight under influence of gravity (2) the impact is instantaneous with the coefficient of restitution capturing the energy loss at impact. (3) The racket movement is assumed to be a pure sine wave with fixed amplitude and frequency. This assumption is not strictly obeyed in experimental data. Human actors move the racket in an approximately sinusoidal trajectory but with a steeper slope just around the impact position. Unfortunately no simple expression is available to faithfully represent the actual waveform. Thus, for mathematical simplicity we use an equivalent sinusoid that is identical to the actual waveform around the time of the impact. The equivalent frequency of this sinusoid is calculated from the period between bounces. The equivalent amplitude of this sinusoid is calculated from the stationary phase of impact θ as:

$$a_r = \left(\pi \frac{1-\alpha}{1+\alpha} \right) \frac{g}{\omega_r^2 \cos\theta} \qquad (1)$$

This equivalent amplitude is about twice as large as the observed amplitude and is such that the slope of the equivalent sinusoid closely matches the observed racket wave form. Based on these assumptions, the deterministic ball bouncing map can be derived as:

$$v_{k+1} = (1+\alpha) a_r \omega_r \cos\omega_r t_{k+1} - \alpha v_k + g\alpha(t_{k+1} - t_k) \qquad (2)$$

$$0 = a_r(\sin\omega_r t_k - \sin\omega_r t_{k+1}) + v_k(t_{k+1} - t_k) - (g/2)(t_{k+1} - t_k)^2 \qquad (3)$$

This is an implicit map with state variables v_k, the ball velocity just after impact, and t_k, denoting the racket time of impact. Parameters are g, the acceleration of gravity, α, the coefficient of restitution and ω_r and a_r the angular frequency and amplitude of racket motion. The map has a period-1 attractor that is stable when the impact acceleration AC fulfills the following constraints:

$$-2\frac{1+\alpha^2}{(1+\alpha)^2} < AC < 0 \qquad (4)$$

The stochastic ball bouncing map is a straightforward extension with a small caveat for the implicit time map (eq 3): we introduce the intermediate state variable τ_{k+1} representing the noiseless time of the next impact. Using this variable, the stochastic ball bouncing map is:

$$v_{k+1} = (1+\alpha)a_r\omega_r \cos(\omega_r t_{k+1}) - \alpha v_k + g\alpha(\tau_{k+1} - t_k) + q_v\xi_k \qquad (5)$$

$$0 = a_r(\sin(\omega_k t_k) - \sin(\omega_k \tau_{k+1})) + v_k(\tau_{k+1} - t_k) - (g/2)(\tau_{k+1} - t_k)^2 \qquad (6)$$

$$t_{k+1} = \tau_{k+1} + q_t\xi_k \qquad (7)$$

where ξ_k denotes a Gaussian white noise process with zero mean and unit variance. The model has two additional parameters, the noise strengths of the velocity and time dynamics, which are denoted by q_v and q_t, respectively. Potentially, the noise in the velocity and time dynamics could be dependent leading to a third parameter capturing the correlation. As we show later in the results section, the experimentally observed correlation is essentially zero and thus we do not bother to introduce a parameter for this correlation.

The correlation structure of the state variables can be calculated by linearizing the stochastic map around the period-1 attractor. To this end, we denote the state of the system by y_k and the noise by ε_k:

$$y_k = \begin{pmatrix} v_k - \hat{v} \\ t_k - \hat{t} \end{pmatrix} \qquad \varepsilon_k = \begin{pmatrix} q_v\xi_k \\ q_t\xi_k \end{pmatrix}$$

with linearized stochastic dynamics:

$$y_{k+1} = Jy_k + \varepsilon_k \qquad (8)$$

where J denotes the Jacobian of the ball bouncing map. The derivation autocovariance function of this vector autoregressive system of order 1 (VAR(1)) is standard (Kendall & Ord 1990, Shumway &Stoffer 2000). $R_y(0)$ at lag-0 of the two state variables can be calculated from the Lyapunov equation:

$$R_y(0) = JR_y(0)J^T + Q \qquad (9)$$

with Q denoting the autocovariance matrix of the noise ε_k. For non-zero lag l the autocovariance function follows from:

$$R_y(l) = J^l R_y(0) \qquad (10)$$

The Jacobian depends on all four parameters of the ball bouncing map. However, in their paper Dijkstra et al (2004) show that in many cases the absolute values of the eigenvalues of the Jacobian depends only on α, the coefficient of restitution. Specifically, the absolute values of the eigenvalues of the Jacobian equal α. Thus, the autocovariance function in eq 10 can be approximated as:

$$R_y(l) = |\alpha|^l R_y(0) \qquad (11)$$

There is one more caveat in using the theory above for comparing theory and experiment: state variable t_k, the time of impact, does not have a stationary state at zero. Specifically, the stationary state is determined by $t_{k+1} = t_k + T$ with $T = 2\pi/\omega_r$. This creates a problem when calculating the stationary state from data as T also has to be estimated from the data and T fluctuates within a trial when actors shift their impact height. In order to avoid this problem, we used the time between impacts $\delta_k = t_k - t_{k-1}$ as state variable. The autocovariance function based on this new state variable ($R_z(l)$) is related to the old one ($R_y(l)$) as follows:

$$R_z(l) = \begin{pmatrix} R_y^{1,1}(l) & R_y^{1,2}(l) - R_y^{1,2}(l-1) \\ R_y^{2,1}(l-1) - R_y^{2,1}(l) & R_y^{2,2}(l) - 2R_y^{2,2}(l-1) + R_y^{2,2}(l-2) \end{pmatrix} \qquad (12)$$

with $R^{1,1}$ denoting matrix element 1, 1. With this background, the stochastic model makes the following predictions:

Prediction 1: Consistent with previous experiments, negative impact acceleration of the racket will be observed.

Prediction 2: Stability, as quantified by the autocovariance function, decreases with increasing coefficient of restitution (α). Particularly, from eq 11 it follows that the autocorrelation function at lag-1 equals α.

Prediction 3: With decreasing passive stability offered by the task dynamics, actors utilize more active control to stabilize their performance if a comparable performance is maintained at higher α conditions. The less dynamic stability the task offers, the more active control should be used. The stochastic model of ball bouncing assumes that oscillatory movement of the racket is purely passive without any modulation according to ball movement. More active control in higher α conditions will make the model deviate more from this assumption about passive dynamics. Therefore there should be an increasing mismatch between the predictions from the passive stochastic model and the real performance from the data with increasingly larger α.

To quantify the amount of active control during steady-state performance, the present paper will calculate covariation between execution variables by using the permutation method introduced by Müller and Sternad (2003, 2004). It has long been recognized that execution variables such as joint angles can compensate for each other to lead to the same movement result such as an end point location. This covariation between execution variables has been proposed as a signature of control for movement coordination (Bernstein, 1967; Arutyunyan et al., 1969; Scholz and Schöner, 1999; Müller & Sternad, 2003, 2004). The present paper will define relevant execution variables for the ball bouncing task and evaluate the covariation in different α conditions.

Prediction 4: Assuming that higher covariation means better or tighter control of execution variables to achieve the performance, then more active control is expected for higher α conditions.

Method

Participants

Eight volunteers, with ages ranging from 21 to 47, participated in this experiment. All participants were 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 Regulatory Committee of the Pennsylvania State University.

Experimental Apparatus

In the virtual reality setup, participants manipulated a real table tennis racket in order to bounce a virtual ball that was projected on a screen in front of them (Figure 1). Participants stood about 0.5 m behind a back-projection screen with width 2.5 m and height of 1.8 m. A PC (2.4 GHz Pentium CPU, Windows XP) controlled the experiment and generated the visual stimuli with an ATI Radeon 9700 graphics card. It also acquired the data using a 16 bit DT322 A/D card (DataTranslation). The images were projected using a Toshiba TLP 680 TFT-LCD projector and consisted of 1024 by 768 pixels with a 60 Hz refresh rate. Accelerations were measured using a solid state piezoresistive accelerometer mounted on top of the racket (T45-10, Coulbourne). The mechanical brake on the rod attached to the racket was controlled by a solenoid (Magnet-Schultz type R 16x16 DC pull, subtype S-07447). A light rigid rod with three hinge joints was attached to the racket surface and ran through a wheel whose rotation was registered by an optical encoder with an accuracy of one pulse for 0.27 mm of racket movement. The pulses from the optical encoder were counted by an onboard counter (DT322). The racket could move and tilt with minimal friction in three dimensions but only the vertical displacement was measured. Images of racket and ball position were shown on-line using custom-made software. The delay between real and virtual racket movement was measured in a separate experiment and found to be 18 ms on average with a standard deviation of 4 ms.

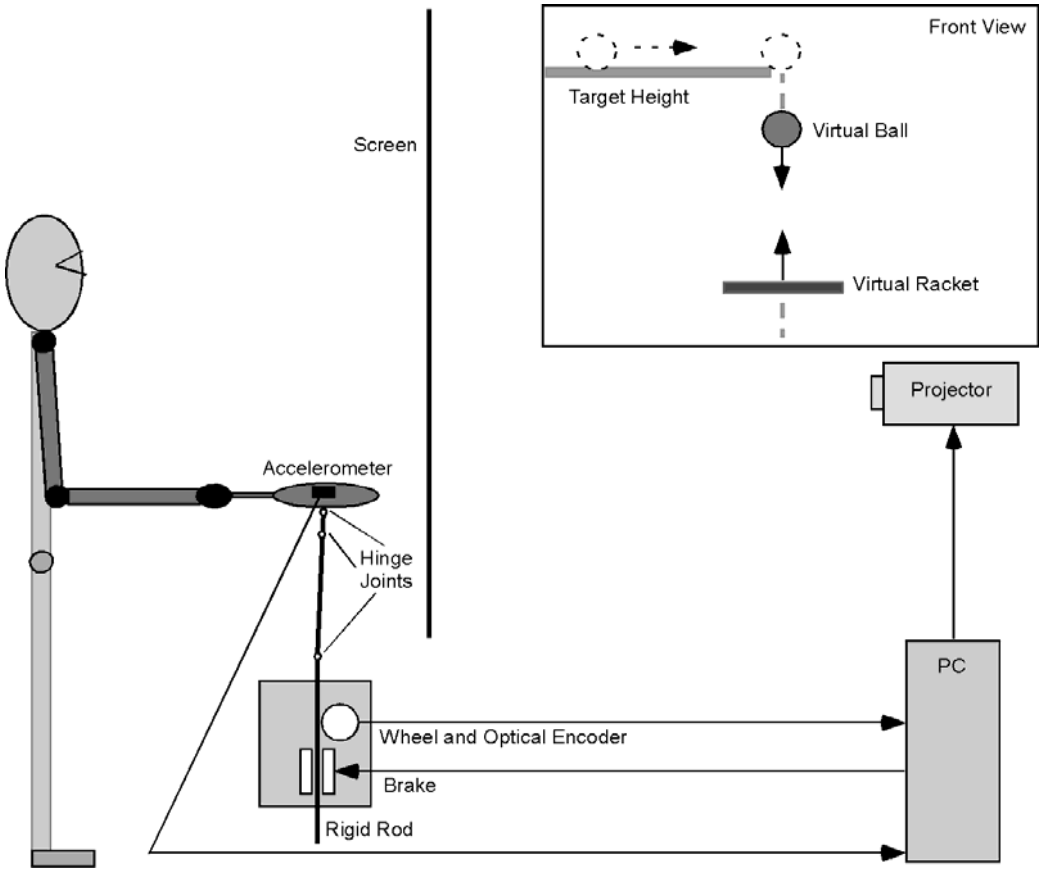

Figure 1: Experimental setup for the ball bouncing task in virtual reality.

The virtual racket was displayed at the same height as the real racket and its displacement was the same as that of the real racket. The ball displayed on the screen was a virtual ball and its movement was governed by ballistic flight and an instantaneous impact event when the virtual racket impacted the virtual ball. Just before the virtual ball hit the virtual racket a trigger signal was sent out to the mechanical brake that was attached to the rod. The trigger signal was sent out 15 ms before the ball-racket contact to overcome the mechanical and electronic delay of the brake. The brake applied a brief force pulse to the rod to create the feeling of a real ball hitting the racket. Thus, the impact caused by the brake and the impact observed on the screen coincided. The duration of the force pulse (30 ms) was consistent with the impact duration observed in a real ball-racket experiment (Katsumata, Zatsiorsky, & Sternad, 2003). The impact force applied was not modulated according to the velocities of the ball and the racket as it stayed the same for all impacts.

To summarize the experimental apparatus, the computer program controlling the

experiment would read the newest racket position from the optical encoder and racket acceleration from the accelerometer. When the racket was far from the ball, the program would update ball and racket positions based on the ballistic flight equation. On the 2.4 GHz computer under Windows XP this lead to an update rate of around 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. The increased computational load leads to a slow-down of the update rate to around 250 Hz in the 30 ms surrounding an impact. The update rate was not fixed because Windows XP is not a real-time operating system and thus timing is not deterministic. To solve this problem all data were time-stamped using the high-resolution timer on the Pentium CPU, with accuracy better than 1 microsec.

Procedure and Experimental Conditions

Each trial began with a ball appearing at the left side of the screen and rolling on a horizontal line extending to the center of the screen. Upon reaching the center, the ball dropped from the horizontal line (0.8 meter high) to the racket. The task is to bounce the ball periodically to the target line (the line the ball started on).

The experiment was conducted in 14 blocks with 7 different α values (0.3, 0.4, 0.5, 0.6, 0.7, 0.8 and 0.9). Each α condition was presented in two blocks with 3 trials in each block. For 4 participants, the blocks were presented first in ascending and then descending order of α values, i.e., from 0.3 to 0.9 then from 0.9 to 0.3. For the other 4 participants, the blocks were presented first in descending order and then in ascending order. Each trial lasted 40 seconds.

Data Reduction and Analysis

The first 5 seconds of data were eliminated from the analysis because subjects stabilized their performance in the initial part of the trial. The raw data of the racket displacement and acceleration were filtered by a 4^{th} order Savitzky-Golay filter (Gander & Hrebicek, 1993) with a window size of 0.01 s on both sides. The filter order and window size were chosen empirically to remove the measurement noise while not excessively smoothing the signals. The Savitzky-Golay filter is superior for smoothing data that have abrupt changes as compared to conventional filters like

Butterworth filters. These abrupt changes occurred in the data when the racket exhibited a sudden drop in acceleration at impact, caused by the brake. The ball displacement data was simulated by the computer so it contained no measurement noise. Therefore, no filtering was performed.

As a verification of our setup, the racket displacement was double-differentiated using a Savitzky-Golay filter and compared with the acceleration data collected by the accelerometer. The comparison showed a good match between the two types of data, supporting the validity of the customized data acquisition program.

Dependent Measures

Figure 2 illustrates the definitions of the primary dependent measures. Performance is evaluated by the ball height error, *HE*, defined as the signed difference between the maximum ball height and the target height. The amplitude of the racket movement (A_R) was calculated as half the distance from the minimum to the maximum of the racket trajectory during one cycle. The period of the ball-racket actions (T) was calculated from the time intervals between impact moments of successive bounces. Impact acceleration (AC) was calculated from the accelerometer signal at the time of impact.

For presentation, we used the grand mean of the dependent measures averaged over occurrences within trials, across trials and subjects. As measure of variability we used the standard deviation calculated from pooled occurrences within trials and across trials but averaged over subjects.

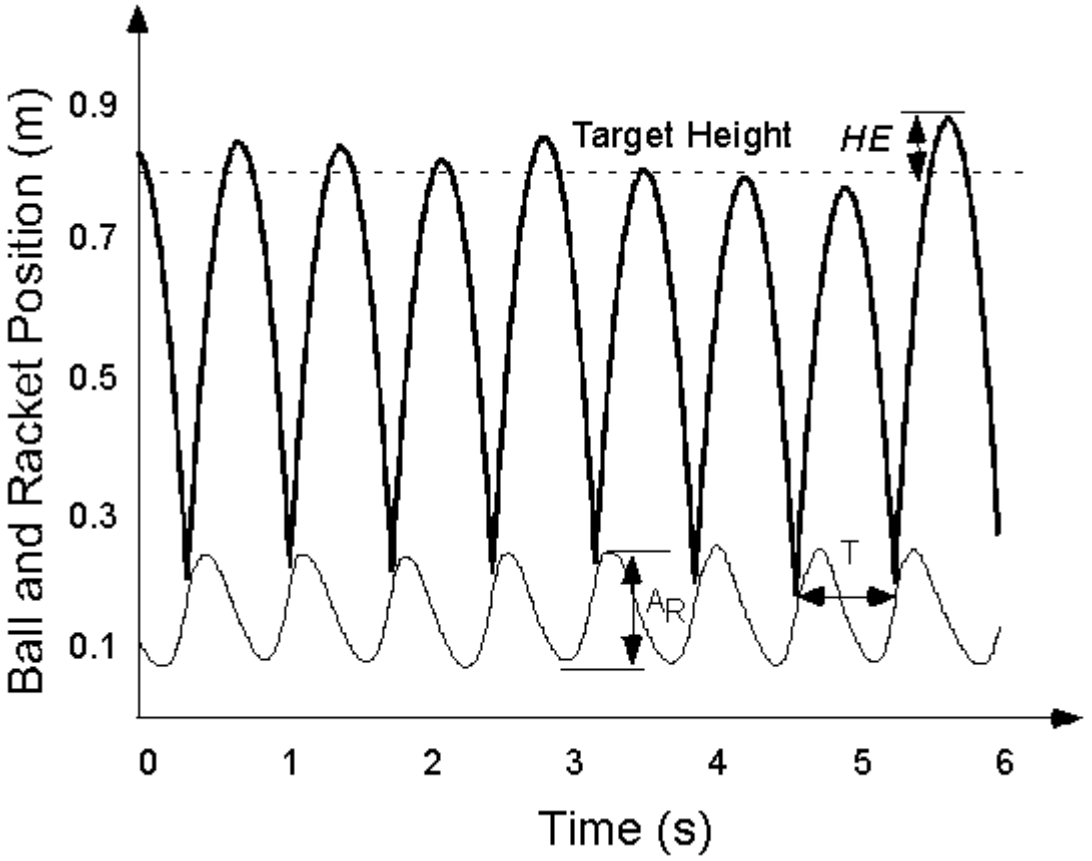

Figure 2: The racket and the ball position from a part of an exemplary trial with dependent measures *HE*, A_R and *T* illustrated.

Correlation Function

The vector autocorrelation function (*ACF*) is calculated from the vector covariance function by normalizing with the lag-0 covariances. The *ACF* captures the correlation of state variables across lags, which correspond to bounces in the model. The vector autocorrelation function encompasses both the autocorrelations and the cross-correlations of the two state variables. We denote vector autocorrelation function by *ACF*, with *ACF* v-v and *ACF* t-t as the autocorrelations for ball release velocity and impact period, respectively. The cross-correlations are denoted by *ACF* t-v and *ACF* v-t. For a given lag, the vector autocorrelation function equals the Pearson product-moment correlation of state variable (say release velocity) and a lagged copy of a state variable (say release velocity lagged by two bounces). At steady state a small deviation of a state variable (such as release velocity) propagates across bounces. If the system is stable, the deviation dies out relatively quickly and

the correlation function shows a quick decrease towards zero. If the system is less stable, a deviation takes longer to die out.

The *ACF* of a single trial was calculated from the detrended release velocity and impact period. To calculated a mean *ACF*, the *ACF's* were averaged over all trials and subjects for a given α. To calculate a measure of variability of the ACF, the standard deviation over trials was calculated for each subject and averaged over subjects.

To calculate the predicted *ACF* from the model, four parameters have to be specified. Two parameters are specified by the experimental design: g is fixed at -9.81 and α varies with condition. The other two parameters are calculated from the performance of each participant: the average bouncing frequency ω_r and the average impact phase θ. From these four parameters, the racket amplitude and the Jacobian can be calculated (See Model section). Then by applying Equations 10, 11 and 12, the *ACF* is calculated. In sum, *ACF* is calculated from the linearized model, by incorporating some performance estimations from individual participants. The comparison of *ACF* from the data and *ACF* from the model serves as a test of the model.

Covariation between Execution Variables

The ball bouncing task can be viewed as a goal directed task such that the goal of every bounce is to propel the ball to the fixed target height. There are three execution variables on the task level that determine the bounce height of the ball: the impact position X_{imp}, the ball velocity immediately before impact V_{bal} and the racket velocity at impact V_{rac}. The bounce height is related to these variables by the ballistic flight equation (Equation 1) and the elastic impact equation (Equation 2). More specifically, the height error (*HE*) can be written as a function of these three execution variables:

$$HE = ((1+\alpha)V_{rac} - \alpha V_{bal})^2/(2g) + X_{imp} - X_{target} \quad (13)$$

where X_{target} is the target height set at 0.8 m in the present study. By applying this equation, a triplet of execution variables (X_{imp}, V_{bal}, V_{rac}) leads to one result variable (*HE*). The covariation of the execution variables is calculated in two steps. First, triplets from all the bounces in one trial are randomly permuted to produce a

pseudo-random data set, which contains no covariation between the three execution variables. By applying Equation 13 the height errors corresponding to this random data set can be computed. The average absolute value of the height errors is devoid of any covariation between the execution variables. By repeating the randomization many times (100 in our case), and averaging the resulting absolute value of the height errors, we calculate a reference value. Second, the difference is taken between this reference value and the observed average absolute value of the height errors without permutation. This difference quantifies the contribution from covariation between execution variables in units of the performance measure.

Results

Performance Measure (HE)

As shown in Figure 3, average *HE* values pooled over trials and participants were positive, indicating that there was a tendency for participants to overshoot the target. The overshoot was larger with higher α. The error bars denote the average standard deviations over all the trials for the different α conditions. As evident from the figure, the variability of *HE* decreased with increasing α.

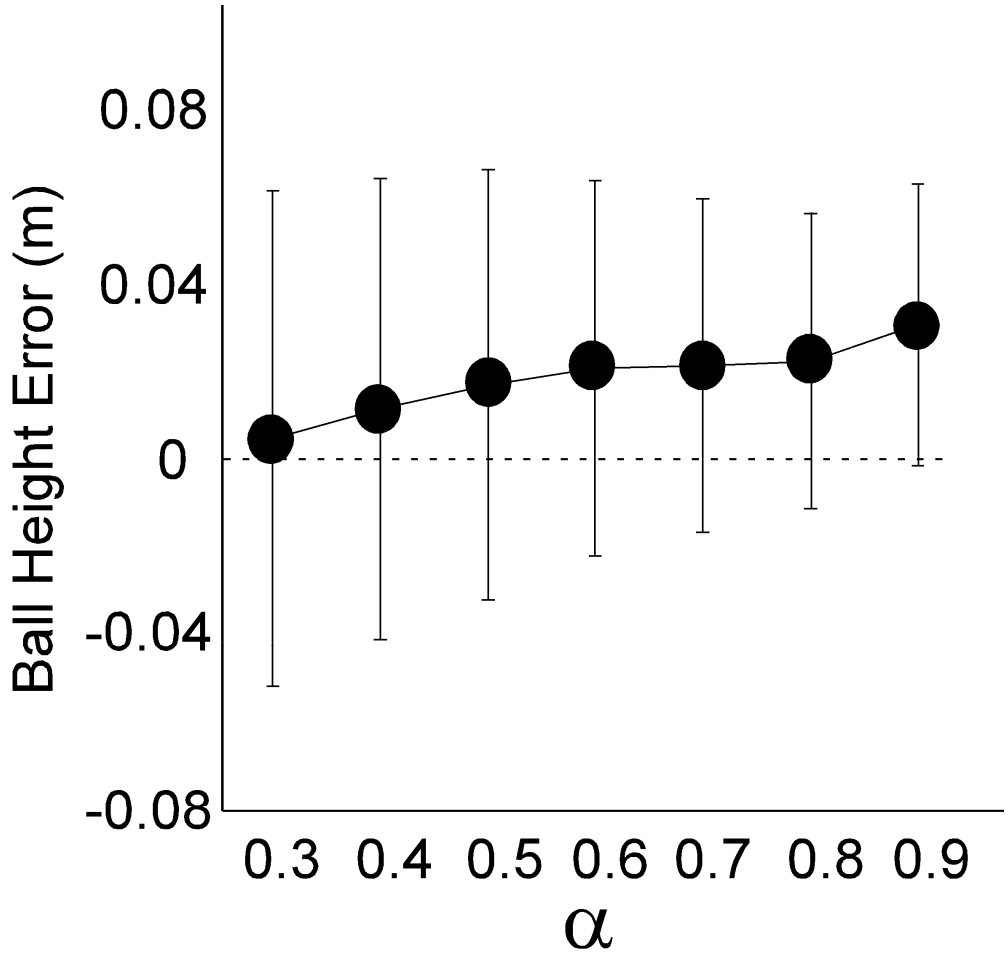

Figure 3: Ball height error (mean ± std) across α conditions.

Racket Kinematics: Racket Amplitude, Racket Period, and Impact Acceleration

The means and standard deviations of racket amplitude (A_R), racket period (T), and impact acceleration (AC) within trials were calculated and then averaged across participants (Figure 4). There are clear changes in these variables with α: A_R decreases and T increases with increasing α. The decrease in A_R can be understood from the physics of the task: the more "bouncy" the racket is, the less racket velocity is necessary at impact to bounce the ball to a fixed target height. Hence, participants moved the racket less with higher α leading to smaller A_R. It was also found for the more bouncy racket that participants tended to lower the racket impact position by a few centimeters (result not shown). This was allowed as participants were free to choose their preferred impact position, even though the target height was fixed. This effectively increased the ball amplitude and hence racket period with increasing α.

Average impact accelerations were negative. This result is consistent with prediction 1 and shows that subjects exploited passive stability of the task for all α conditions. The range of impact acceleration also fell within the boundaries as predicted by the model (Equation 4). The variability of all these racket kinematic variables decreased with increasing α. To elucidate this tendency, the standard deviations of these variables within each trial were calculated and then averaged over participants for each α condition, as shown in Figure 4. A_R, T and AC show a clear decrease in variability with increasing α.

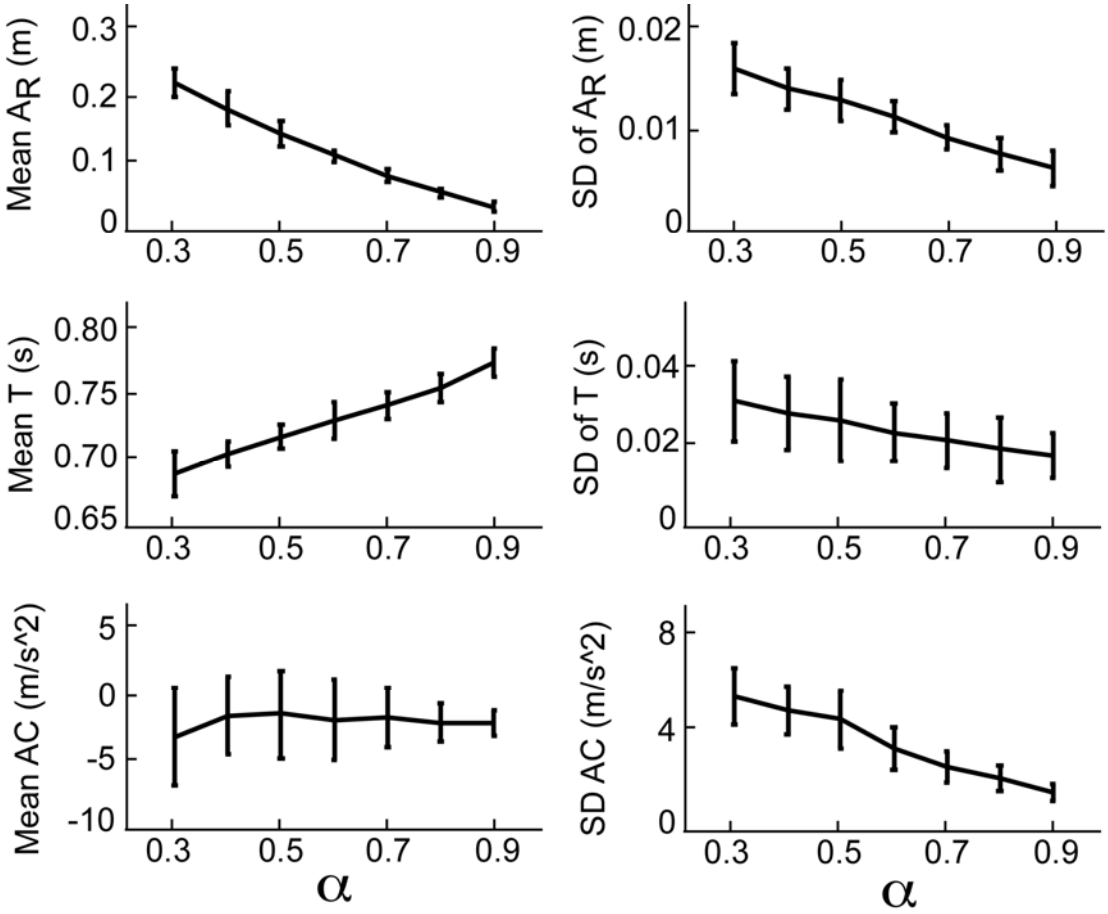

Figure 4: The mean and standard deviation of racket amplitude, racket period and impact acceleration. The error bars denote the standard deviations across participants.

Covariance Functions

Figure 5 illustrates lag-1 *ACF*s of v-v, v-t, t-v and t-t as a function of α. The first thing to notice is the similarity between all four *ACF*s, both the experimentally

observed ones and the theoretically predicted ones. This is caused by the strong coupling between the two state variables, ball release velocity and time between contact. A positive fluctuation in one state variable (say ball release velocity) necessarily leads to a positive one in the other (a larger time between contact). Secondly, both the lag-1 *ACF*s from the data and the model increased with increasing α, indicating a decreasing stability. Thus, model prediction number 2 about the α dependency of stability was confirmed by the data. *ACF* values from the model and from the data matched closely for $\alpha \leq 0.5$, but not for higher α conditions. In these conditions, the data exhibit smaller lag-1 *ACF* than the model, indicating that better dynamic stability was achieved by the actor than the model predicted. This observation is consistent with prediction 3.

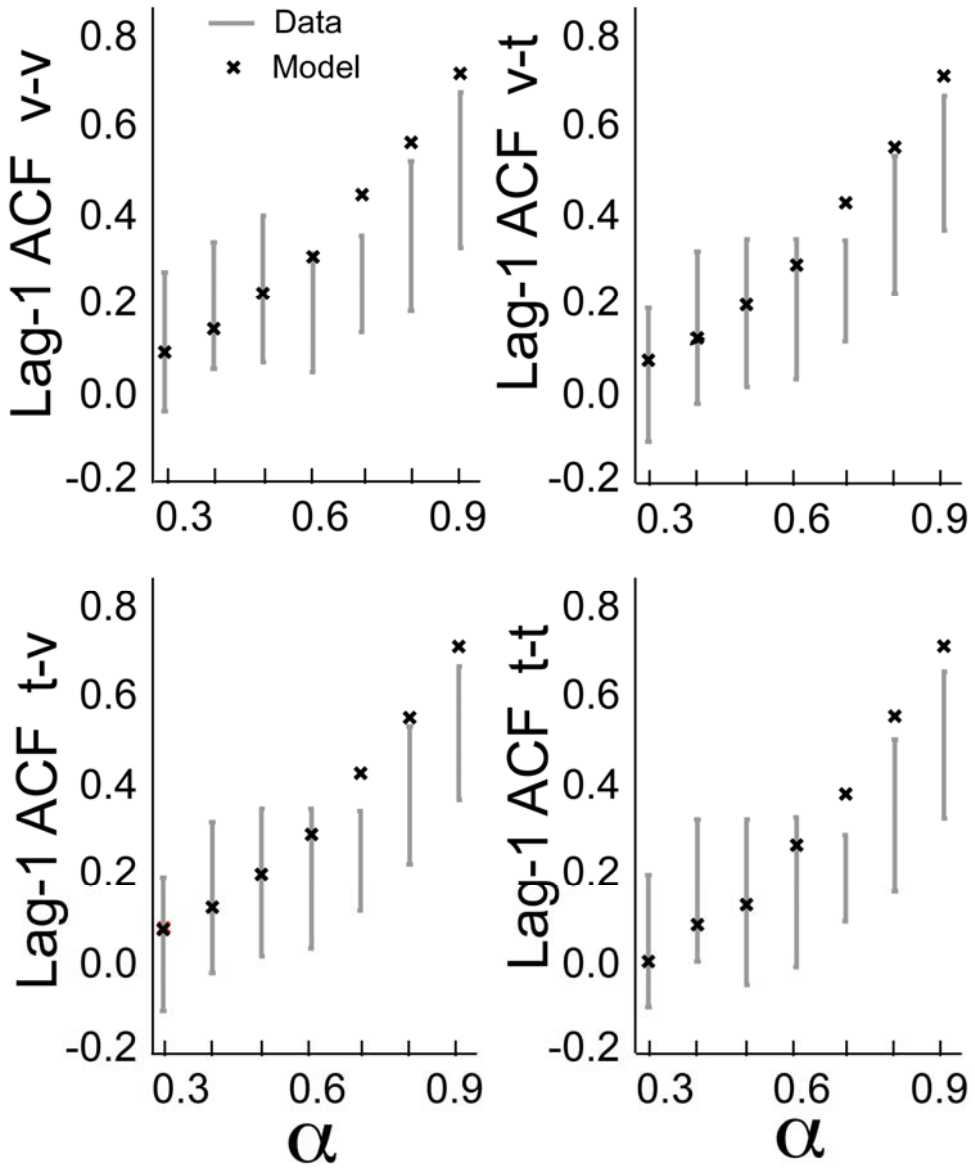

Figure 5: Lag-1 auto- and cross- correlation functions of release velocity and impact period as a function of α. Derived both from the data (grey lines) and the model (black crosses). Error bars denote the standard deviations across trials.

Figure 6 shows the autocorrelation function between release velocities (*ACF* v-v) for different α conditions, both from the model and the data of a typical participant. From Figure 5 one might get the idea that there is a difference between observed *ACF*s and the theoretical ones for larger coefficients of restitution but that the difference is not that large: for $\alpha > 0.6$ the theoretical predictions are outside of a band of 1 std from the observed mean but less than say 1.5 times the std. Figure 6 shows

the difference to be larger especially for $\alpha > 0.7$. Since the eigenvalues of the Jacobian are complex (imaginary part is non-zero), the theoretical *ACF*s display oscillatory behavior quite different from the experimentally observed *ACF*s. The autocorrelation of time between impact and the crosscorrelations between release velocity and period revealed similar patterns as *ACF* v-v, therefore are not shown.

As discussed in prediction 3, we interpret the increasing difference between the predicted and observed *ACF* as indirect evidence for an increasing contribution of active control. To uncover potential contribution of active control more directly, we performed a covariation analysis.

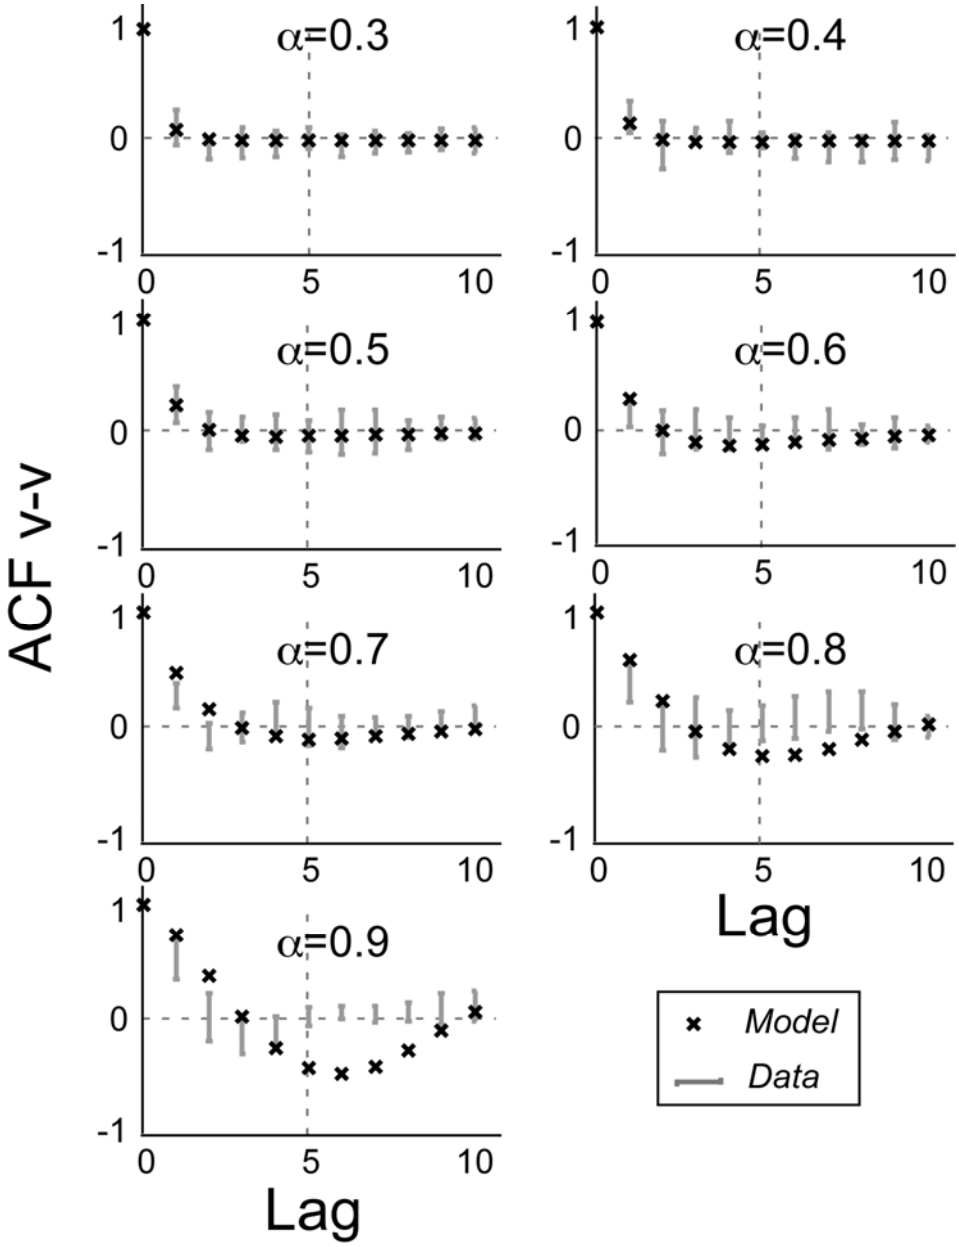

Figure 6: Autocorrelation functions of release velocity as a function of lag in different α conditions. Derived from the data (grey lines) and the model (black crosses). Error bars are the standard deviations across trials.

Covariation between Execution Variables

Covariation captures the contribution of mutual compensation between execution variables in reducing the variability of performance measures. Larger covariation values indicate more control. The statistical framework of covariation analysis assumes independent execution variables. As seen in the previous section subsequent bounces are statistically correlated, with correlations dying out after two bounces. In order to eliminate the observed correlation between adjacent bounces, all the bounces within one trial were down-sampled by 3, i.e., every fourth bounce was grouped together. Thus, the original data set from a trial was separated into 3 data sets for the purpose of the covariation analysis. Covariation was calculated separately for each data set and averaged to obtain a mean covariation for that trial. Covariations from all trials within each α condition were then averaged.

To demonstrate that covariation is indicative of active control rather than a consequence of the task dynamics, covariation was calculated from the stochastic passive stability model. With noise at a level comparable to human performance, the model was simulated with 100 bounces in one trial. A total of 100 simulated trials were obtained for each α condition and covariation was calculated. As shown in Figure 7 for real human performance covariation increases with increasing α, indicating that mutual compensation between execution variables was more pronounced in the dynamically less stable conditions. Covariation from the simulated data set presents a very different picture: the values are significantly smaller and stay close to zero for the different α conditions. This indicates that the model with passive dynamics can only produce negligible levels of covariation between execution variable. Therefore, the large covariation from the actual performance was not an artifact from the task dynamics, but stems from active control of the actor.

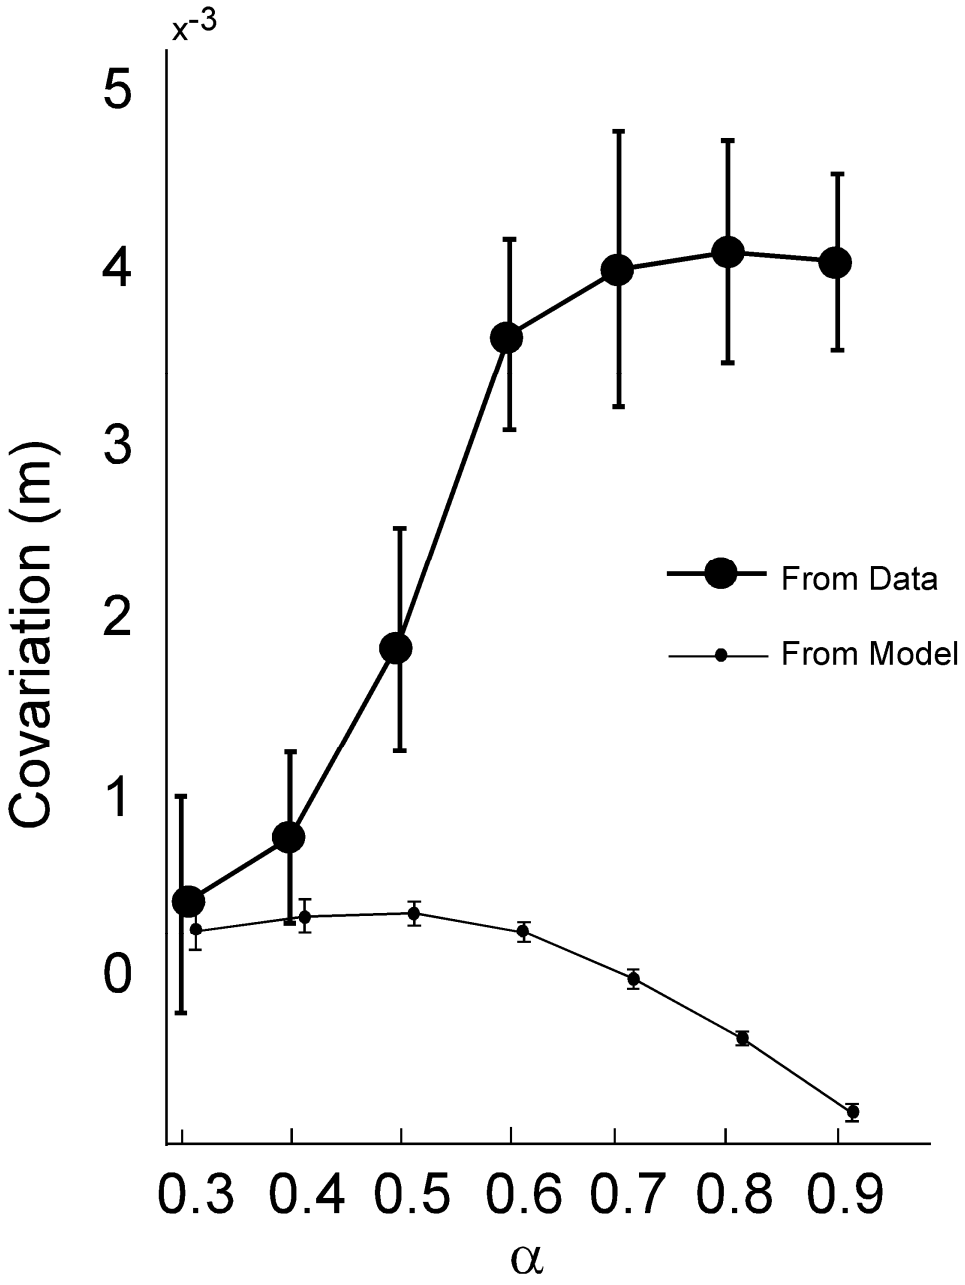

Figure 7: Covariation between three execution variables (X_{imp}, V_{bal} and V_{rac}), averaged for each α condition. Error bars stand for the standard deviation between trials within each α condition. The bold line is from the data and the thin line is from the model simulation.

Discussion

The present study aimed to reveal the influence of passive stability and active control on the performance of a rhythmic ball bouncing task at steady-state

performance, where exploitation of passive stability had been found to be a preferable action strategy (Schaal et al., 1996; Dijkstra et al., 2004). This strategy alleviates the necessity for cycle-to-cycle adjustments of racket movements based on the perceived changes of the ball movement. However, whether active control of the racket movement exists at steady-state performance has been left untested. We recognized that the racket movement is stereotypical with little variations from cycle to cycle and subtle adjustments are hard to test as evidence for control. The present study chose the covariance structure of two state variables of the task dynamics to evaluate the correlation from one cycle to the following cycles. To make predictions about the covariance structure, the original deterministic model was extended to a stochastic model with the addition of noise terms to the dynamics (see Dijkstra et al., 2004). The coefficient of restitution of the ball-racket contact was identified as a key parameter influencing stability in the model. Dynamic stability of the system was predicted to decrease with increasing α such that lag-1 correlations should increase indicating a higher correlation between one cycle to the next. We speculated that with less dynamical stability of the system, the actor would turn to employ more active control in order to achieve a comparable stable performance. As a consequence, with more contribution of active control in the higher α conditions, lag-1 correlations should have smaller values than the model predicted since the model only includes passive stability. Further, there should be increasing mismatch between the data and the model predictions with increasing α.

The major performance measures are the means and standard deviations of height error (*HE*) in each individual trial. It turned out that actors consistently overshot the target, as indicated by positive mean values of *HE*. Mean *HE* increased when the racket became "bouncier" at larger α. More importantly, the standard deviations of *HE* decreased with increasing α, showing a decreased variability in performance. This seems to be contradictory to the expectation based on the passive model, which predicts that dynamic stability should decrease with increasing α. Less stability of the system should lead to more variability in the human performance. Regarding the racket kinematics, it was found that the racket period increased and the racket

amplitude decreased with increasing α. The increased racket period can partly be explained by the increasing overshoot of the ball over the target and partly by a lowered impact position. The decreased racket amplitude is expected as the racket moves less to lead to a smaller velocity at impact when the racket becomes bouncier. Thus, the observed means of racket period and racket amplitude can be understood by the physics of the task. However, all these variables showed less variability with increasing α. It appears that actors managed to reduce their variability in movement execution when stability of the task deteriorated with changes in α. Another important finding is that impact acceleration was negative in all α conditions. A previous Lyapunov analysis suggested that mean impact acceleration between -2 and -5 m/s^2 is optimal for exploitation of passive stability (Schaal et al., 2006). The mean results for each α condition were in this range, signifying that utilizing passive stability was still a preferable strategy adopted by the actor when faced with different degrees of dynamic stability of the task.

The covariance structure of the state variables revealed that stability of the task was indeed reduced when α became larger. Lag-1 correlations within and between the two state variables all showed an increasing tendency (Figure 5). As lag-1 correlations reflect the degree of dissipation of small fluctuations from one bounce to the next, this result shows that stability of the task decreases with increasing α, as predicted by the model. However, the actor's performance showed smaller lag-1 correlations for higher α, indicating that actors achieved better and more stable performance for these conditions by using a mechanism other than relying on passive stability. Moreover, this mismatch between the model prediction and the actual performance increased with increasingly larger α. The results on autocorrelation and cross-correlation over different lags reinforced this impression. The model prediction and the actual performance showed similar pattern for $\alpha < 0.6$ conditions; for increasingly larger α, these two departed from each other more. These results suggested that the performance for low α conditions can be well accounted by the passive model. However, with less dynamic stability in high α conditions, actors employed other strategies, presumably by using more active control.

One piece of evidence for more active control in high α conditions is that variability of racket kinematics, including racket period, racket amplitude and impact acceleration, decreased with increasing α as stated above. This result shows that actors more carefully controlled their racket movement when confronted with less stability of the task. Except impact acceleration, these variables are descriptive about the *continuous* racket movement. Focusing on the discrete impact events, we recognized that three execution variables at impact determine the ball height following that bounce. This many-to-one relationship between execution variables and the result variable is related to the redundancy problem (Bernstein, 1967). Many researchers suggested that compensation between multiple execution variables to achieve a single goal is one of the solutions to the redundancy problem (Bernstein, 1967; Arutyunyan et al., 1968, 1969; Scholz & Schöner, 1999; Müller & Sternad, 2003, 2004). Using the permutation method suggested by Müller and Sternad (2003), we calculated the covariation of execution variables in different α conditions. The covariation is regarded as a signature of active control as it reflects the degree of mutual compensation between execution variables at impact. It turned out that covariation increased gradually and consistently with increasing α. This result is a second piece of evidence that more active control was utilized for higher α conditions. To elucidate that larger covariation for higher α conditions was not caused by the passive dynamics of the task, the stochastic model was numerically simulated to yield predictions about covariation as a function of α. The passive model without any active control component produced much smaller covariation than exhibited by the actual performance, suggesting the increasingly large covariation was indeed a reflection of more active control.

The present study reveals that when the stability property of a task is modified the actor can flexibly adopt two different control strategies: exploitation of passive dynamics and perception-guided active control. When the dynamic stability of the task is high in low α conditions, active control is not pronounced and exploitation of passive stability is the primary strategy. The support for the first claim is that covariation between execution variables was relatively small and variability in

movement execution was large. The support for the second claim comes from the good match between the model prediction and the actual performance in terms of covariance structure of state variables, suggesting that the passive model suffices to explain the actor's performance. More support for the second claim comes from the impact acceleration, an indicator of exploitation of passive stability, that was in the optimal range as predicted by the model.

When the dynamic stability of the task was reduced with increasingly larger α, there was evidence that active control became more pronounced while passive stability was still exploited. Covariation increased and variability in racket kinematics dropped with increasing α values, indicating a better and tighter control of action. Impact acceleration remained negative in the optimal range. It appears that actor resort to active control when passive stability becomes insufficient in maintaining comparable stable performance as in low α conditions.

The stability the present study examined is the consistency of kinematics of an external object that is manipulated repetitively by the actor. This stability is different to what is usually defined as stability in human control system where a stable movement of the human body itself is at concern. Hasan (2005) reviewed many studies about the response of human control system to mechanical perturbations and made a long list of various phenomena that are defined as stable behaviors. These stable behaviors included staying static though free to move, moving but within certain boundaries, successful maintaining articular relationships without damage or irritation, reliable movements, moving but showing little random or chaotic motion and so on. Despite the diversity of these operational definitions of stability in the motor control literature, stability of a dynamical system can be rigorously defined: Jordan and Smith state that "a dynamic system is stable in the Lyapunov sense if infinitesimally small variations in the state of the system remain infinitesimally small forever" (Hasan, 2005; Jordan & Smith, 1987).

In motor control, stability is usually equated with the inverse of variability. However, stability is more than invariance or a reduced level of variability in certain performance measures. Specifically, dynamical stability is a formal and rigorous

concept in mathematics and physics as mentioned above. It can be quantified independently from measured variability in human movement. This point is illustrated in the present study: when the passive stability of the movement system decreased with increasing α, the variability of ball amplitude actually decreased, opposite to common expectation. The analysis on the execution of the task, instead of merely the result variable, revealed that actors adopted a higher degree of active control of the racket when dynamic stability was reduced. This result suggests that evaluation of stability of overt performance measure based on examination of its variability is not sufficient to understand the control that leads to performance changes. By modeling the physics of the task system and by analyzing execution details, the relative contribution of open-loop dynamics and closed-loop control can be distinguished and provide more insights into the underlying reasons for variability changes in performance.

References

Arutyunyan, G. H., Gurfinkel, V. S., & Mirskii, M. L. (1969). Organization of movements on execution by man of an exact postural task. *Biophysics, 14,* 1162–1167.

Bernstein, N.A. (1967). *The coordination and regulation of movements.* London: Pergamon Press.

Bühler, M., Koditschek, D.E., & Kindlmann, P.J. (1994). Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research, 13*, 101-118.

Collins, S.H., Ruina, A., Wisse, M., & Tedrake, R. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science*, *307*, 1082-1085.

de Rugy, A., Wei, K., Müller, H. & Sternad, D. (2003). Actively tracking 'passive' stability in a ball bouncing task. *Brain Research, 982,* 64-78.

Dijkstra, T.M.H., Katsumata, H., de Rugy, A. & Sternad, D. (2004). The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies*, *11, 3*, 319-345.

Gander, W. & Hrebicek, J. (1993). *Solving Problems in Scientific Computing using Matlab and Maple,* Verlag: Springer.

Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transition in human hand it, movements. *Biological Cybernetics, 51*, 347-356.

Hasan, Z. (2005). The human motor control system's response to mechanical perturbation: Should can it and does it ensure stability? *Journal of Motor Behavior*, *37, 6*, 484-493.

Jordan, D. W., & Smith, P. (1987). *Nonlinear ordinary differential equations*. New York: Oxford University Press.

Katsumata, H., Zatsiorsky, V., & Sternad, D. (2003). Control of ball-racket interactions in the rhythmic propulsion of elastic and non-elastic balls. *Experimental Brain Research, 149,* 17-29.

Kendall M., & Ord, K. (1989) *Time series*. Oxford University Press.

McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research 9, 2*, 62–82.

Müller, H. & Sternad, D. (2003). A randomization method for the calculation of covariation in multiple nonlinear relations: Illustrated at the example of goal-directed movements. *Biological Cybernetics, 89,* 22-33.

Müller, H., & Sternad, D. (2004). Decomposition of variability in the execution of goal-oriented tasks – Three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance, 30, 1*, 212-233

Schaal, S., Atkeson, C.G. & Sternad, D. (1996). One-handed juggling: A dynamical approach to a rhythmic movement task. *Journal of Motor Behavior, 28*, 165-183.

Scholz, J.P., Kelso, J.A.S., & Schöner, G. (1987). Nonequilibrium phase transitions in coordinated biological motion: Critical slowing down and switching time. *Physics Letters A, 123*, 390-394.

Schöner, G., Haken, H., & Kelso, J.A.S. (1986). A stochastic theory of phase transitions in human hand movement. *Biological Cybernetics, 53*, 442-452.

Shumway, R.H., & Stoffer, D.S. (2001) *Time series analysis and its applications* Springer.

Sternad, D., Duarte, M., Katsumata, H. & Schaal, S. (2001). Bouncing a ball: Tuning into dynamical stability. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 1163-1184.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2000). Dynamics of a bouncing ball in human performance. *Physical Review E, 63*, 011902-1 – 011902-8.

Tufillaro, N.B., Abbott, T., & Reilly,J. (1992). *An experimental approach to nonlinear dynamics and chaos*, Redwood City, CA: Addison-Wesley.

Wei, K., Dijkstra, T.M.H., & Sternad, D. (submitted). Perturbations of passive stability in a ball bouncing task. *Journal of Neurophysiology.*

CHAPTER 5

Variability and Stability during Learning and Adaptation

Chapter 5 contains an original manuscript prepared for submission to a journal.

"Variability and Stability during the Acquisition and Adaptation of a Rhythmic Skill"

Tjitske Boonstra, Kunlin Wei, Dagmar Sternad

Department of Kinesiology and Integrative Biosciences, Pennsylvania State University

Abstract

It has been argued that acquisition of a skill is characterized by the exploration and exploitation of variability within a given task. Actors first explore the space of solutions to find a stable strategy for a given task and then further decrease variability. One of our previous studies had illustrated that during acquisition of the ball bouncing skill the increasing attunement to passive stable is accompanied by a reduction of variability in performance. The present study aims to confirm that skill improvement is indeed characterized by utilizing passive stability but additional to show that the structure of variability changes independently. Applying a novel analysis method the learning process will be evaluated in the execution space, which is spanned by relevant execution variables that lead to the movement outcome variable. A set of individual events will be represented as a distribution in the execution space that contains a manifold of solutions. The study also studied the adaptation of the learned skill in the execution space when the original preferred solutions became unstable. The process of adaptation was characterized by exploration in execution space and simultaneous re-establishment of the passively stable regime.

Introduction

While many basic behaviors are developed early in life, such as walking upright and grasping an object, humans continually face new perceptual-motor tasks ranging from tying shoelaces or hammering a nail into the wall to the numerous "unnecessary" challenges created by leisure time activities. Often they are invented exactly *because* they pose new challenges to the perceptual-motor system. Apparently, humans like to and constantly have to modify and adapt their learnt movements to the ever-changing demands from the environment. Acquisition and adaptation of skills is often viewed as the same process although they may happen at different time scales and may be of different degree.

In studies on motor learning the typical indicators for documenting performance improvement is a decrease in error or in the variability of a performance variable, such as distance from a target or period and amplitude in a rhythmic performance (Darling & Cooke, 1987; Gottlieb, Corcos, Jaric, & Agarwal, 1989; Worringham, 1991, 1993). Reducing variability has often been equated with an increase in stability. This simple inverse relationship, however, obscures the fact that empirical variability can be indicative of many different facets ranging from the obvious "lack of control" to more beneficial aspects, such as compensatory variation between parameters, and exploration of new tasks. Dysfunctional behavior is not only evidenced by abnormal levels of variability, as seen in tremor or excessive movements, but also by the absence of fluctuations in stereotypy and, consequently, the absence of adaptability of the system.

Stability is a similarly broad concept and the commonsense use of the word as invariance of a task parameter over repetitions does not capture its full meaning. Specifically in dynamical systems and control theory stability is a well-defined concept. With a view to biological systems, stability is defined as the ability of an organism to persevere in a particular behavior in spite of perturbations from an internal

or external environment (Jagacinski & Flach, 1999; Strogatz, 1994). More formal definitions hold for specific aspects of stability, such as global, neutral, or Lyapunov stability (Jordan & Smith, 1987; Strogatz, 1994). Stability was probably first recognized in its central role for biological movement systems by Bertalanffy who postulated that goal-directed behavior is governed by equifinality, i.e., equilibrium states (Bertalanffy, 1950, 1973). In a similar vein the dynamical systems perspective has proposed that the high-dimensional movement system is organized around stable states. The most fertile line of research has focused on rhythmic uni- and bimanual coordination, which was modeled as coupled nonlinear oscillations which show limit cycle behavior (Amazeen, Amazeen, & Turvey, 1998; Kelso, 1995). Important for the present argument is that in these studies fluctuations were only captured in terms of stochastic noise processes which represented processes operating at another spatio-temporal scale of the system, e.g., faster frequencies at the neural or molecular level. Further, variability was only analyzed in its inverse relationship with stability. The present study will extend this approach and connect the analysis of stable behavior with an analysis of the structure of variability.

In a series of previous studies Sternad and colleagues investigated the skill of bouncing a ball rhythmically on a racket as an exemplary perceptual-motor task in which dynamical stability plays a central role (Katsumata, Zatsiorsky, & Sternad, 2003; Schaal, Sternad, & Atkeson, 1996; Sternad, 2000; Sternad, Duarte, Katsumata, & Schaal, 2000, 2001). This task has received attention in both the robotics and the motor control literature (Aboaf, Drucker, & Atkeson, 1989; Bühler, 1990; Buhler, Koditschek, & Kindlmann, 1994). The actor (or actuator) holds a racket in his/her hand and hits a ball into the air keeping a consistent period and amplitude. The cyclic ball-racket interactions have been modeled by a system of nonlinear discrete equations in modification of the “hopping particle model” that was originally developed in the applied mathematics literature. There, it provided a study case to investigate dynamical stability and the period-doubling route to chaos (Guckenheimer & Holmes, 1983; Tufillaro, Abbott, & Reilly, 1992). The application of this model to

human performance was first advanced by Sternad and colleagues in a series of studies using an experimental set-up that closely mimicked the physical model of the ball bouncing on a periodically moving table (Dijkstra, Katsumata, de Rugy, & Sternad, 2004; Schaal et al., 1996). Stability analyses of the task dynamics indicate that there exists a stable regime if the racket hits the ball with negative acceleration. Tuning to this stable regime, small fluctuations in performance can be dissipated without active error corrections. This strategy is optimal in terms of control effort and it stands in contrast to the more computationally expensive approach of classical control theory in which a deviation of the ball has to be compensated for by an explicit change of the actuator trajectory. Empirical studies showed that indeed experienced actors are sensitive to the stability properties of the task and bounce the ball with negative racket acceleration (Dijkstra et al., 2004; Schaal et al., 1996; Sternad et al., 2001).

How do humans learn to exploit this stability? It has been proposed that humans first try to find the stable points of a system and from there on further refine their skill. For example, (Beek & van Santvoord, 1992) suggested that learning a new motor skill involves the discovery of invariance or fixed points in the perceptual-motor workspace of the skill. (Schöner, Zanone, & Kelso, 1992) suggested that the dynamics of the task evolves with practice to develop an attractor solution near the to-be-learned pattern. (Milton, Small, & Solodkin, 2004) reported that the first step in learning a motor task consists of an increasing reliance on the self-regulatory aspects of the motor task, i.e., the dynamics of the task, followed by a minimization of intentionally directed corrective movements. However, none of these studies have combined a systematic quantification of both stability and variability.

One previous study on ball bouncing already addressed skill acquisition as a process of tuning to stable regime (Dijkstra et al., 2004; Sternad et al., 2001). It was found that novice bouncers started their practice with positive racket impact acceleration, which was unstable solution according to the stability analysis of the task. With practice, the impact acceleration gradually decreased to negative values

within an optimally stable range as predicted by the model. The variability of the ball amplitude decreased simultaneously with this tuning to passive stability regime. While this finding is consistent with the common belief that variability is equal to the inverse of stability, it does not establish the causality relationship between variability and stability in human performance. The present study aims to firstly replicate that skill improvement is indeed characterized by utilizing passive stability but additional to show that the structure of variability changes independently. In order to reveal changes in the *structure* of variability during skill acquisition, the so-called TNC decomposition as proposed by (Müller & Sternad, 2003, 2004) will be adopted. This method parses changes in variability during learning into three conceptually different components, tolerance *T*, noise reduction *N*, and covariation *C*.

The TNC-analyses are performed in so-called execution space. The execution space is spanned by the execution variables that determine the task outcome completely. For example, in a throwing task, the release angle, velocity, and position determine whether the ball is going to hit the target. If the goal of performance is to achieve a given ball amplitude, it follows from basic mechanics that the ball amplitude is completely determined by three execution variables: racket velocity ($\dot{x}_R$), ball velocity before impact ($\dot{x}_B^-$) and impact position (x_B). The result, i.e., ball amplitude, can be calculated with the equations of ballistic flight in a gravitational field given a coefficient of restitution capturing the energy loss at the moment of contact. Each bounce is thus regarded as a separate event or data point in execution space. An example of execution space of the ball bouncing is given in Figure 1. The 3-dimensional surface is the solution manifold containing all the possible combinations of the execution variables for perfect hits that lead to a target ball amplitude of 0.8m. Further, all points in this space can be assigned a value of the result variable. Figure 2 depicts this by showing a 2-dimensional section of the three-dimensional space, taken at the impact position of 0.285m. In each panel, the white areas represent combinations of execution variables that result in zero error (solution manifold) and the increasingly darker shades indicate increasingly larger

errors. The black data points stand for individual bounces. Different clouds of data points in these 4 subplots schematically illustrate performance at different stages of learning.

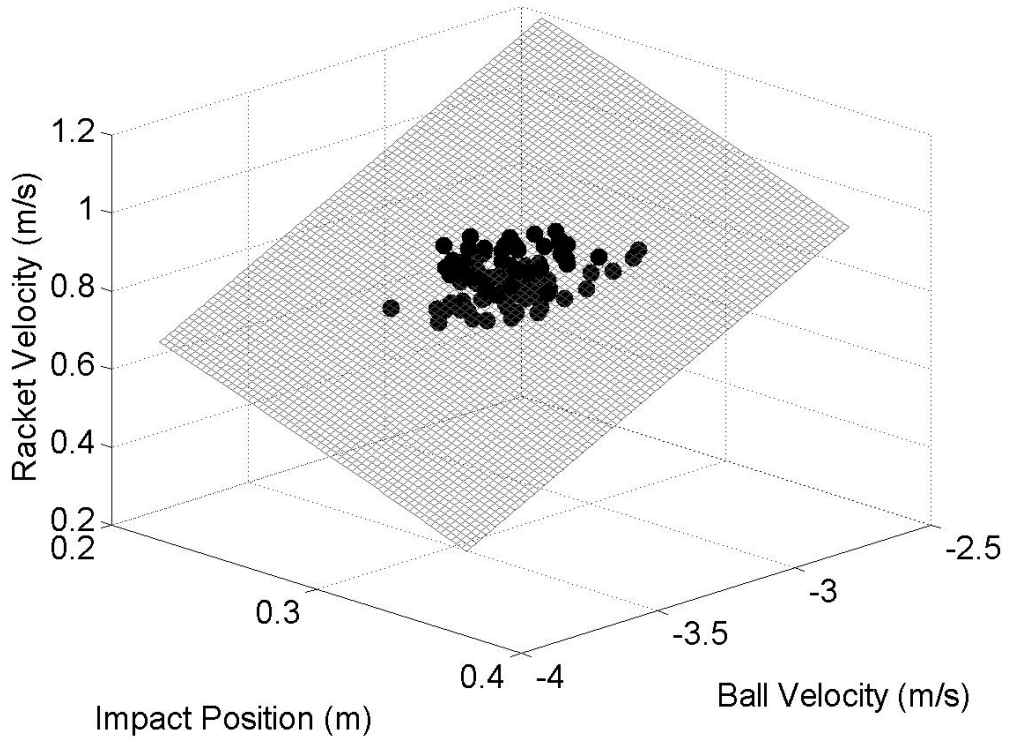

Figure 1: Execution space for ball bouncing with the surface representing all possible solutions with zero error, i.e., hitting the target exactly. The dots denote the bounces of one trial.

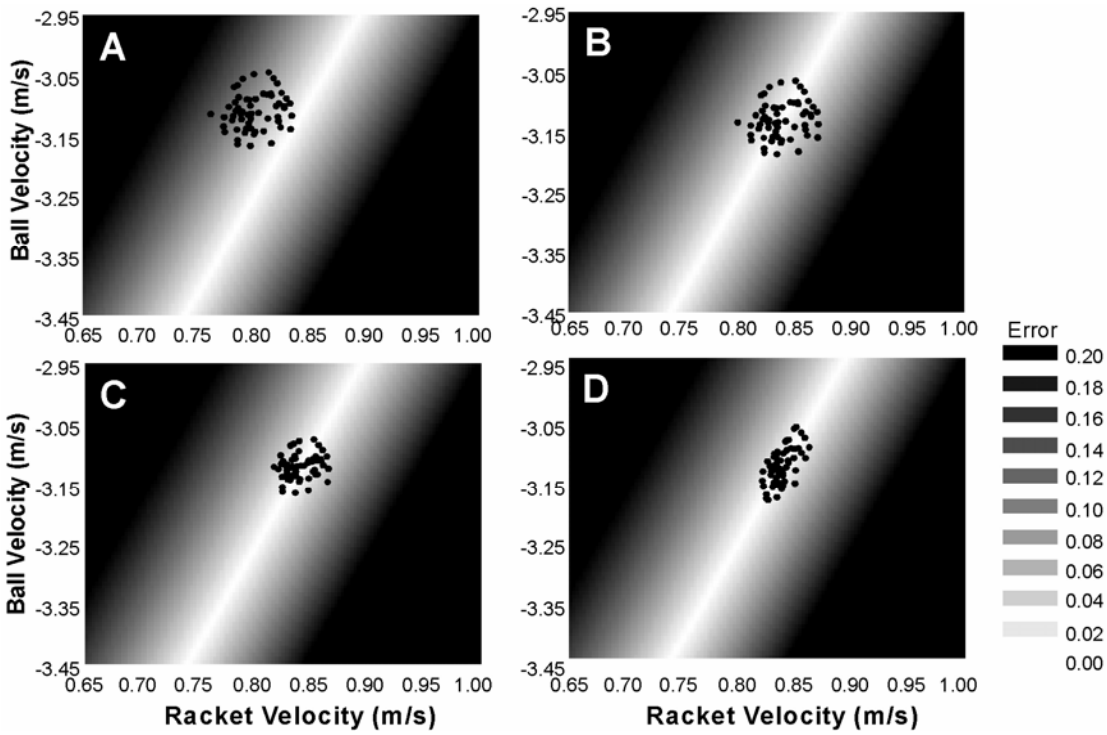

Figure 2: Execution space and solution manifold in ball bouncing taken at a section of the impact position at 0.285m. The white shading indicates all possible solutions (ball velocity and racket velocity combinations) with zero error in ball height. The grey shades show the corresponding error (see legend). Four panels represent performances at different stages of learning (details in text). Each black dot denotes one bounce.

Task Tolerance (T). The first component is task tolerance which accounts for the search for areas with better solutions in execution space. As shown in Figure 2, data points at the average location denoted by A move to location B. One can intuit that a broader solution space is less sensitive or more tolerant with respect to perturbations or errors. Therefore, finding the areas in the execution space that are more tolerant to errors will contribute to the improvement in performance. This part of improvement in performance is captured by *T*.

Noise Reduction (N). It is commonly known that a certain degree of fluctuations or noise is inevitable in movement performance and can never be completely eliminated. Reduction of such fluctuations as a consequence of practice is widely

acknowledged as an important reflection of skill improvement (e.g., Darling & Cooke, 1987; Gottlieb, Corcos, Jaric & Agarwal, 1989; Worringham, 1991, 1993). The TNC-method can separate the influence of noise reduction (reduced data scatter in execution space) on the performance improvement from other influences. Returning to Figure 2 consider data sets B and C. Both sets are at the same location in execution space; however, the dispersion of the two data sets differs significantly. This change in variability from B to C is referred to as reduction of Noise, *N*.

Covariation (C). The difference between data sets C and D can be viewed as a third factor contributing to decreases of the variability in the result. Data set D clusters along the direction of the solution manifold, whereas C does not. The two execution variables covary in the data set D such that the data distribution is aligned with solution manifold. It has long been recognized that execution variables such as joint angles can compensate for each other to lead to movement results of less variability, such as more consistent trajectories. Several studies have shown that human actors indeed use covariation to decrease the variability in the result. This covariation between execution variables has been proposed as a signature of control for movement coordination (Arutyunyan, Gurfinkel, & Mirskii, 1969; Bernstein, 1967; Müller & Sternad, 2003, 2004; Scholz & Schöner, 1999). The improvement in variability by covarying the execution variables is referred as Covariation, *C*.

The present study will simultaneously investigate the changes in stability and variability during learning of the ball bouncing task. While a model of the task and its stability analysis about dynamic stability has already been developed and shown to be insightful, the execution of bouncing will also be analyzed for its structure of variability, i.e., relative contributions of variability components. Changes in variability and stability and their relationship can be independently determined and examined during the process of skill acquisition.

Further, in order to study how stability and variability of performance change during adaptation to a novel execution space, the present study will introduce perturbations in the execution space after actors already stabilized their performance. This change of task properties was conducted by manipulating the solution manifold

of ball bouncing in a virtual set-up. From previous studies, it is expected that after practicing the task, the actor will adopt the strategy to utilize the passive stability with negative racket impact acceleration. By experimental manipulations, the introduced perturbations will effectively destabilize the originally preferred executions forcing the actor to migrate from the perturbed region to a new region in execution space. Will the actor change their strategy to employ more perception-guided active control as the errors introduced by perturbations become obvious? Or alternatively, will the actor continue to utilize passive stability while changing location in execution space? If the actor maintains the strategy of tuning to passive stability, how does the structure of variability change with this adaptation? During relocation in execution space, how far will the actor move to avoid the perturbed region? These questions will be addressed by analyzing the impact acceleration, distributional properties of the variability in the execution space and TNC variability decomposition.

Methods

Participants. Eight right-handed participants (4 male, 4 female) from the undergraduate and graduate population at the Pennsylvania State University volunteered to participate in this experiment (mean age: 31.2 years, SD: 7.3 years). They were informed about the purpose of the experiment, but were naïve about the nature of the manipulations in Part II of the experiment. Prior to the experiment the participants were instructed about the experiment and signed an informed consent form in agreement with the Institutional Review Board of the Pennsylvania State University. They were paid for their participation.

Apparatus. A virtual set-up was used to conduct the experiment (Figure 3). This set-up consisted of a real table tennis racket which had a rigid rod attached underneath. The rod itself was 1.5m long and consisted of three segments that were connected by hinge joints. With these linkages it was possible to freely move and tilt the racket in 3 dimensions. Even though the task did not require movements in 3D, this configuration of the rod ensured that displacements of the racket away from vertical movements did not encounter resistance. The lower segment of the rod ran through a noose in the

floor piece. On the side of the noose was a wheel that was rotated by the vertical displacements of the rod. The revolutions of this wheel were measured with an optical encoder. The signal of the optical encoder was transformed by a digital-to-analog converter and sent online to the computer that controlled the visual display. Note that only the vertical movements of the racket were measured. In principle the racket could move in all three dimensions, but these movements were negligible as the virtual ball movements only required vertical displacements of the racket (see below).

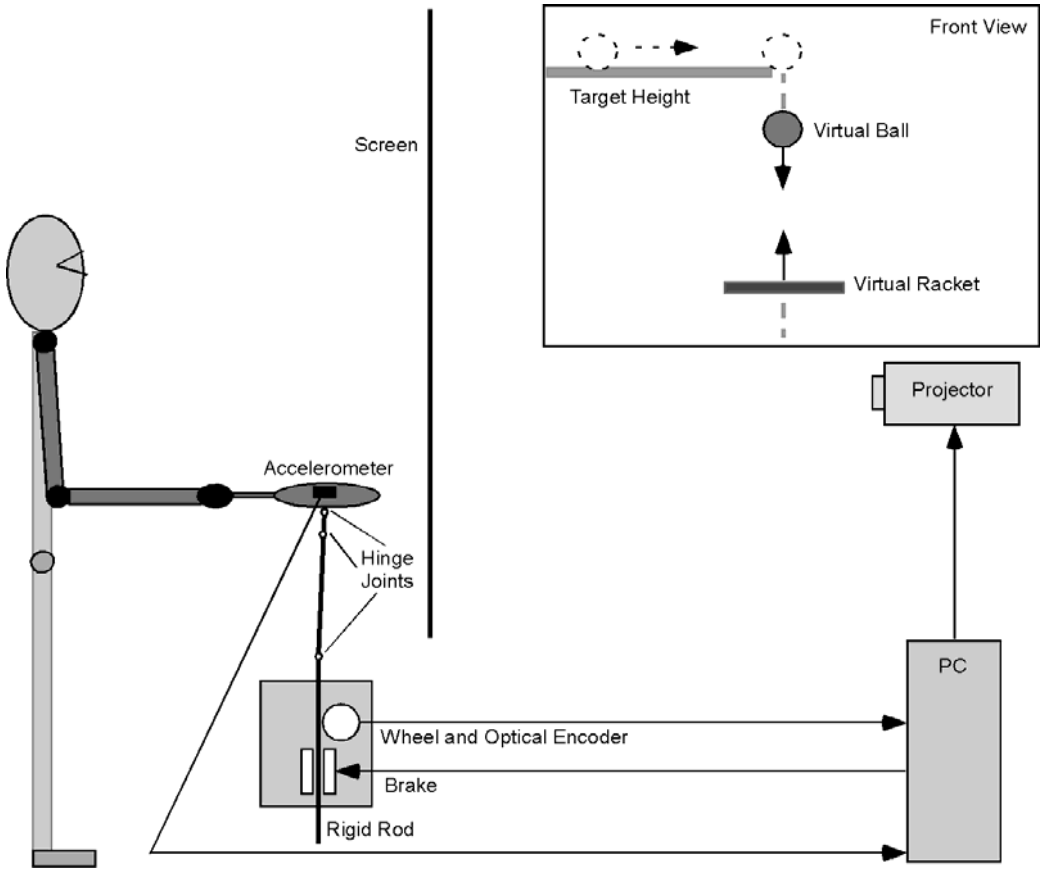

Figure 3: Virtual reality setup for the ball bouncing task

The participant stood behind the apparatus on a platform that could be adjusted for different heights. The criterion for adjustment was that in the rest position the participant held the racket in a comfortable arm position with the elbow and upper arm at right angle (shown in Figure 3). A back projection screen with width of 2.5m and height of 1.8m was used to project the movements of the racket and the ball. The participant was positioned approximately 1m behind this screen, with the projected racket directly in front of the real racket at the same height. Movements of the virtual

racket corresponded to the online recorded movements of the real racket such that 10cm of real racket displacement would be shown as 10cm of virtual racket displacement on the screen. The total latency between the recorded and the displayed signal of the racket movements was 11ms. Both the ball and the racket movements were confined to the vertical direction only.

In order to simulate the impact of the ball, a mechanical brake pulled on the rod at every ball-racket impact producing a force that was similar to the force imparted by a tennis ball. This braking action was produced by a solenoid and was applied for 30ms. Participants moved the real racket to interact with the virtual ball, while both the racket and the ball were shown on the screen. The mean sampling rate was about 800Hz. The refresh rate for the visual display was 75Hz. Custom-written software was developed in Visual C++ 6.0 (Microsoft) to collect the racket displacement data, compute the ballistic flight of the ball and the impact event, and subsequently display both the ball and the racket.

As the acceleration of the racket movements at the ball contact was a central variable in the analysis, an additional accelerometer was attached to the base of the real racket (uni-axial V94-41, range ±10g, resolution .01g, Colbourn, PA). Its analog signal was collected and converted to digital at the same sampling frequency as the displacement of the racket.

The ball's trajectory was calculated using the equations of ballistic flight and elastic impact. If the ball is in the air, its vertical motion is only influenced by gravity:

$$x_B(t) = -1/2gt^2 + \dot{x}_{B,0}t + x_{B,0} \qquad (1)$$

$x_B(t)$ is the vertical position of the ball, t is the time which is reset to zero at every impact, g denotes gravity, $x_{B,0}$ is the position of the ball at impact, $\dot{x}_{B,0}$ is the velocity of the ball immediately after impact. $\dot{x}_{B,0}$ is determined by the impact relation:

$$(\dot{x}_B^+ - \dot{x}_R) = -\alpha(\dot{x}_B^- - \dot{x}_R) \qquad (2)$$

$\dot{x}_B^+$ denotes the velocity of the ball after contact, equivalent to $\dot{x}_{B,0}$ in Eq.1; $\dot{x}_B^-$ is the velocity of the ball before the contact. α is the coefficient of restitution that captures the energy loss during the impact event. Due to the assumptions that the impact is instantaneous and that the racket has a much larger mass than the ball, racket velocity $\dot{x}_R$ remains the same before and after impact. For the experiment α was set to 0.6. The ball displacement was calculated and sampled at the same frequency as the racket displacement was sampled.

Procedure and Design. The participants were instructed to bounce the virtual ball rhythmically for the duration of one trial. They were instructed to bounce the ball so that the peaks of the ball trajectory reached the target line as accurately as possible. The target line was at a height of 0.8m measured from the lowest possible racket position. At the beginning of each trial, the ball appeared on the right side of the screen and rolled across a horizontal target line, which extended to the middle of the screen directly above the racket position. When the ball reached the end of the line, it dropped down onto the center of the racket (Figure 2). This starting procedure was designed to prepare the participants for each trial.

The experiment consisted of two parts, each consisting of two sessions of data collection. In Part I, participants completed 24 trials on the first day and 24 trials on a consecutive day. Each trial lasted 60 seconds. In Part II, participants similarly performed two sessions with 24 trials each on consecutive days, but now under perturbed conditions. To provide a transition the first session in Part II began with an additional 10 trials where participants performed with the same unperturbed condition as in Part I.

In the first session of Part I, participants were introduced to the experiment with a standardized instruction. To familiarize them with the set-up, they first practiced the task for 3 trials without any data recording. At the first session of Part II, participants were informed about the presence of perturbations and were encouraged to adapt as best as they could. No further information was given as to the nature of the perturbations.

Data Processing and Calculation of Dependent Measures. As the sampling frequency for the racket and ball displacements and racket acceleration was slightly variable depending on the performance of the computer, these data were re-sampled at a uniform frequency of 500 Hz using a spline interpolation algorithm. While racket and ball displacements did not require filtering, the re-sampled acceleration signals were filtered with a Savitzky-Golay filter (order 0, window size 0.01). All analyses were performed with Matlab (Version 7.0.1, The Mathworks).

Performance Measures. For every trial, the mean absolute error (*AE*) and the standard deviations of ball amplitude (*SDB*) were calculated. For each bounce, the absolute difference between the peak position of the ball trajectory and the height of the target line was calculated. The average of the absolute errors, *AE,* and its standard deviations, *SDB,* were calculated over the individual bounces of each trial. The first 8 seconds of each trial were not included in the analyses to eliminate possible start-up transients before the participant had established a stable bouncing pattern. As the entire trial included approximately 80-100 bounces, the performance measures were calculated over approximately 70-80 bounces.

Acceleration at Impact. In order to determine the racket acceleration directly prior to impact (*AC*) from the continuous accelerometer signal, the time of impact was defined as the moment when the brake started to pull on the rod. This moment was determined by the time when the computer sent the signal to trigger the break including a correction for the delay. This delay was measured separately and was about 6ms. The *AC* values were then determined from the accelerometer signal at one sample prior to impact. For every trial, the mean *AC* was calculated over the estimates for the 70-80 bounces per trial, similarly eliminating the first 8s of the trial.

Execution Variables at Impact. The three execution variables impact position (*IP* equivalent to $x_{B,0}$ in Eq. 1), velocity of the ball before impact (*BV*, equivalent to $\dot{x}_{B,0}$ in Eq. 1) and velocity of the racket at impact (*RV,* equivalent to $\dot{x}_R$ in Eq. 2) were registered by the data acquisition program directly when a racket-ball impact happened. (The notations for three execution variables were changed from now on for

clarity.) The mean values and standard deviations of the execution variables were calculated over all 70-80 bounces per trial.

Variability Components. Based on the execution variables at each bounce, the TNC analyses were performed as described in Appendix A. One trial was treated as a data set consisting of approximately 70-80 data points, i.e., bounces. The calculations of *T*, *N*, and *C* were conducted as follows: In Part I, the first trial from the first session was taken as a reference trial. Each subsequent trial was then compared with this first trial such that the changes and contributions of *T* and *N* were calculated in a pairwise comparison manner. Hence, the values of *T* and *N* for subsequent comparisons reflected the changes over practice in a cumulative fashion. *C* was calculated by comparing the data set from each trial to the shuffled data set. In Part II, the last trial of the 10 unperturbed trials served as the reference trial. Except this change, all the other calculation steps were the same as in Part I. For the calculation of covariation in the respective data pairs, the data was permuted 50 times. A detailed explanation is provided in Appendix A and in Müller and Sternad (2003).

Distortion of Execution Space. The objective of Part II was to evaluate how participants adapted to changes of the topology of the solution manifold. Given that participants had established their individual solution strategies in Part I, a distortion of the manifold was designed in an individualized way to afford comparisons across participants. The individual's strategy as established in Part I was assessed from the first 10 trials under normal conditions in Part II. Before this procedure was applied, we verified that the means and distributions of the first 10 trials were not significantly different from the last 10 trials in Part I. The individual's performance corresponded to a data distribution in execution space, where the means and dispersions of the execution variables characterized an individual's strategy: means $\overline{IP}, \overline{BV}, \overline{RV}$ and their respective standard deviations *SDIP*, *SDBV*, *SDRV*. Note that the variables in general had a normal distribution such that the mean was an appropriate statistic. Figure 4 shows the execution space with the large ellipsoid representing the 3D distribution of the pre-distortion trial bounded by two standard deviations in the

directions of the three execution variable. The small ellipsoid represents the distribution of data at the end of the session.

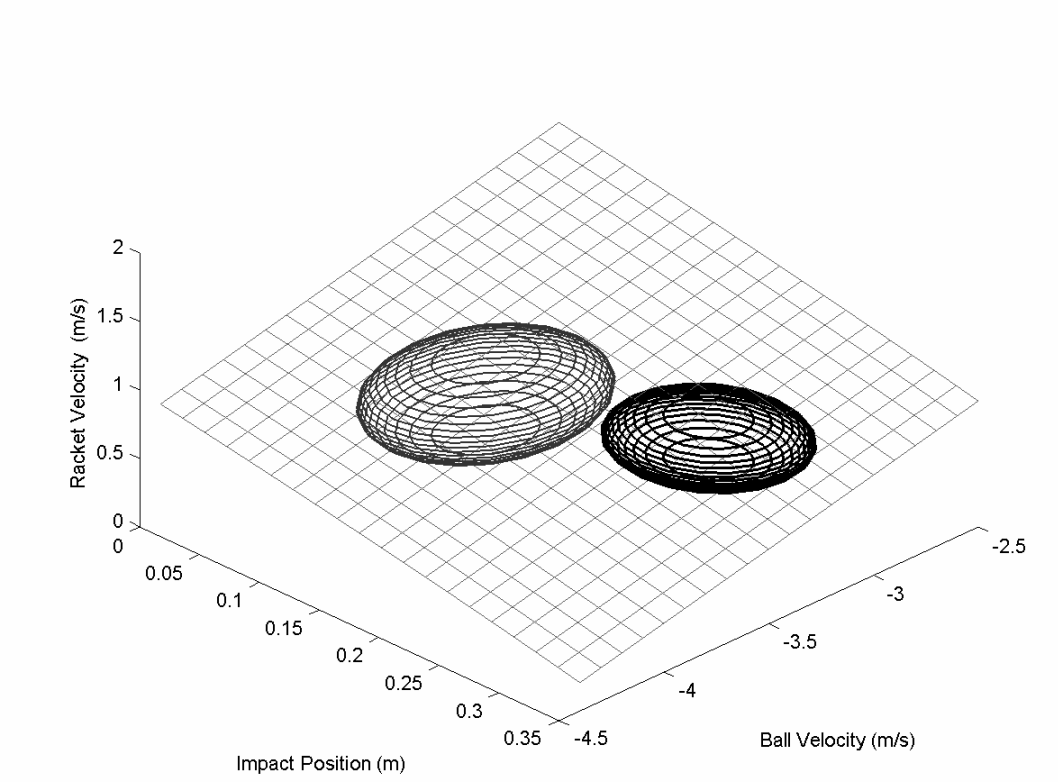

Figure 4: The dispersion of three execution variables in the execution space. The center of the ellipsoid is the mean performance ($\overline{IP}, \overline{BV}, \overline{RV}$), the radii are the 2 times standard deviation of each variable (2**SDIP*, 2**SDBV*, 2**SDRV*). The bigger ellipsoid is from the 10th trial under normal condition in Part II, the smaller ellipsoid is from the last trial in Part II.

These parameters provided the baseline to calculate the distortion of the solution manifold for each individual participant. A perturbation was applied when a bounce was within the 3D ellipsoidal distribution determined from the 10 pre-distortion trials. The magnitude of the perturbation was maximum when the bounce occurred at exactly the mean preferred solution ($\overline{IP}, \overline{BV}, \overline{RV}$), the center of the ellipsoid. The perturbation magnitude linearly decreased with increasing distance from the preferred

solution and became zero at the surface of the ellipsoid. The perturbation consisted of adding some value to the ball velocity BV at release of that particular bounce i, leading to a deviation from the expected and physically correct ball amplitude (see Eqs. 1 and 2). Thereby, the originally preferred solution in execution space was no longer successful. Further, each bounce i could be now characterized by its geometric distance D_i to the preferred solution ($\overline{IP}$, $\overline{BV}$, $\overline{RV}$). For the calculation of the perturbation magnitude the distances D_i from the preferred solution had to be normalized because the three variables had different units and magnitudes. D_i was computed as the Cartesian sum over the three normalized distances in each dimension, D_{IP_i}, D_{BV_i}, and D_{RV_i}:

$$
\begin{aligned}
D_{IP_i} &= (IP_i - \overline{IP})/(2 * SDIP) \\
D_{BV_i} &= (BV_i - \overline{BV})/(2 * SDBV) \\
D_{RV_i} &= (RV_i - \overline{RV})/(2 * SDRV) \\
D_i &= \sqrt{D_{IP_i}^2 + D_{BV_i}^2 + D_{RV_i}^2}
\end{aligned}
\qquad (3)
$$

Hence, for each bounce D_i was computed. If the data point was inside the ellipsoid, i.e., D_i was between 0 and 1, a perturbation was applied. The magnitude of the perturbation was such that for D_i = 0, the perturbation was maximum, corresponding to a deviation in ball amplitude of 0.15m. For $D_i \geq 1$, perturbations were zero. For $0 \leq D_i \leq 1$ the perturbation scaled linearly producing amplitude deviations between 0 to 0.15m. Note that D_i was always positive. However, if the original bounce would have undershot the target, the applied perturbation was assigned a negative direction, amplifying the undershooting even more. Conversely, if the ball overshot the target, perturbations were assigned a positive direction. This sign convention avoided a canceling of ball amplitude errors and reduced the possibility of

confusing feedback to the participants.

These perturbations led to an effective distortion of the solution manifold, as shown in Figure 5 for a typical subject. Figure 5A displays a 2D slice of the undistorted manifold; Figure 5B shows a 2D slice through the distorted manifold, illustrating the oblique orientation of the ellipsoid calculated for an individual's performance. The 2D slice was selected at $\overline{IP}$, the mean performance for that subject. This corresponded to slice through the center of the ellipsoid. Grey shades denote absolute errors of the ball amplitude with white denoting zero error. Note that while zero error solutions disappear within the ellipsoid, a hint of slightly better solutions close to the original manifold remain, visible as a thin light grey line. This is due to the fact that the perturbations that modified the ball amplitude are not proportional to the original ball amplitude.

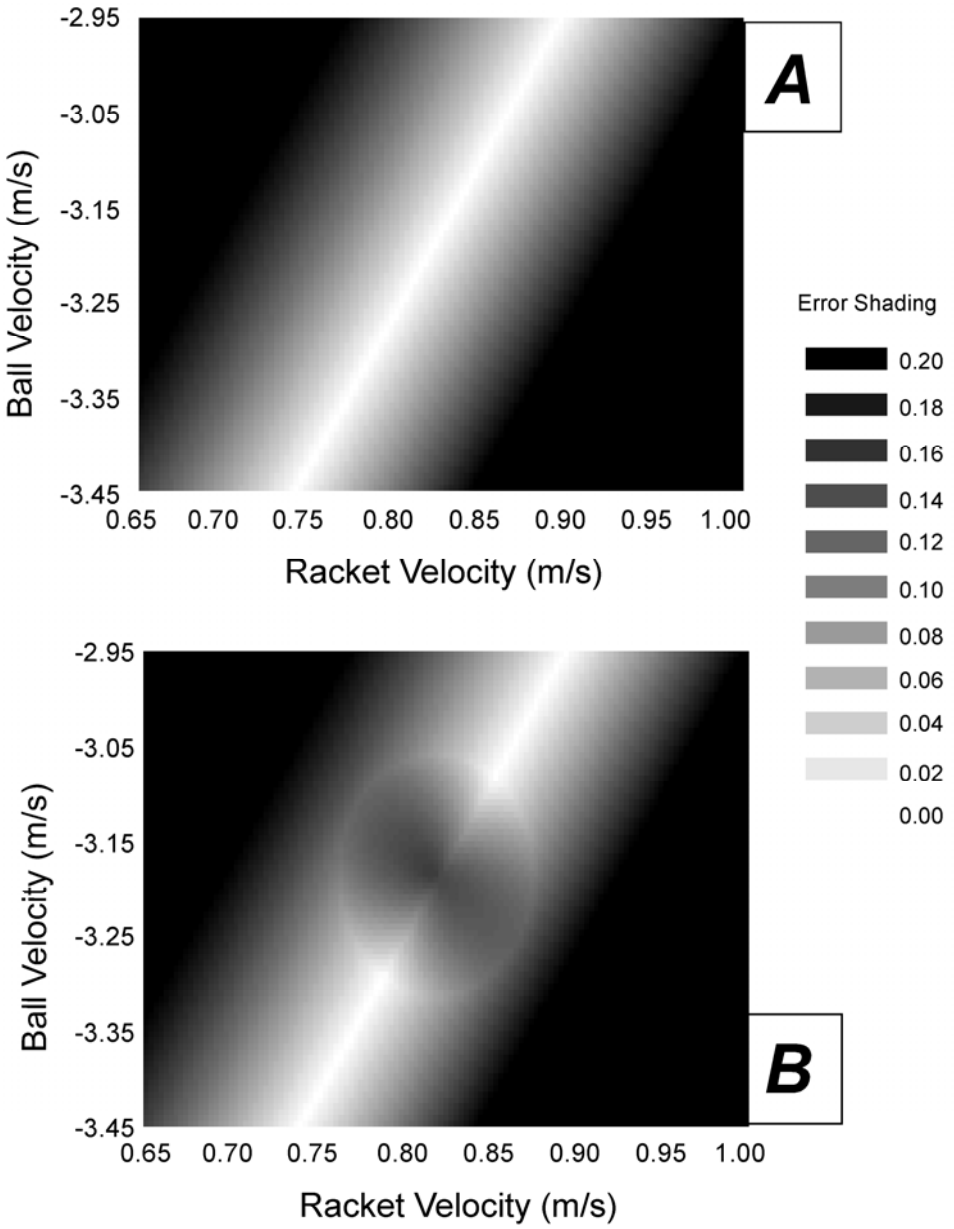

Figure 5: Solution manifold in ball bouncing before distortion (**A**) and after distortion (**B**) It shows all possible solutions (ball velocity and racket velocity combinations) for impact position at 0.27m. The associated *AE*s with the solutions are shown in different grey shadings.

Curve Fitting. The time course of dependent measures *AE, SDB* and *AC* over the trials in Parts I and II were fitted by an exponential regression equation:

$$y(t) = ae^{-bt} + c$$

where y is the dependent measure, t denotes the independent measure which was trial

number in the present case, and *a*, *b* and *c* were fitting parameters. The nonlinear Levenberg-Marquardt regressions were performed by using the curve fitting toolbox in Matlab 7.01.

Elimination of Trials. In Part I inspection of the raw data made it necessary to discard one participant's (P5) accelerometer data from the analyses because it showed too much noise. This was traced to a problem with the accelerometer, which was subsequently repaired. As the signal of the encoder was unaffected and all other measures based on its position readings could be determined they were included in the analyses. In Part II of the experiment the data of Participant 4 had to be eliminated from the analyses, because of an error in the calculation of the distortion of the solution manifold. Additionally, some selected trials had to be eliminated as participants failed to perform a regular pattern. The criterion for elimination was that the ball amplitude was two standard deviations below the mean amplitude for three consecutive bounces. In Part I four such trials, in Part II six such trials were detected. When analyzing the time course of the trial means within participants these missing data points were filled in by taking the average of the two adjacent data point.

Results

Part I:

Task Performance. For a first evaluation of task performance, the absolute error of the ball to the target *AE* was examined. Figure 6A depicts the group averages over the course of the 48 trials together with an exponential fit as shown by the solid line. For each trial the *AE* means of all 8 participants were averaged. The overall R^2 of the exponential fit was 0.92. The exponential decrease in *AE* indicated that participants indeed became more accurate at bouncing the ball to the target height. The average error decreased from 0.08m to 0.035m. Figure 6B shows the individual exponential fits for each of the 8 participants' data, showing an exponential decrease in error for all but one participant. The data points of each trial are not shown in the figure for

clarity of presentation. The R^2-scores for the 8 regression fits ranged between 0.60 and 0.88. The dashed line corresponds to the participant with the lowest R^2-value of 0.60. This participant (P8) had a lot of experience in recreational table tennis, which might explain his relatively good performance level already at the beginning.

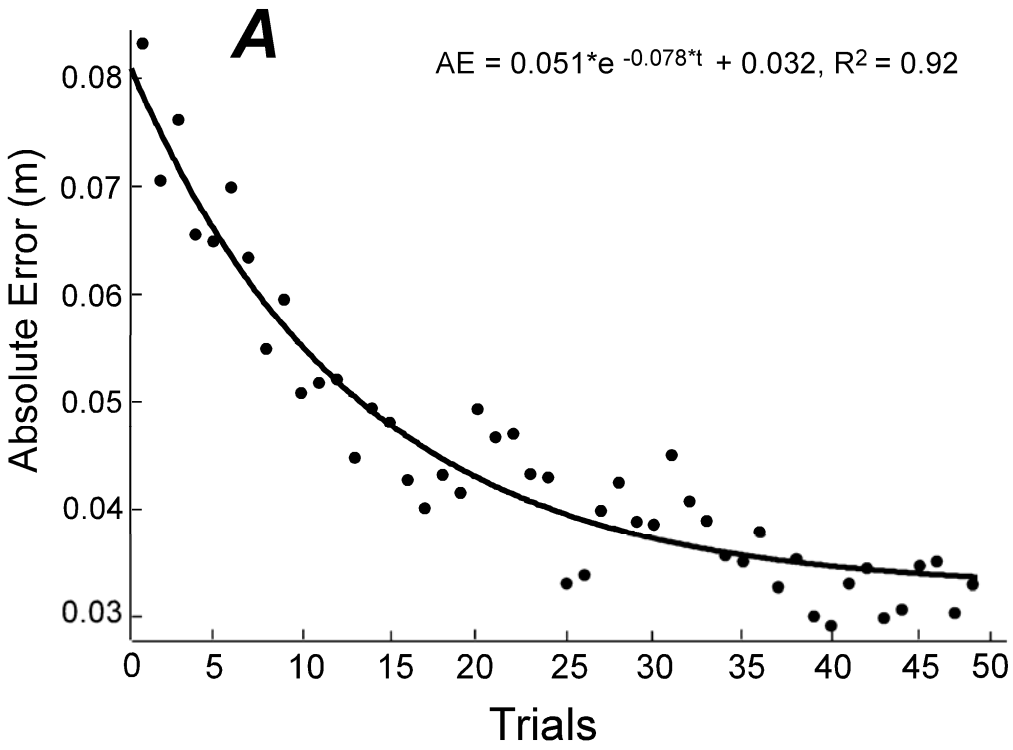

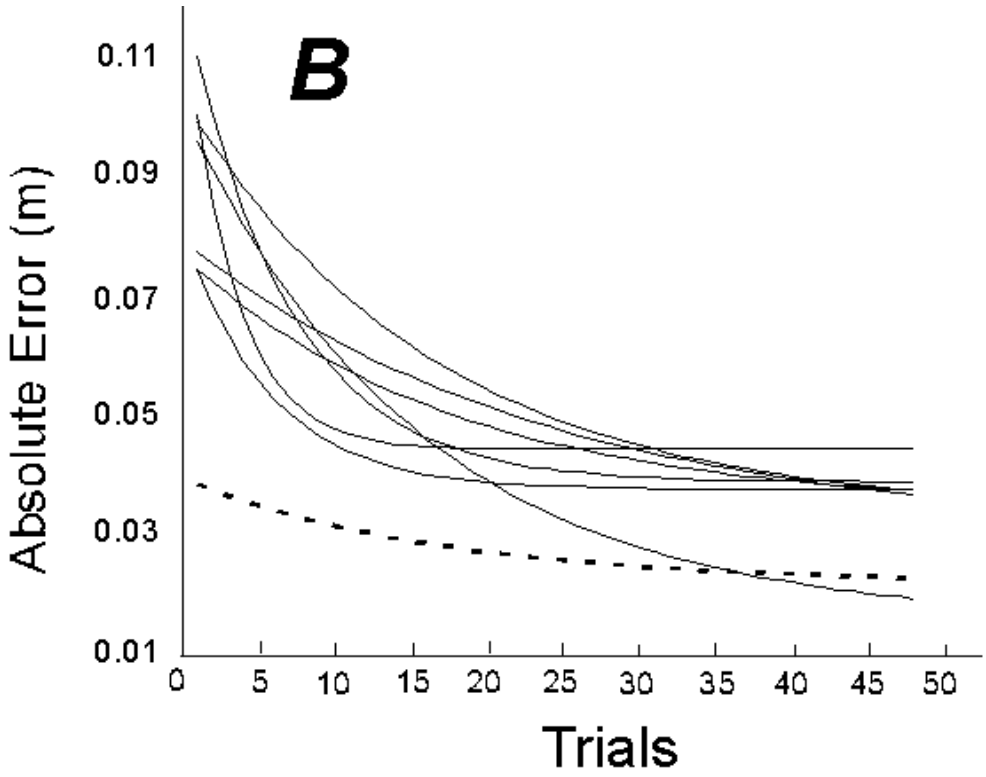

Figure 6: Absolute error of ball height across the series of 48 trials of Part I of the experiment. **A:** Group averages calculated for each trial number. The solid line represents the exponential fit. The fitted exponential function is shown with r-square value. **B:** Exponential fits for individual participants. Each line represents one participant. Data points for individual trials were omitted for clarity.

Figure 6A shows that the participants still had an absolute error of 0.035m even after 48 trials, indicating that the participants remained to either under- or overshot the target. Therefore, it was also useful to examine variability in terms of variable error or standard deviations of ball amplitude, *SDB*. The same exponential regressions were performed on both the group's and the individual participants' data. A comparison of the fit parameters of the exponential regressions of the *AE* and the *SDB* revealed that the rate of decrease was very similar. The R^2-values for *SDB* were slightly smaller. Table 1 summarizes the parameters of the individual fits and their R^2-values.

Table 1:

Parameters of the exponential fits of the individual participants' and group average data of absolute error of ball height (*AE*), standard deviations of ball height (*SDB*) and accelerations at impact (*AC*) in Part I of the experiment. For Participant 5 there were no acceleration data available.

Participants	AE				SDB				AC			
	a	b	c	R^2	a	b	c	R^2	a	b	c	R^2
P1	0.068	0.058	0.032	0.65	0.081	0.053	0.041	0.58	202.8	0.0007	-197.70	0.68
P2	0.096	0.309	0.036	0.65	0.134	0.319	0.044	0.58	15.65	0.017	-14.680	0.79
P3	0.070	0.183	0.041	0.65	0.076	0.175	0.536	0.57	9.721	0.035	-6.332	0.72
P4	0.085	0.063	0.015	0.88	0.107	0.07	0.022	0.86	6.402	0.039	-0.896	0.77
P5	0.081	0.145	0.038	0.60	0.099	0.197	0.046	0.57	–	–	–	–
P6	0.096	0.189	0.042	0.61	0.107	0.168	0.051	0.61	5.836	0.035	-3.176	0.64
P7	0.075	0.115	0.039	0.74	0.077	0.108	0.049	0.58	4.644	0.179	-3.203	0.62
P8	0.018	0.051	0.021	0.60	0.026	0.024	0.0182	0.57	5.109	0.144	-1.941	0.63
Average	0.051	0.078	0.032	0.92	0.058	0.071	0.039	0.91	6.306	0.048	-3.395	0.93

Dynamical Stability. To investigate the change in performance stability, the racket accelerations immediately before impact, *AC*, were analyzed. Figure 7A shows the group averages over 48 trials. *AC* decreased exponentially over practice ($R^2 = 0.93$) from initially positive to negative values. The dashed horizontal line at zero highlights this change. The individual exponential fits in Figure 7B show that each participant had a tendency to decrease their *AC* values from positive to negative throughout the 48 trials. The R^2-values for the individual fits ranged between 0.62 and 0.79. The data points are not shown for clarity of presentation. Table 1 lists the fit parameters of both

the individual participants and their group average. On average the *AC* values turned negative from trial 13 onwards. Note that the change in the initial part of the trial sequence is more rapid in the *AC* values than in the *AE* and *SDB* values. Further, the *AE* and *SDB* values tended to asymptote, whereas the *AC* values did not seem to have reached a plateau after 48 trials.

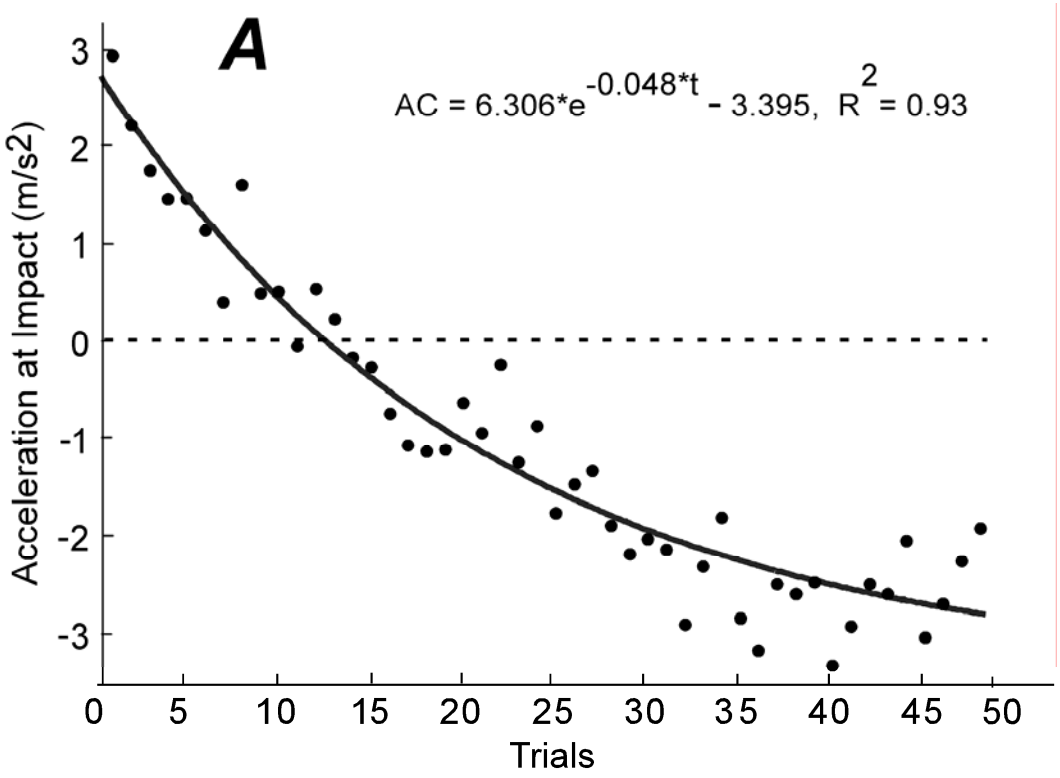

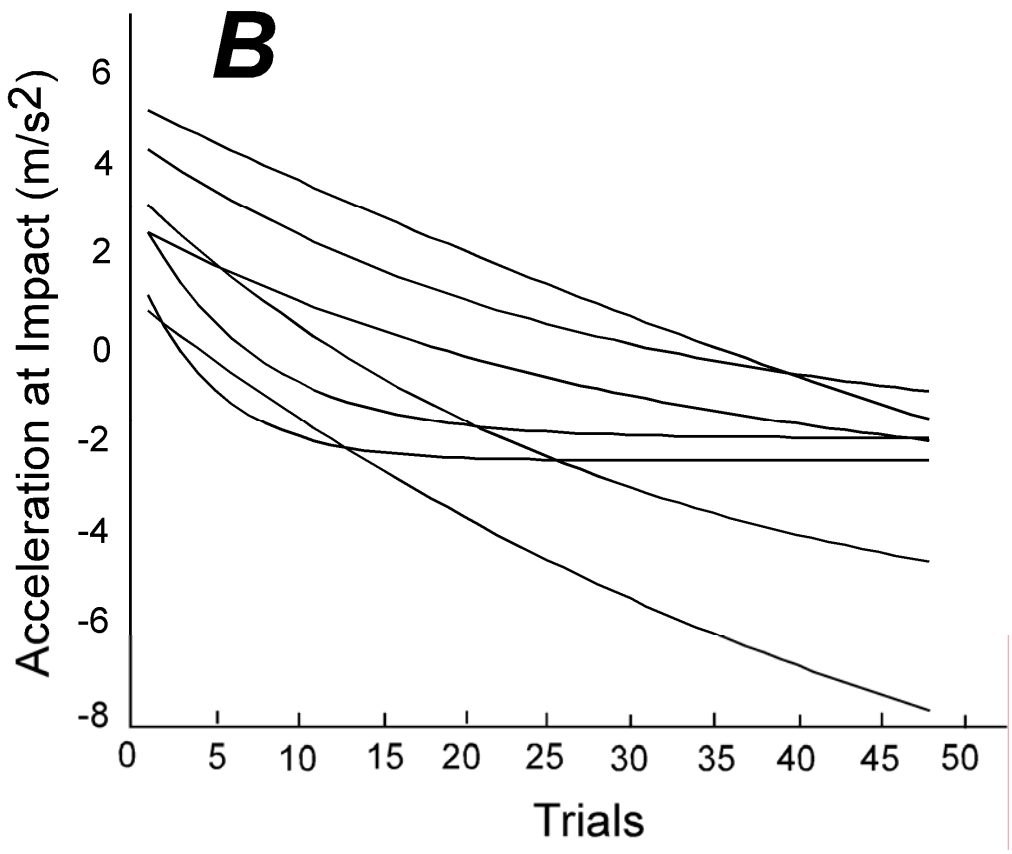

Figure 7: Racket accelerations at impact *AC* across the series of 48 trials. **A:** Group average of racket accelerations *AC*. Each point represents the average of 8 participants in a trial. The solid line represents the exponential regression fit. The dashed horizontal line at $AC = 0$ is plotted to highlight the transition from positive to negative values. **B:** Exponential regressions for individual participants.

Variability Decomposition. Figure 8 illustrates the contributions to performance improvement by the three components, tolerance *T*, noise reduction *N*, and covariation *C*, plotted over the 48 trials. The change in *AE* from trial 1 to trial *n* ($n = 2, 3, \ldots 48$) was parsed into the three components, which are indicated by the different shadings of the bars in Figure 7. The decomposition was performed for each individual participant for each of the 48 trials. Subsequently, each of the three components were averaged over 8 participants and plotted per trial. It can be seen that the improvement in performance was mainly brought about by *N*, which continued to increase until the last trial. There was also a small positive contribution of *T* and *C,* which was

relatively invariant throughout practice.

The same decomposition was also performed on the second performance measure *SDB*. The results were very similar. This is partly to be expected as the changes of *SDB* and *AE* with practice had a very similar time course.

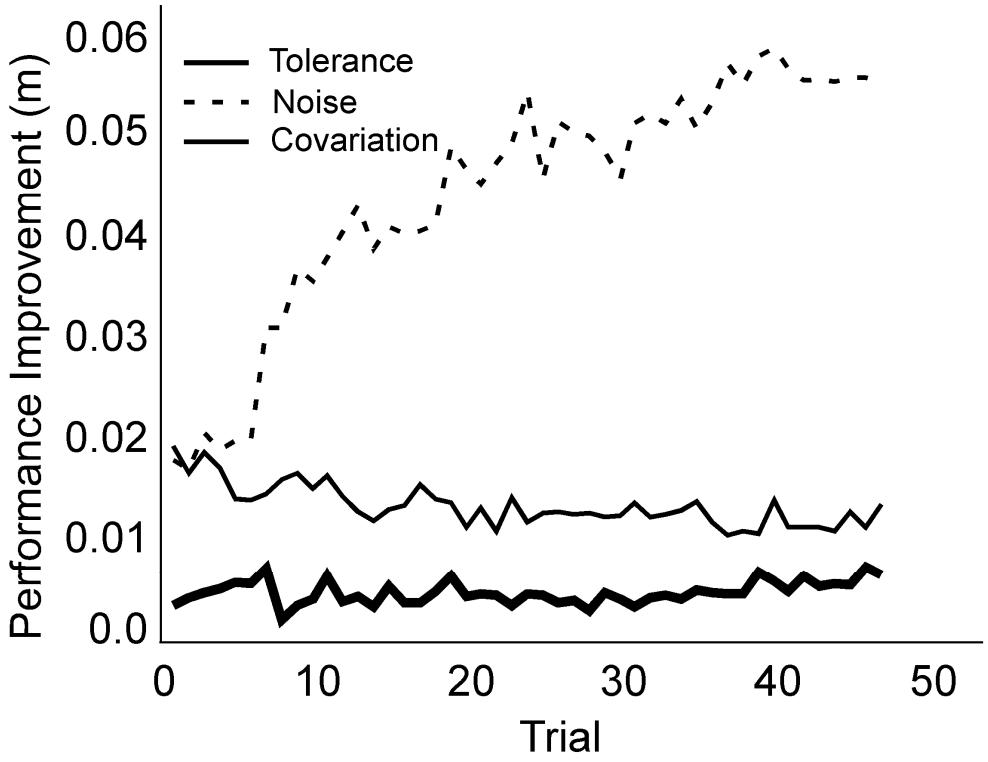

Figure 8: Contribution of the three components *T*, *N* and *C* to performance improvement. Each line represents the group averages per trial.

Part II:

In Part I, all 8 participants completed 48 trials in which the measured performance indicators reached a plateau, suggesting that they learnt the task. In Part II of the experiment, participants were exposed to distortions of execution space to study adaptation as the search for new solutions. At the beginning of Part II participants were informed about the upcoming distortion but not about the nature of the distortion.

They were told that they could and should adapt and improve their performance as much as possible. At the end of Part II participants were asked whether they understood how their performance became perturbed and how they thought that they had changed their strategy. Except for one participant (P3), nobody was able to explicitly tell how they had adapted to the perturbations. P3 remarked that he noticed that changes of his impact position had an effect on the perturbations.

Task Performance. The values of *AE* averaged over 7 participants are shown in Figure 9A. The first 5 data points represent the 5 trials before the distortion was implemented; the following 48 trials depict the performance under distorted conditions. Before distortion onset, *AE* was on average 0.03m; at the first distorted trial, *AE* increased significantly to 0.075m. However, shortly after the distortion onset *AE* decreased again and from trial 35 onwards, *AE* had returned to pre-distortion levels. Figure 9B shows the exponential regressions of all 7 individuals.

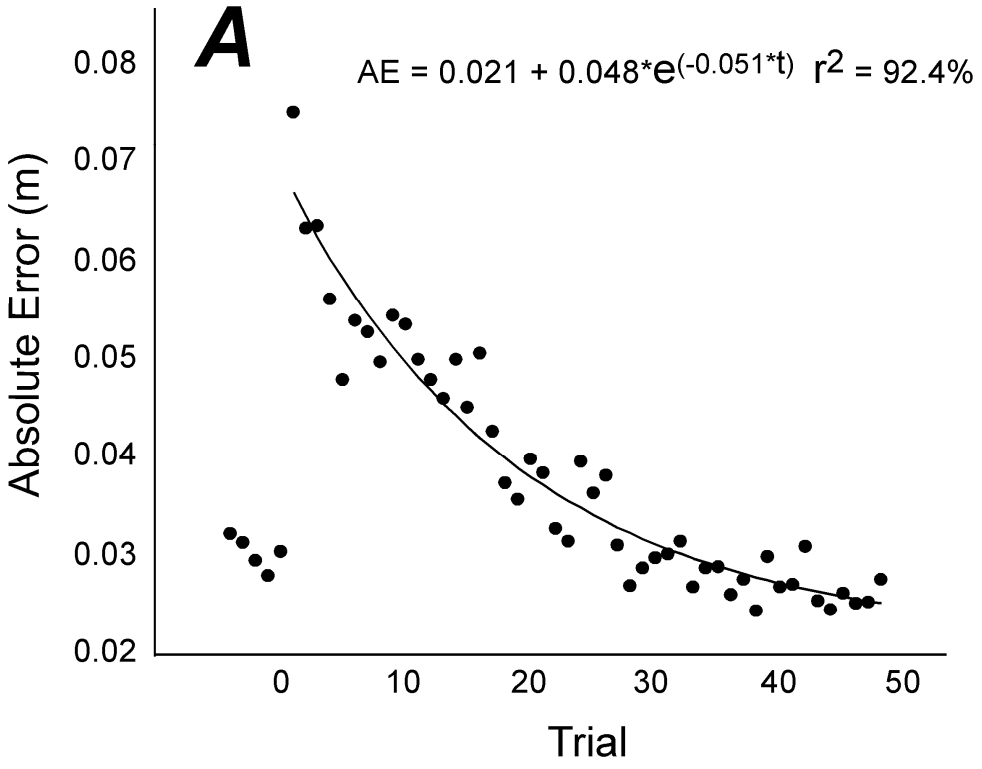

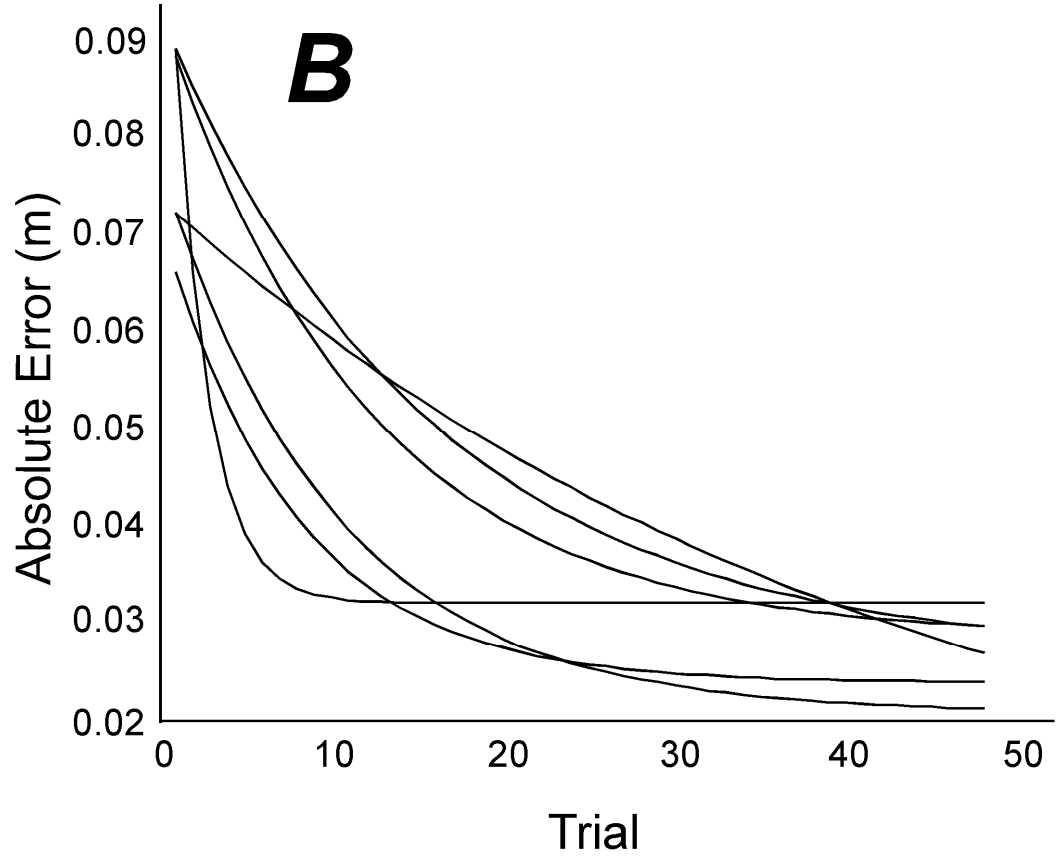

Figure 9: Absolute error of ball height across the series of 48 trials of Part II of the experiment. **A:** Group averages calculated for each trial number. The solid line represents the exponential fit. The fitted exponential function is shown with r-square value. **B:** Exponential fits for individual participants. Each line represents one participant. Data points for individual trials were omitted for clarity.

To further test whether participants returned to pre-distortion levels at the end of the session, pairwise *t*-tests compared the 5 trials that directly preceded distortion onset and the 5 last trials in this session. With the exception of P8, all participants performed better at the end of the session, decreasing their absolute error by .002m down to .013m. However, in only three participants these comparisons were significant (P1, P2, P3). Hence, following the adaptation to the distortion, participants still improved their performance.

Racket Accelerations at Impact. The change of racket accelerations at impact *AC* over trials before and during the distortion are graphed in Figure 10. Four of the seven

participants showed similar exponential patterns and their average is displayed in Figure 10A: prior to distortion the impacts were performed with negative values similar to the ones seen at the end of the preceding sessions. With distortion onset *AC* became significantly more negative but then gradually increased towards the pre-distortion values. Interestingly, the negative *AC* values were re-established from more negative values, opposite to the change seen across learning in Part I where participants began with positive values. The exponential regression performed over the distorted trials had a R^2-value of 87.8%. Figure 10B highlights this pattern of return by the respective individual regression fits for those 4 participants. The remaining three participants showed less systematic patterns and the R^2 values of their regressions were very low. Their individual data are plotted in Figure 11C. Although their *AC* trial data showed considerable scatter without a clear pattern, overall the data still show convergence towards a negative *AC* value.

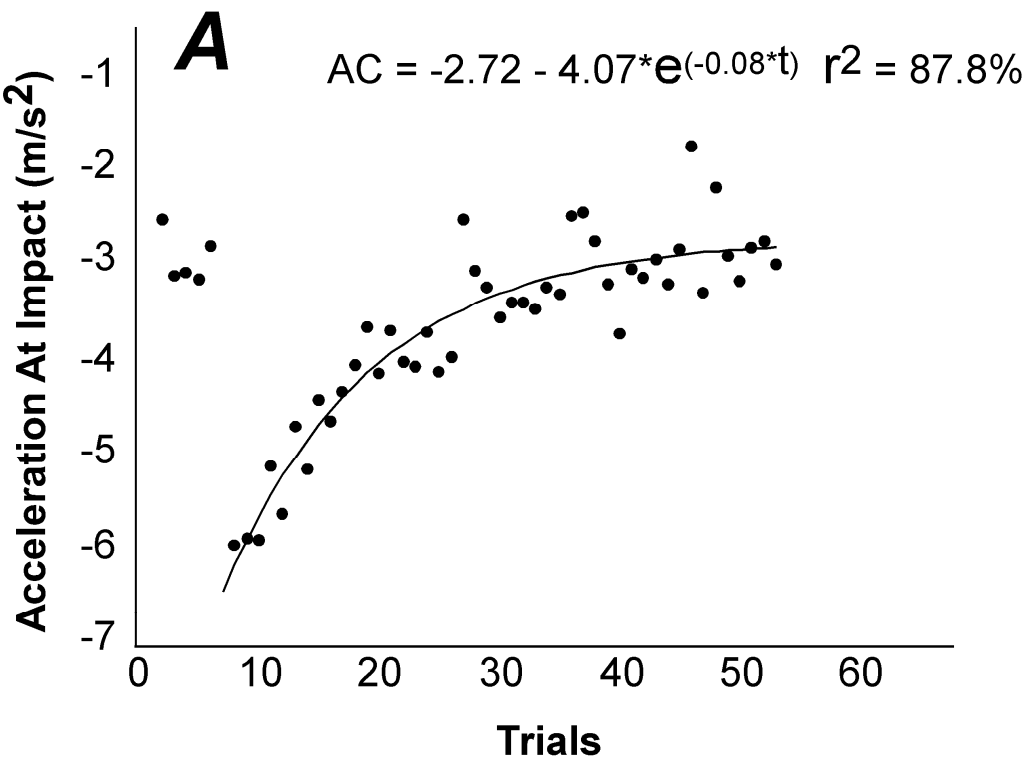
A
AC = -2.72 - 4.07*e(-0.08*t) r2 = 87.8%
Acceleration At Impact (m/s2)
-1
-2
-3
-4
-5
-6
-7
0
10
20
30
40
50
60
Trials

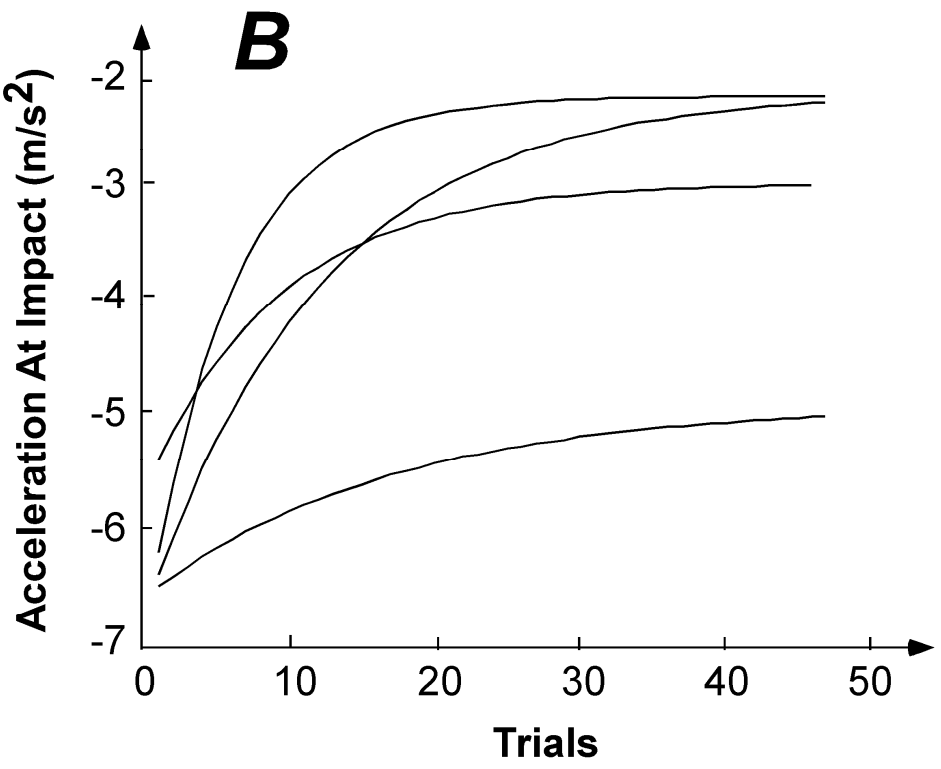
B
Acceleration At Impact (m/s2)
-2
-3
-4
-5
-6
-7
0
10
20
30
40
50
Trials

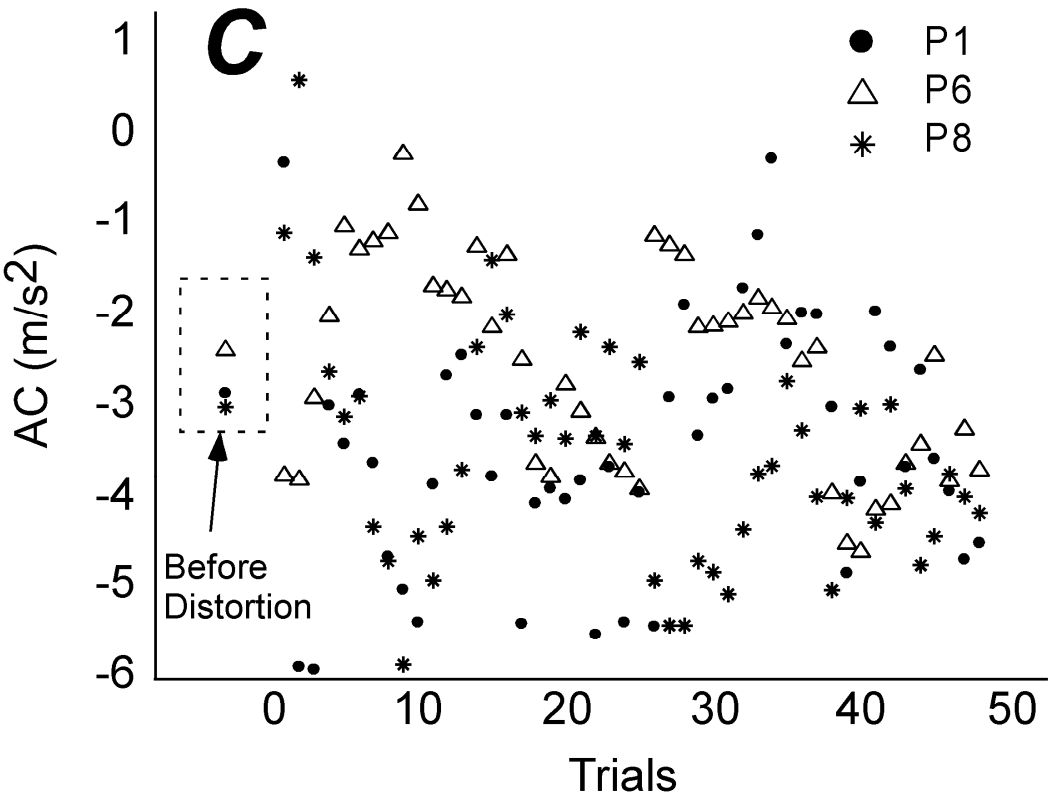

Figure 10: Racket accelerations at impact *AC* across 48 trials under distortion condition in Part II. **A:** Group average of racket accelerations *AC*. Each point represents the average of 4 selected participants with good exponential fitting. The solid line represents the exponential regression fit. **B:** Exponential regressions for the selected 4 individual participants. **C**: The data for the three participants that were not included in the group average due to their lack of an exponential return pattern.

Execution space Analysis. For a first inspection of how participants adapted to the distortions, their data were plotted in execution space. Figure 11 shows the execution space with the planar solution manifold and the distorted ellipsoid of P2. The black dots show the means of execution variables of all 48 trials with distortion. As can be seen, performance was changed away from the ellipsoid to another more successful location on the solution manifold. Interestingly, the change in location appears relatively large. With a view to Figure 4, which also shows the distribution of execution variables of the last trial, it can be seen that the relocation takes the variance

of the trial into account such that there is no overlap of the scattered data with the distorted ellipsoid. It is further striking that the mean data move along the planar surface rather than a random search in all three dimensions. This may be due to the fact that there was some local "ghost" of the solution manifold left (see Figure 5).

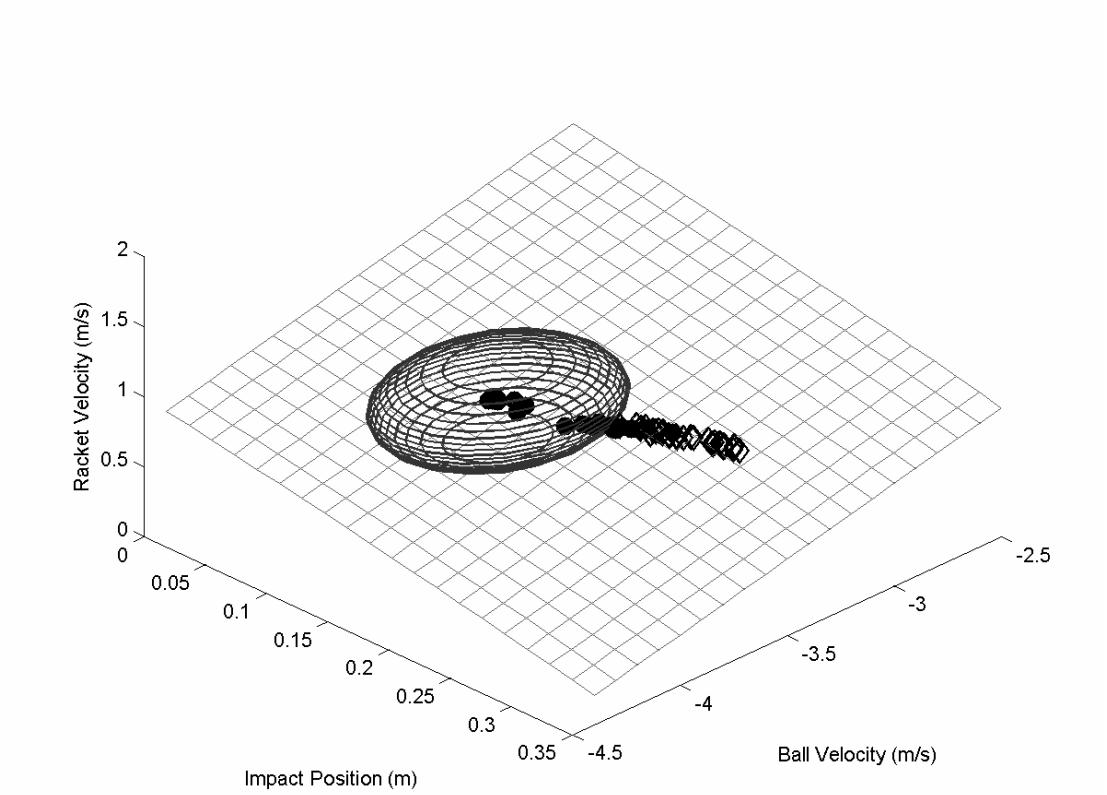

Figure 11: The perturbation ellipsoid (ellipsoid) and the mean performance from individual trials (diamond markers) in Part II from P2.

To test whether participants successfully moved out of the distorted area in execution space, the average D was calculated for each trial and then averaged across participants. Figure 12 shows this increase in D, the error bars indicate standard deviation across subjects. Note that once $D > 1$, the triplet of execution variables have moved outside the distorted ellipsoid. While participants are very fast in leaving the location of maximum perturbation, it takes until trial 20 where also the variance is no longer within the distorted ellipsoid. Note that D did not increase further from trial 20

onwards. This indicated that for most of the bounces, the execution was just outside of the distortion ellipsoid.

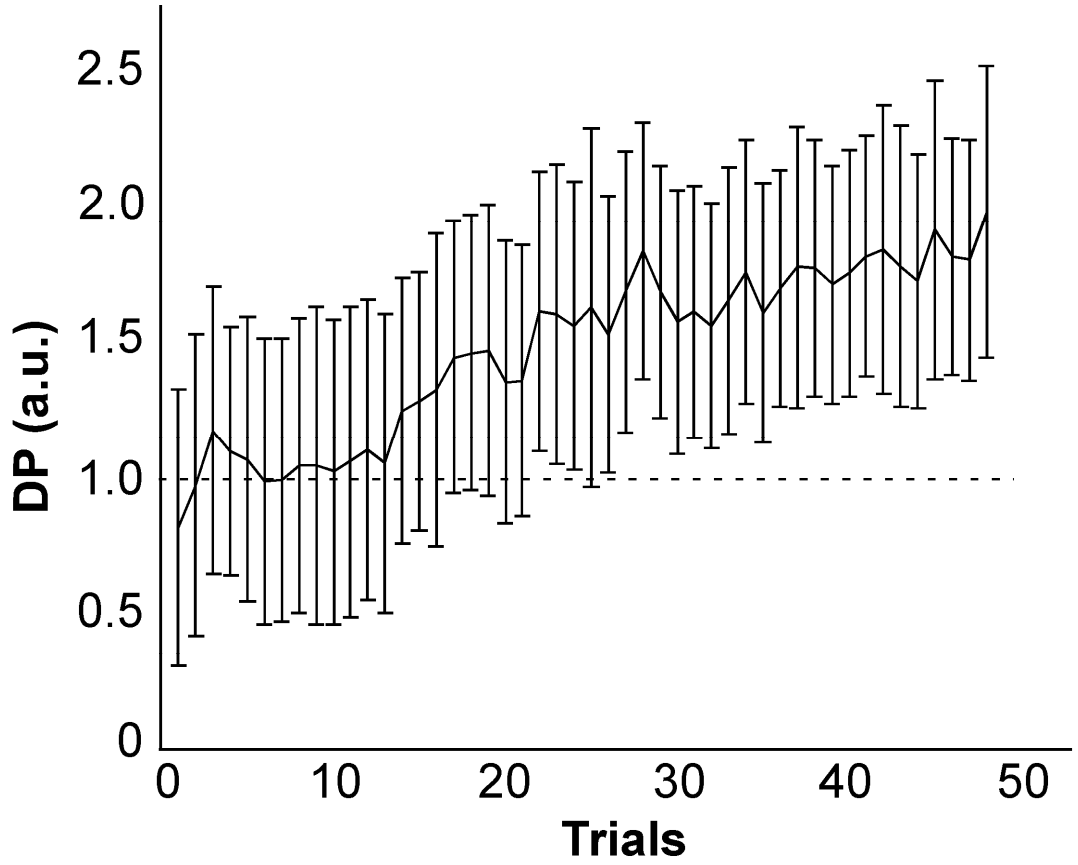

Figure 12: The average *DP* across participants in Part II. The error bar stands for the average standard deviation for each trial.

Variability Decomposition. To apply the TNC-decomposition the first trial of the distorted condition was compared to all the other trials in the distortion condition. Figure 13 presented the results for all participants. It was evident that *T* contributed most to the performance improvement. From trial 20 onwards the contribution of *T* plateaued at a constant level. *N* markedly increased in contribution from trial 13 onwards and then remained relatively constant until trial 48. *C* had relatively large contribution at the beginning and turned smaller from trial 14 onwards. The major difference between Part I and Part II was the use of tolerance, i.e., exploration of execution space.

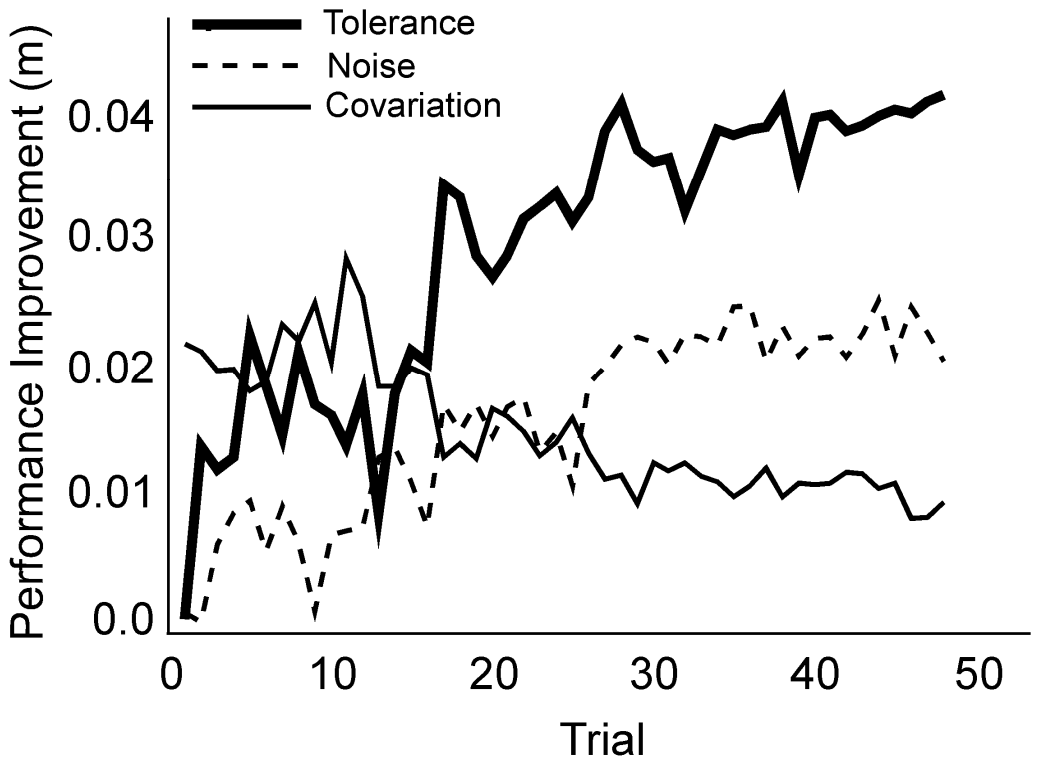

Figure 13: Contribution of the three components *T*, *N* and *C* to performance improvement in Part II. Each line represents the group averages per trial.

Discussion

Ball bouncing task has been investigated as an exemplary task to study how humans exploit the properties of the task and adopt coordination patterns under task constraints. Stability analysis of a simple mechanical model for the task has revealed that there exist a set of dynamic solutions when the racket hit the ball with negative acceleration. Using this stable solution effectively makes the bouncing system operate round a period-1 attractor, thus small fluctuations in the performance die out without the necessity to make corrective movement of the racket. This stability property is termed passive stability, as it requires no active control and comes from a pure mechanical system deprived of adaptive neural control. Humans were found to exploit this passive stability and hit the ball with negative racket acceleration (Schaal et al., 1996; Sternad et al., 2001; Dijkstra et al., 2004). However, this strategy was not an

intuitive solution for the actor, as a previous learning study found out that novice actors first hit the ball with positive racket acceleration, then after practice they gradually changed to use negative impact acceleration (Dijkstra et al., 2004). This learning process was accompanied by performance improvements. How the actor discovered the solutions of using passive stability through practice is still a question. To unravel the learning process, we adopted the TCN method to study the changes in performance variability over practice.

The TNC method partitions the variability improvement in performance into three components: tolerance, noise and covariation. Tolerance captures the improvement gained from locating the data set of execution variables to a different place in the execution space where the area of successful solutions is larger and more tolerant to noise/fluctuations in execution. In ball bouncing task we found out the solution manifold is more or less the same (approximately homogeneous) at different locations in the execution space, thus the relocation of solutions will not help the performance improvement too much. In Part II of the present study, a distortion of the solution manifold was introduced to force the actor move away from their preferred region in the execution space. Hence the relative contribution of tolerance components can be studied during the adaptation. Noise captures the improvement gained from reducing the variability of all possible execution variables. While a decreased level of noise has been universally recognized as a result of learning, it is usually assessed in term of variability reduction of the movement result, for example, the variability of ball amplitude in the ball bouncing task. The TNC method recognizes that noise reduction is only one of the three components in the performance improvement and it comes from the stochastic nature of all execution variables. It is calculated after the other two components are subtracted from the total performance improvement. Covariation captures the improvement from covarying the execution variables according to the relationship between the execution variables and the movement outcome. Geometrically, it can be shown as alignment of the data set along with the solution manifold in the execution space. It was recognized as a

signature of control in other studies on variability (Latash, Scholz, Danion, & Schöner, 2002; Shöner & Scholz, 1999). In the TNC method, covariation is calculated by first computing within-trial covariation, which is the difference in performance between the original data set and the permutated data set, then computing the improvement in within-trial covariations across trials.

It was found all participants showed improvement in performance with practice under normal condition, and the improvement was accompanied by a change in impact acceleration from positive values to negative values. This result was consistent to previous findings (Dijkstra et al., 2004). However, the previous study only had limited number of subjects to show the simultaneous change in variability measure and in impact acceleration. The present study confirmed that utilizing passive stability is an acquired strategy and the process of tuning to stable regime may assist the reduction in variability of overt performance measure. Note this result was from the present virtual reality setup. Its consistency to the result from real ball bouncing provide a support for validity to study ball bouncing task in virtual reality (see also de Rugy, Wei, Müller, & Sternad, 2003).

In Part 1, TNC method revealed that tolerance component was relative small and its contribution was almost the same throughout the learning. This result is expected as the solution manifold for the ball bouncing task under normal condition is homogeneous in the execution space, i.e., different locations in the execution space offer equivalent tolerance to noise. Thus even actors' solutions (the data cloud of execution variables from one trial) migrate in the execution space, as long as average solution for that trial is on the solution manifold, tolerance would be the same. From this aspect, the finding of a nearly constant tolerance showed actors learned the solution manifold in the early phase of practice and stayed around the manifold (instead of moving away from it) during the learning process. It appeared that the solution manifold, as a reflection of physics of the ball bouncing task, is easy to comprehend by the actor. Noise reduction contributed the most to the performance improvement. Geometrically, noise reduction can be visualized by a reduction in

scatter of execution variables in the execution space. Though participants made their average performance (the center of the scatter) stay on the solution manifold, they further improved their performance by reducing the scatter, as evidenced by a continuing improvement in noise reduction. Covariation was almost zero at the beginning then remained small throughout the practice session. This may related to the relatively flat solution manifold for ball bouncing task, as compared to complicated nonlinear solution manifold as in skittle throwing task, for example (Müller & Sternad, 2004). As a result, the control aspect that is reflected by covariation, i.e., how to combine the three execution variables, is easy for the ball bouncing task. Covariation did not reduce as a result of reliance on passive stability.

Passive stability was exploited such that small fluctuations could die out without active corrections in the racket movement. How much did this strategy contribute to the overall performance improvement as compared to the noise reduction? We can not make a conclusive remark on this matter given current results. The impact acceleration decreased to negative values from about 10^{th} trials onwards, however, noise component continued to improve until the end of Part I. There is no casual relationship between adopting negative acceleration and reducing noise in movement execution. Both of these two aspects contributed to the performance improvement.

Upon the introduction of distortion, the performance error was significantly increased. Trial by trial, participants adapted to the distorted solution manifold and improved their performance to pre-distortion level. Racket impact acceleration showed different changes for individual participants. For 4 participants, impact acceleration became more negative in the early phase of adaptation, and then gradually returned to pre-distortion level. Note this change was in the opposite direction as the one in Part I when the impact acceleration changed from positive values to more negative values. Thus, there appeared to have a preferred range of impact acceleration for individual participants. For the remaining three participants, impact acceleration did not show a clear changing pattern, instead, it scattered more within the negative range. These results indicated passive stability was still exploited

in the process of adaptation. No matter whether the distortion brought an abrupt change in impact acceleration or it was kept negative within a bounded range, hitting the ball with negative acceleration was always guaranteed. This suggested exploitation of passive stability was still the preferred solution for the ball bouncing task with distorted solution manifold.

Consistent to the prediction, tolerance contributed most to the performance improvement during adaptation. While not unintuitive, the analysis highlighted this conception of exploration well acknowledged in the motor control literature. It increased to reach some plateau at approximately 20^{th} trial. The inspection of execution variables in execution space found out that at the 20^{th} trial, the average solution moved out of the distorted area (Figure 1). This showed that the participant found his/her new solution in the execution space that was distinct from the original preferred location. This relocation in execution space was a continuous process on the solution manifold, not a random search in any direction possible (Figure 11). Noise reduction was relatively small when the distortion was introduced. After new solution was established in the execution space at approximately 20^{th} trial, noise reduction increased much further. Covariation was relatively large in the early phase of adaptation, but it became smaller with the other two components getting larger. This might suggest that actors resorted to better combinations of execution variables when perturbations were applied. It is possible that the actor strived to align their variability along with the original solution manifold, which is shown as a thin light grey line in Figure 4. This location offers slightly better solutions compared to other locations in the perturbed region. As a result, this alignment helped to improve the performance as compared to the performance with shuffled data set. However, after moving out of the perturbed region in the execution space, the solution manifold became flat and familiar to the actor again. As such the covariation changed back to its original values before perturbations were introduced. Taken together, results from the TNC analysis suggested that the changes in variability components during adaptation happened in sequence: confronted with a strange and unfamiliar solution manifold, the actor

remained on the original solution manifold and simultaneously relocated to a more tolerant area along the solution manifold; noise in execution was reduced to further improve the performance.

The actor seemed to take the variance of performance into consideration when relocating their new solution, since after adaptation the scatter of execution variables did not overlap with the distortion ellipsoid (Figure 4). The other evidence is that the distance to the center of the distortion ellipsoid (*DP*) was slightly larger than 1 plus one standard deviation of *DP* after 20^{th} trials (Figure 12). This effectively made most of the bounces stay out of the distortion ellipsoid. After the solution (with its scatter) just moved out of the distortion ellipsoid, it remained there without further movement on the solution manifold, as shown by a nearly constant *DP* after 20^{th} trial. The fact that actors just barely moved out of the perturbed region and not further, suggests that they not only change their performance (many other locations in execution space would have been possible), but they are also attracted by the previously preferred location in the task space. This resonates with the proposition that learning a new coordination pattern is governed by the interaction between initial coordination dynamics and the to-be-learned coordination pattern (Zanone & Kelso, 1992). The final location outside of the perturbation ellipsoid indicates the adaptation is an outcome of the compromise between the initial and the to-be-adapted patterns.

References

Aboaf, E. W., Drucker, S. M., & Atkeson, C. G. (1989). *Task-level robot learning: Juggling a tennis ball more accurately.* Paper presented at the Proceedings of IEEE International Conference on Robotics and Automation, Scottsdale, Arizona.

Amazeen, P. G., Amazeen, E. L., & Turvey, M. T. (1998). Dynamics of human intersegmental coordination: Theory and research. In D. A. Rosenbaum & C. E. Collyer (Eds.), *Timing of behavior: Neural, computational, and psychological perspectives* (pp. 237-259). Cambridge, MA: MIT Press.

Arutyunyan, G. H., Gurfinkel, V. S., & Mirskii, M. L. (1969). Organization of movements on execution by man of an exact postural task. *Biophysics, 14*, 1162-1167.

Beek, P. J., & van Santvoord, A. A. M. (1992). Learning the cascade juggle: A dynamical systems analysis. *Journal of Motor Behavior, 24*(1), 85-94.

Bernstein, N. (1967). *The coordination and regulation of movement*. London: Pergamon Press.

Bertalanffy, L. v. (1950). An outline of a general systems theory. *British Journal for the Philosophy of Science, 1*(2), 134-165.

Bertalanffy, L. v. (1973). *General systems theory: foundations, development, applications*. New York: Braziller.

Bühler, M. (1990). *Robotic tasks with intermittent dynamics*. Yale University, New Haven.

Buhler, M., Koditschek, D. E., & Kindlmann, P. J. (1994). Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research, 13*, 101-118.

Darling, W. G., & Cooke, J. D. (1987). Changes in the variability of movement trajectories with practice. *Journal of Motor Behavior, 19*(3), 291-309.

de Rugy, A., Wei, K., Müller, H., & Sternad, D. (2003). Actively tracking "passive" stability. *Brain Research, 982*(1), 64-78.

Dijkstra, T. M. H., Katsumata, H., de Rugy, A., & Sternad, D. (2004). The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies, 11*(3), 319-345.

Gottlieb, G. L., Corcos, D. M., Jaric, S., & Agarwal, C. G. (1989). Practice improves even the simplest movements. *Experimental Brain Research, 73*, 436-440.

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

Jagacinski, R. J., & Flach, J. M. (1999). *Control theory for humans*. Hillsdale, NJ: Erlbaum.

Katsumata, H., Zatsiorsky, V., & Sternad, D. (2003). Control of ball-racket interactions in the rhythmic propulsion of elastic and non-elastic balls. *Experimental Brain Research, 149*, 17-29.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.

Latash, M. L., Scholz, J. P., Danion, F., & Schöner, G. (2002). Motor control strategies revealed in the structure of variability. *Exercise and Sport Sciences Reviews, 30*, 26-31.

Milton, J. G., Small, S. S.; Solodkin, A. (2004). On the Road to Automatic: Dynamic Aspects in the Development of Expertise. *Journal of Clinical Neurophysiology. 21, 3,*134-143.

Müller, H., & Sternad, D. (2003). A randomization method for the calculation of covariation in multiple nonlinear relations: Illustrated at the example of goal-directed movements. *Biological Cybernetics, 39*, 22-33.

Müller, H., & Sternad, D. (2004). Decomposition of variability in the execution of goal-oriented tasks – Three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance, 30*(1), 212-233.

Schaal, S., Sternad, D., & Atkeson, C. G. (1996). One-handed juggling: A dynamical approach to a rhythmic movement task. *Journal of Motor Behavior, 28*(2), 165-183.

Scholz, J. P., & Schöner, G. (1999). The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research, 126*, 289-306.

Schöner, G., Zanone, P. G., & Kelso, J. A. S. (1992). Learning as change of coordination dynamics: Theory and experiment. *Journal of Motor Behavior, 24*(1), 29-48.

Sternad, D. (2000). Juggling and bouncing balls: Parallels and differences in dynamic concepts and tools. *International Journal of Sports Psychology*.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2000). Dynamics of a bouncing ball in human performance. *Physical Review E, 63*, 011902-011901 - 011902-011908.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2001). Bouncing a ball: Tuning into dynamic stability. *Journal of Experimental Psychology: Human Perception and Performance, 27*(5), 1163-1184.

Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Reading, MA: Addison-Wesley.

Tufillaro, N. B., Abbott, T., & Reilly, J. (1992). *An experimental approach to nonlinear dynamics and chaos*. Redwood City, CA: Addison-Wesley.

Worringham, C. J. (1991). Variability effects on the internal structure of rapid aiming movements. *Journal of Motor Behavior, 23*(1), 75-85.

Worringham, C. J. (1993). Predicting motor performance from variability measurers. In K. M. Newell & D. M. Corcos (Eds.), *Variability and motor control* (pp. 53-63). Champaign, IL: Human Kinetics.

Appendix A

For a trial with n bounces, each bounce i (i = 1, .. n) is described by the vector $\boldsymbol{e_i}$ = (X_{IP}, V_B, V_R), where X_{IP} is the impact position, V_B the ball velocity immediately before impact, and V_R is the racket velocity at impact. Thus, all bounces from a trial produce a matrix of execution variables $\boldsymbol{E}$ with a dimension of $n \times 3$. Each row of this matrix corresponds to one bounce and one data point in execution space. Calculating the resulting ball amplitude and the resulting error from the target line according to Equation 1, $\boldsymbol{E}$ produces a vector with n values of AE. As average of AE is regarded as the primary measure of performance, the mean of AE within one trial are calculated.

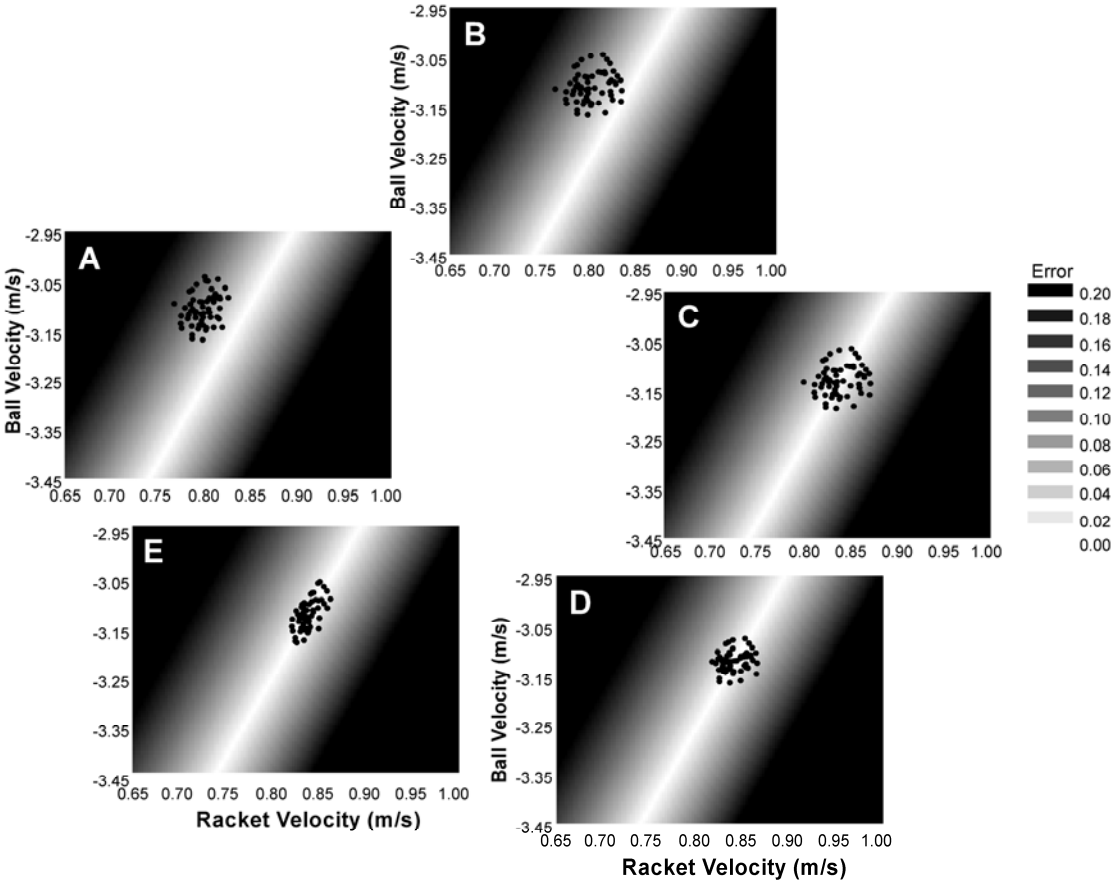

Figure A1: Schematic illustration of the computational steps of the TNC decomposition with fictive trials. Panels A and E illustrate a fictive trial early and late in the learning process, respectively. For details of the computation see the Appendix.

Using two fictive trials, the five panels in Figure A1 illustrate the computational steps for the TNC-decomposition. Panel A represents one trial at the early stage of learning, panel E at the late stage of learning. The five data sets in panels A-E are labeled as $\boldsymbol{E_A}$, $\boldsymbol{E_B}$, $\boldsymbol{E_C}$, $\boldsymbol{E_D}$ and $\boldsymbol{E_E}$ for clarity. The average result of *SD-AE* of $\boldsymbol{E_A}$ is worse than $\boldsymbol{E_E}$ as it occurs earlier in the learning process.

Step 1: Permutate each column of $\boldsymbol{E_A}$ and $\boldsymbol{E_E}$ separately. This permutation rearranges the data in the three execution variables such that a different result is obtained for each bounce. From these new results *SD-AE* are calculated. The permutations of Sets A and E are shown by $\boldsymbol{E}_B$ and $\boldsymbol{E}_D$, respectively.

Step 2: The mean locations of $\boldsymbol{E_B}$ and $\boldsymbol{E_D}$ are calculated by taking the means of each execution variable, separately. Move $\boldsymbol{E_B}$ to the mean location of $\boldsymbol{E_D}$ by adding the distance between $\boldsymbol{E_B}$ and $\boldsymbol{E_D}$ for each execution variable to $\boldsymbol{E_B}$. This step creates the new data set $\boldsymbol{E_C}$.

Step 3: Calculate *SD-AE* for each set by applying Equation 3.

Step 4: Calculate the following differences in *SD-AE* to obtain the three components:

- The difference between $\boldsymbol{E_E}$ and $\boldsymbol{E_D}$ is called *Covariation*, *C*, which quantifies the performance improvement due to co-varying execution variables in a

goal-relevant way. Covariation can also be calculated between $\boldsymbol{E_A}$ and $\boldsymbol{E_B}$.

- The performance difference between $\boldsymbol{E_B}$ and $\boldsymbol{E_C}$ is called *Tolerance*, *T*, which quantifies the performance improvement obtained from a better location that is more tolerant to errors or noise. The only difference between $\boldsymbol{E_B}$ and $\boldsymbol{E_C}$ is the location in execution space.
- The performance difference between $\boldsymbol{E_C}$ and $\boldsymbol{E_D}$ is termed *Noise*, *N*, which quantifies performance improvement due to decreased scatter in the data. Note that both $\boldsymbol{E_C}$ and $\boldsymbol{E_D}$ are at the same location in execution space but have no covariation.

CHAPTER 6

Variability and Stability in Learning and Interlimb Transfer

Chapter 6 contains an original manuscript prepared for submission to a journal

"Variability and Stability in Learning and Interlimb Transfer"

Mohamed Tlili[1], Kunlin Wei[1], Hiromu Katsumata[2], Dagmar Sternad[1]

[1]Department of Kinesiology and Integrative Biosciences, The Pennsylvania State University

[2]Department of Movement Sciences, Daito-Bunka University, Tokyo, Japan

Abstract

Using the task of rhythmically bouncing a ball with a racket held in one hand, it was demonstrated earlier that actors can employ a "passive" strategy where explicit error corrections are not necessary. This study pursues several aims: first, in two days of practice it tested whether stability changes concomitant with result measures of performance improvement. Second, it tested whether this strategy was equally expressed in the dominant and non-dominant hand. Third, a transfer design tested whether this strategy is transferred across hands. Fourth, a variability decomposition method aimed to uncover components of change in the structure of variability. Two groups of 10 right-handed participants performed the ball bouncing task on two successive days. On each day participants practiced the task for 30 trials with one hand, followed by 10 more trials performed with the second hand. Group 1 performed the sequence starting with their preferred hand, group 2 reversed the order. Results showed that stability had an exponential change towards more negative values indicating a strategy utilizing more passive stability. This change persisted even when performance variability plateaued to steady values, reflecting continued reorganization of strategy. Learning and continued reorganization was seen in both hands, regardless of dominance. The acquired passive strategy showed positive transfer to the opposite limb with no asymmetry in direction of transfer. Variability decomposition revealed significantly more exploration in the learning phase of the dominant hand. Finally, there appeared to be an asymmetry in the continued change in the dominant hand such that after continued improvement transfer to the non-dominant hand was no longer evident.

Introduction

Confronted with the ever-changing environment humans need to continuously learn and adjust their behaviors. Hence, when learning a new skill it is most central that this skill can be adapted and generalized to other situations. One example for such generalization is to perform an acquired skill with another limb. Motor learning and adaptation and, more specifically, interlimb transfer are long-standing issues in motor control that have received a lot of attention, particularly from more practically oriented research. Both for the training of athletes and in rehabilitation it is essential to know whether such interlimb transfer exists and how it can be used for optimal practice. Is training with the dominant hand concomitantly training the nondominant hand? Can training with the nondominant limb simultaneously improve skill with the dominant limb?

Significant evidence for such positive transfer has been provided in many different task contexts. The typical indicators by which adaptation and transfer has been evaluated and quantified are error measures and their variability. Yet, most studies in motor learning have had a very practical flavor, asking whether certain practice conditions lead to positive transfer or not. The questions about *what* is learnt and *what* it transferred are still insufficiently understood in more basic research terms. Several conceptual frameworks have been brought to bear to understand learning and transfer: Research couched in the framework of schema theory suggested that the learner acquires a generalized motor program that is subsequently parameterized for new tasks or new effectors, while the temporal structure of the movement is preserved across different effectors (Schmidt, 1975). A similar conceptualization of learning and transfer has been adopted in more recent accounts that propose internal models as the core concept of their theorizing (Imamizu, Uno, & Kawato, 1995; Kawato, 1999; Wolpert, Ghahramani, & Jordan, 1995). In these accounts the internal model, which is a nonlinear mapping between performance and intrinsic variables, is developed and parameterized during learning. For adaptation to new task demands or to a new effector, only parameters need to be tuned.

From a dynamical systems perspective, learning and transfer has been described as creating a dynamical model or a coordination dynamic (Mitra, Turvey, & Amazeen, 1997). In the context of learning new phase relationships between two rhythmic movements (Zanone & Kelso, 1992, 1997) showed how existent stable attractor states were reshaped and a new focal task was developed. Due to the effector-unspecific nature of the dynamic model, transfer of the learned skill would also be expected to occur across different effector systems (Kelso & Zanone, 2002). Stability is the core concept … Despite their different conceptual frameworks these accounts agree that there exists some higher-order structure that is acquired during skill practice which facilitates skill transfer to untrained tasks or untrained effectors.

In a series of recent studies Sternad and colleagues demonstrated the prominent role of dynamical stability in human performance at the example of a rhythmic ball bouncing task (de Rugy, Wei, Müller, & Sternad, 2003; Dijkstra, Katsumata, de Rugy, & Sternad, 2004; Schaal, Sternad, & Atkeson, 1996; Sternad, Duarte, Katsumata, & Schaal, 2000, 2001). Based on a mathematical model of the mechanical interactions of the ball and racket during rhythmic bouncing, stability analyses revealed that there are passively stable solutions to this task. Specifically, the criterion for a passively stable solution was that the racket contacts the ball in an upwardly decelerating phase. If humans adopt this performance strategy, then small perturbations will not require active error compensation. Empirical studies provided supportive evidence that actors indeed exploited this stable regime offered by the physics of the task and coordinated their limb movement to perform racket-ball impacts with negative acceleration (Schaal et al., 1996; Sternad et al., 2001). Furthermore, a previous study on learning revealed that novice subjects started off with positive impact acceleration (the unstable strategy from the model prediction) and with practice adopted the passively stable strategy. During the acquisition process, the impact acceleration was gradually tuned to the negative range (the predicted stable strategy), together with reduced performance variability (Dijkstra et al., 2004; Sternad et al., 2000). Hence, the stable solution for the ball bouncing task is not an intuitive solution. Instead, it has to be

"discovered" by the actor through practice. Given its subtle nature, one open question is whether this strategy is transferred across different effector systems. This is the first question that will be tested in the present study.

In most studies on motor learning and transfer, the primary indicator for performance improvement is the error or the variability of a given performance variable. Acquiring a skill is characterized by a reduction of errors and variability with practice. Decreasing variability is also often equated with increasing stability in performance. This simple inverse relationship, however, obscures the fact that empirical variability can be indicative of many different facets, ranging from the lack of control to more beneficial aspects, such as compensatory variation between variables, and exploration of new tasks (Riccio, 1993)To further the understanding of the variability and its changes over practice, Müller, Sternad and colleagues developed the so-called TNC-decomposition method that decomposes variability in redundant tasks into three components, tolerance, noise and covariation (Müller, Frank, & Sternad, 2007; Müller & Sternad, 2003, 2004). Central to this approach is the parsing of the task into execution and result variables. For example, reaching with a multi-joint arm to a target is separated into execution space spanned by the degrees of freedom of the joint variables and the endpoint error.

TNC variability decomposition decomposes the improvement of variability in result variables into three components: improvement originated from exploration and migration in execution space (*T*olerance), reduction of purely stochastic components in execution (called *N*oise), and improvement due to tighter covariation between execution variables (*C*ovariation).

Given redundancy in the motor system, the mapping between execution variables and the result variable is many-to-one, i.e., different combinations of execution variables can lead to the same result. This mapping can be conceptualized as a solution space with a family of solutions associated with different levels of "success" in the result variable. Depending on the nature of the solution space and inherent noise, different locations in the solution space have different probabilities of

success and different “tolerance” to noise. Improvements in performance, after some exploration of the task, are then related to the migration of performance to a better location in the solution manifold. This part of improvement is captured by the tolerance component. The second component, noise reduction, is a widely acknowledged phenomenon in skill learning. For the third component covariation, it has also been discovered for long that variability of execution variables can compensate for each other to achieve a better movement outcome. This covariation between execution variables has been interpreted as a signature of control.

By applying TNC variability decomposition, the learning process can be unraveled so as to identify the relative contribution of different components at different learning stages. For example, the aforementioned tuning to the stable regime offered by the task dynamics can be related to the tolerance and/or covariation changes at the early stage of learning, while noise reduction is a continuous process until certain asymptote is reached. More specific for the present study, is the question whether after initial acquisition with one effector, whether different variability components will change their contributions when transferring to another effector. This question will further the understanding of *what* is transferred under the surface of apparent improvement in a single result variable. Further, if variability is a signature of performance, does the structure of this variability change during learning and transfer?

The present study investigates the learning and transfer of the single-handed ball bouncing task in novice participants. The participant will start practice by using the dominant or nondominant hand, then switch to the contralateral hand. The performance error, variability in execution and the stability strategy adopted will be evaluated over practice and transfer. Positive transfer in performance variability between arms is expected. More importantly, it is expected that after subjects comprehend the task dynamics and exploit its stability property, they will transfer this strategy to the untrained arm. This transfer of the action strategy should be symmetrical between arms due to its abstract nature. However, some kinematic

measures of the task might be transferred asymmetrically due to the effects of handedness and proficiency of hemispheres. The learning and transfer process will also be scrutinized in terms of changes in variability by using TNC decomposition method.

Method

Participants. There were 20 volunteer young adults (10 males and 10 females) between 24 to 27 years of age. The participants declared themselves right-handed. They were also tested with the handedness questionnaire by Hull and showed values between 89 and 99%. None of the participants had previous experience with the experimental task. All of them had normal or corrected-to-normal vision. Prior to the study the participants were informed about the experimental procedure after which they signed the consent form in compliance with the Regulatory Committee of the Pennsylvania State University.

Experimental Apparatus and Data Collection. The task was performed in a virtual reality set-up the subject manipulated a real tennis racket in front of a large screen (2.5 m wide and 1.8 m high) onto which the visual display was projected (Figure 1). The subject stood upright at a distance of 1.0 m from the screen and held the racket vertically at a comfortable height. A rigid rod with three hinge joints and one swivel joint was attached underneath the racket and ran through a noose that rotated a wheel, whose revolutions were measured by an optical encoder. The digital signal from the optical encoder was transformed by a digital board (16 bit DT322 A/D card, Data Translation) and sent online to the computer that controlled the experimental set-up. Thus, only the vertical displacement of the racket was measured, even though the joints of the rod permitted the racket to move and tilt in three dimensions to avoid friction.

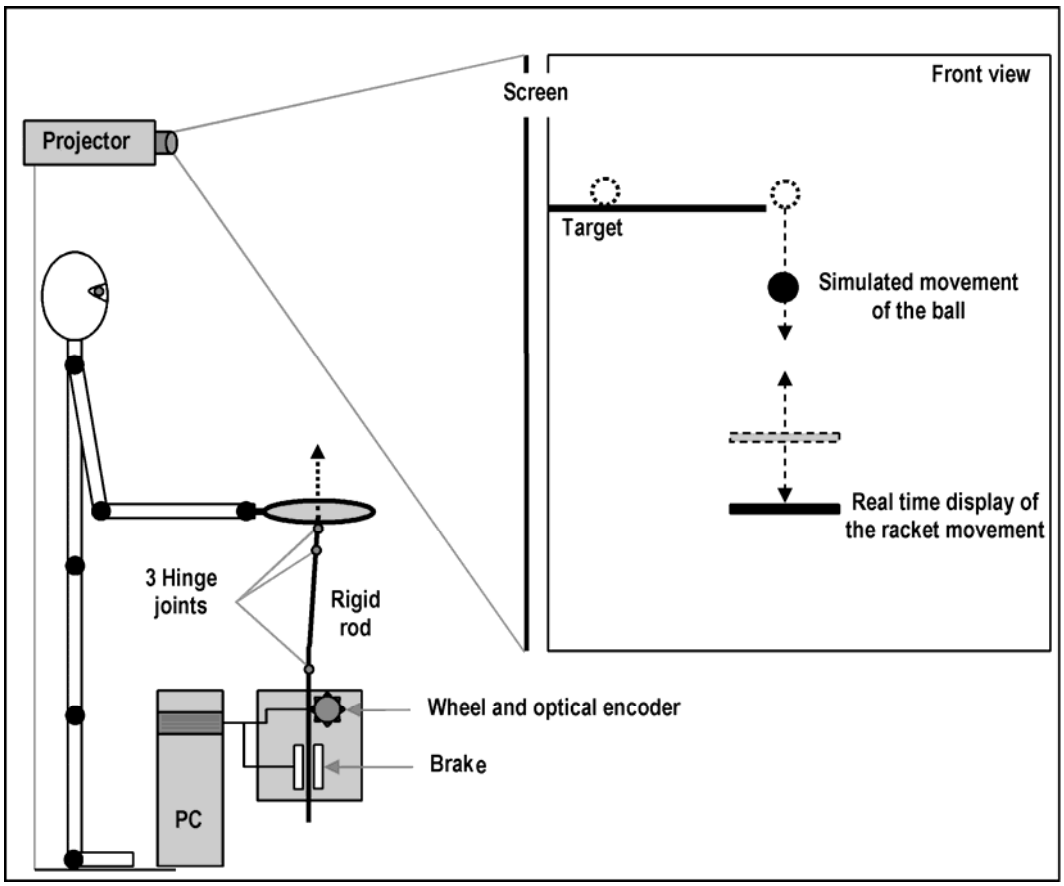

Figure 1. The virtual reality set-up and data collection for the ball-bouncing task.

The data acquisition program projected the displacement of the racket on the screen in nearly real time with an electromechanical delay of 11 ms. The virtual racket was displayed at the same height as the real racket and its displacement was the same as that of the real racket. The participant manipulated the real racket to control the movement of the virtual racket, which bounced a virtual ball on the screen. This movement of the virtual ball was determined by ballistic flight and the racket-ball impacts. In order to simulate the impact force between the racket and the ball, a mechanical brake was attached to the rod. The brake applied force to the rod at each impact for a duration of 30 ms, which was close to the real racket-ball contact duration (Katsumata, Zatsiorsky, & Sternad, 2003). The brake would decelerate the motion of the rod and the attached racket, which was sensed by the participant. The force developed by the brake was adjusted to mimic the amount of impact that

produced by a tennis ball falling on the racket. However, due to the hardware constraints, the force was constant for all contacts irrespective of the racket/ball velocity at each contact. The data acquisition program was written in Visual C++ (Microsoft, v6.0) and the virtual display was realized by Open GL Graphics. The sampling frequency of the racket and the ball movement was 800 Hz, but the visual display was updated at 75 Hz.

The ball's trajectory was calculated using the equations of ballistic flight and elastic impact. When the ball was in the air, its vertical displacement was only influenced by gravity:

$$x_B = -1/2gt + \dot{x}_{B,0}t + x_{B,0} \qquad (1)$$

where x is the vertical position of the ball, g is gravity and $x_{B,0}$ and $\dot{x}_{B,0}$ are the initial position and velocity of the ball immediately after impacts, t is the time elapsed from the last impact. Assuming instantaneous elastic impact, $\dot{x}_{B,0}$ was determined by the impact relation:

$$(\dot{x}_B^+ - \dot{x}_R) = -\alpha(\dot{x}_B^- - \dot{x}_R) \qquad (2)$$

where $\dot{x}_B^-$ and $\dot{x}_B^+$ denote the velocity of the ball immediately before and after contact, respectively; $\dot{x}_R$ denotes the velocity of the racket at impact. The coefficient of restitution α captures the energy loss at impact with a value ranging between 0 and 1. 0 means complete energy dissipation and 1 means no energy loss during impacts. In the present experiment, α was set to 0.5. It was also assumed that the mass of the racket was sufficiently larger than the mass of the ball such that the effect of the impact on the racket trajectory could be ignored.

Procedure and Experimental Conditions. Each trial began with the ball appearing on the left side of the screen and rolling on a horizontal line that extended to the middle of the screen. When the ball reached the end of the line, it dropped down towards the racket (Figure 1). This starting procedure visually prepared the participant

for the beginning of the task. The virtual racket was located directly below the drop position of the ball and it could move in the vertical direction only. Participants held the racket with one hand and waited for the ball to drop. The participant was instructed to rhythmically bounce the virtual ball such that the peaks of the ball trajectory were as close as possible to the target line (the same line that the ball started on, 0.7m above the initial racket position). Participants had visual information about their performance from the continuous visual display of the racket and ball on the screen.

Participants were divided into two groups of 10 individuals each. Each group performed two sessions of a total of 44 trials each; the two sessions were scheduled on two consecutive days. The first four trials in each session tested performance in the preferred and non-preferred hand to evaluate the participant's initial performance in each hand (pre-test). This was followed by a practice phase of 30 trials and a transfer phase of 10 trials. Group 1 performed the pre-test first with their preferred hand (PH) for two trials and then with their non-preferred hand (NPH), again for two trials. The practice phase was performed with their preferred hand and the transfer phase with their non-preferred hand. On Day 2 the order of hands was switched (Figure 2). Group 2 performed all parts of the two sessions with the order of hands reversed hand. Before the formal data collection on the first day, participants were given two extra trials with each hand to familiarize themselves with the virtual task. These trials were not included in the analysis.

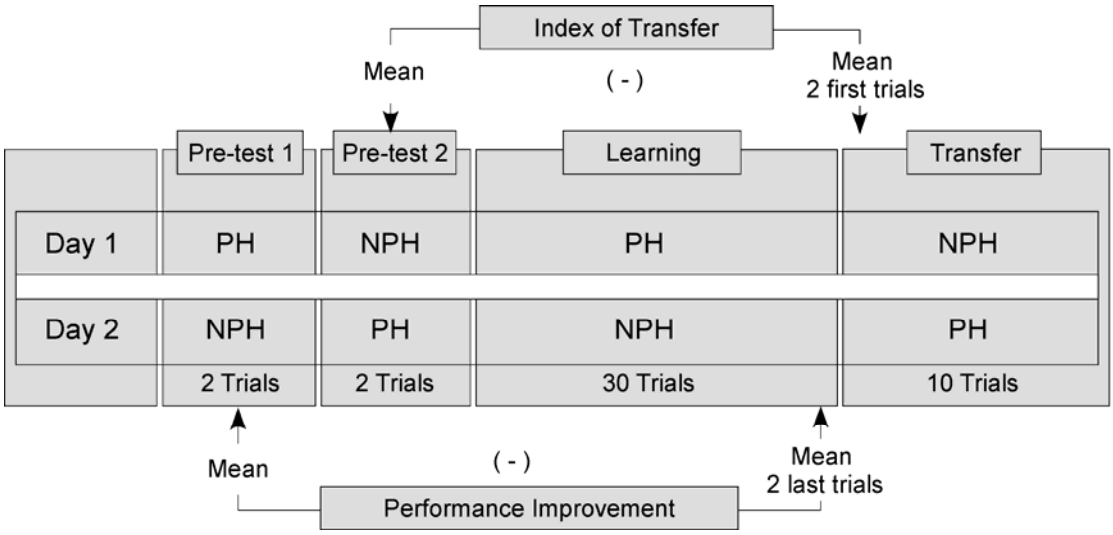

Figure 2. Experimental design for Group 1. PH refers to the performance with the preferred hand; NPH refers to the non-preferred hand. Group 2 followed the same pattern of sessions only the sequence of preferred and non-preferred hand was reversed.

Data Reduction and Analyses. The displacement and the acceleration data of the racket movements were low-pass filtered using the Savitzky-Golay filter with a cutoff frequency of 20 Hz. The displacement of the ball did not require further processing as it was simulated and, hence, contained no measurement noise. Figure 3 shows a typical time series showing the approximately sinusoidal displacements of the racket and the ballistic flight trajectory of the ball. The ball-racket contacts are highlighted by the symbols and the vertical dashed lines. The target line at 0.70m is also displayed. The difference between the apex of the ball trajectory and the target height is the error defined as the absolute difference between the peak of the ball trajectory and the target height, *AE*. The acceleration trace of the racket shows a more noisy signal with a very marked sudden change at the time of the ball-racket contact. At each bounce the acceleration of the racket one sample directly before the impact was used for the measure *AC*. For each trial the values of all bounces were averaged their standard deviations were determined, *SD*.

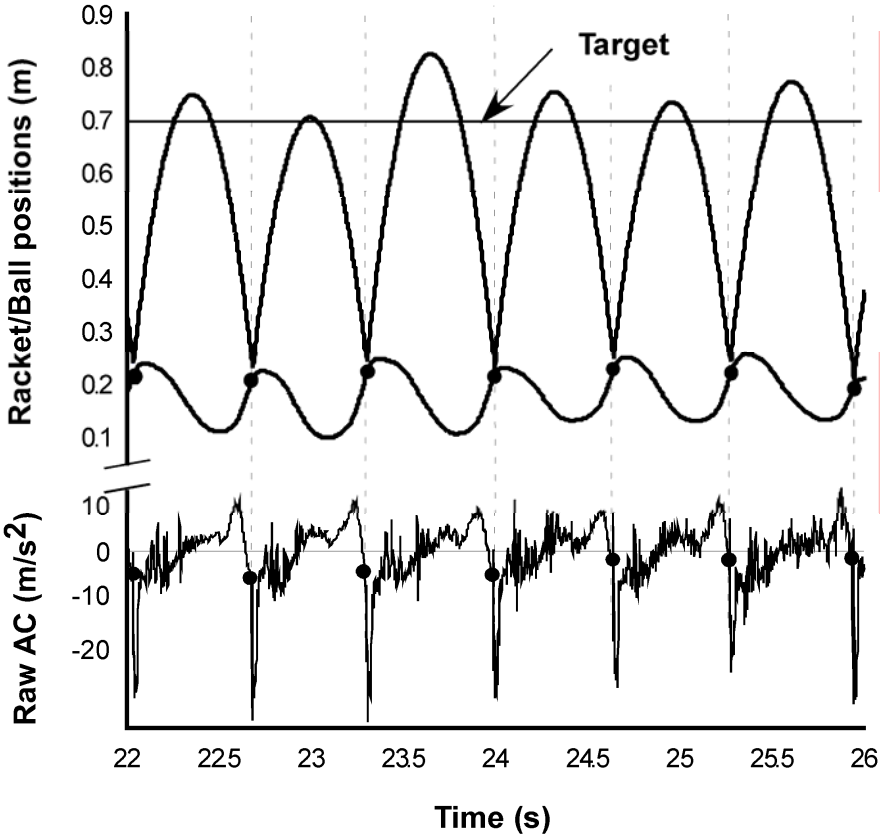

Figure 3. Segment of an exemplary time series of the ball and racket trajectory and the acceleration signal measured separately by an accelerometer. The grey dots and the dashed vertical line correspond to the moments of the impact of ball-racket. The solid horizontal line represents the target height.

To assess the changes in performance during practice and transfer sessions the time course of both *SD-AE* and *AC* were quantified by fitting exponential functions over the sequence of trials. Using the Levenberg-Marquardt method in Matlab v.7 for curve fitting the rate of change was estimated by the following function:

$$f(t) = a + b * e^{-\frac{t}{\tau}} \quad (3)$$

where t is the trial number, f is the dependent measure, a and b are fitting parameters for the regression to estimate. The parameter t denotes the learning rate, as it is the relaxation time of the exponential function.

To evaluate the amount of performance improvement during practice and transfer, changes in *SD-AE* and *AC* were quantified by two indices (Figure 2). The index of performance improvement, *PI*, was calculated by the difference of the mean of the two trials of the pre-test (trial 1 and 2) and the mean of the two last trials of the practice phase (trial 33, 34). To evaluate the transfer across hands, the index of transfer, *IT*, was calculated as the difference between the mean of the two pre-test trials (trials 3, 4) and the mean of the two first trials in the transfer session (trials 35, 36). Each index compared performance of the same hand.

Statistical Analyses. The dependent measures were subjected to mixed-effect 2 (day) x 2 (group) ANOVA and where considered significant when $p < 0.05$. Only significant effects are reported.

Task Analysis. Different to the model-based analysis of stability which analyzes the task performance as a dynamical system, the ball bouncing task can also be viewed as a sequence of individual bounces, where each bounce propels the ball to the given target height. The height of the ball trajectory is determined by three execution variables: the racket velocity at ball impact V_R, the ball velocity immediately before impact V_B, and the impact position X_{IP}. The bounce height is related to these variables by the ballistic flight equation (Eq 1) and the elastic impact equation (Eq 2). More specifically, the performance measure absolute error *AE* can be written as a function of these three execution variables:

$$AE = \left| ((1+\alpha) * V_R - \alpha * V_B)^2 / (2 * g) + X_{IP} - X_T \right| \qquad (4)$$

where X_T is the target height set at 0.7 m in the present study. By applying this equation, a triplet of execution variable (V_B, V_R, X_R) can lead to one result variable *AE*. Importantly, different combinations of the three execution variables can lead to the

same result, i.e., the task is redundant. This relation can be viewed by spanning a space of execution variables and plotting the results for all combinations of the execution variables. Figure 4A displays the execution space and the surface represents the set of solutions that have zero error (AE = 0m). This manifold of solutions is approximately planar for ball bouncing but may take on any nonlinear shape depending on the task. There are an infinite number of other manifolds that are defined by different results, i.e., AE = 0.05, 0.10m, etc.; however, for clarity, only the zero solution manifold is plotted. The points represent the bounces of one trial.

To better graphically illustrate the conceptual steps of the variability decomposition of variability the three-dimensional space is "sliced" at one value of X_{IP} to obtain the two-dimensional depiction in Figure 4B. The different grey shades denote different levels of error, AE = 0.03m, with white denoting AE = 0m (see legend). Each data point is one bounce, such that the set of points in Figure 4A and 4B represent one trial. Once, the task is depicted in this fashion, the variability of trials and its change across trials can be further decomposed.

A

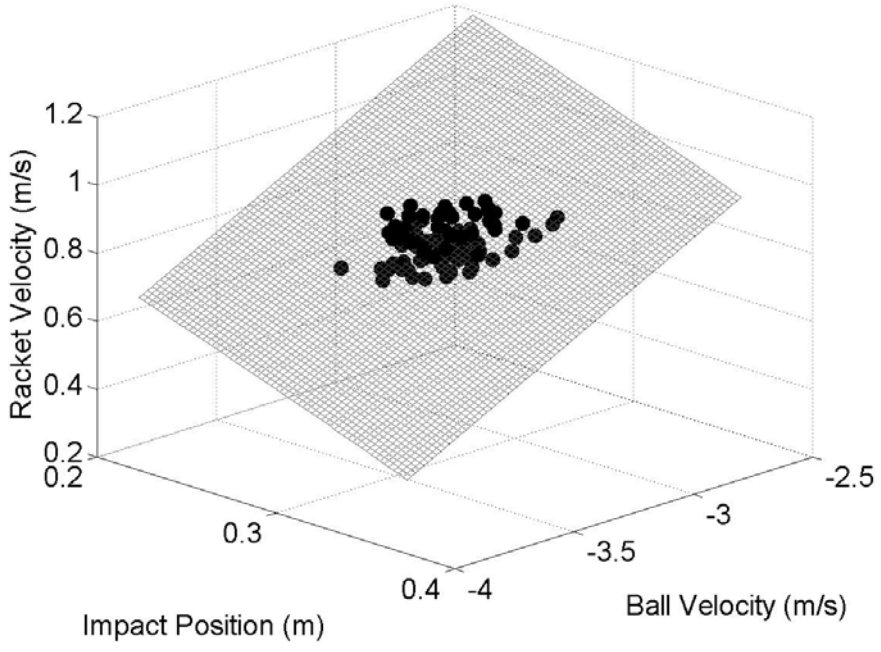

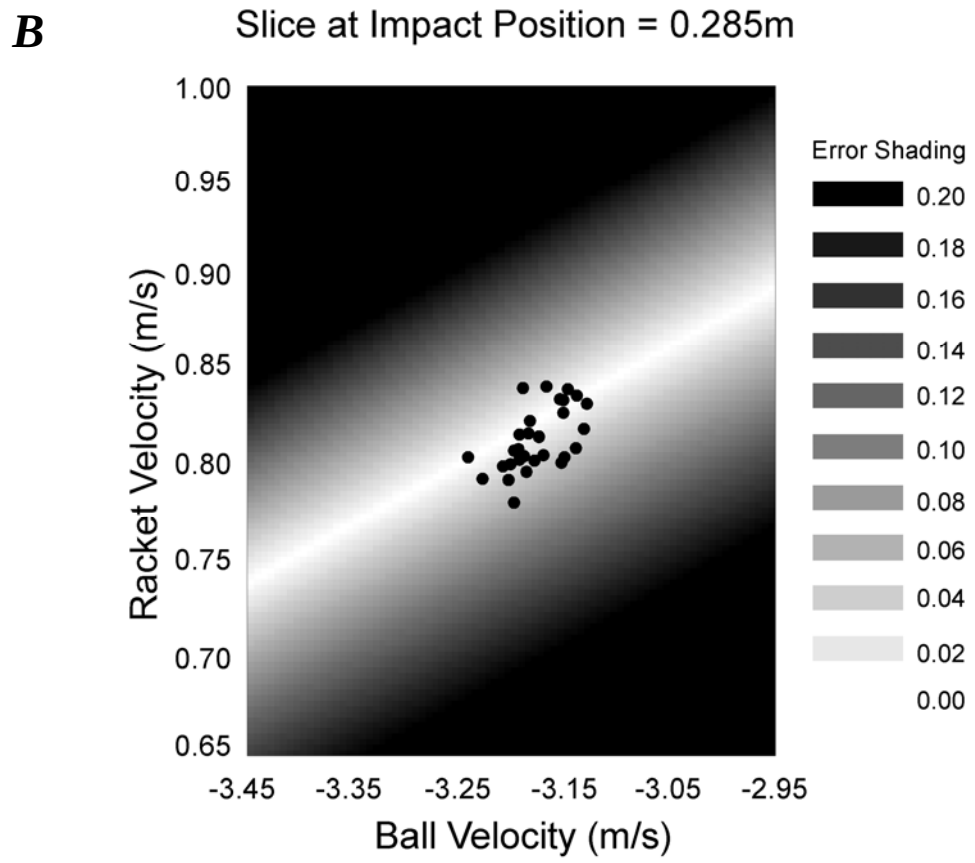

Figure 4: **A**: 3D illustration of the execution space for the ball bouncing task. The surface indicates the composite of all solutions that lead to exact hit of the target. Each dot stands for a single bounce. **B**: 2D displace of the execution space with impact position fixed at 0.285m. The grey shades stands for the error the target.

Quantification of the Components of Variability. If the task can be parsed into execution and result variables with a known functional relation (Müller & Sternad, 2003, 2004) proposed a method to decompose the variability in performance into three components: tolerance *T*, noise *N*, and covariation *C*. The application of the method requires a parsing of the performance into sets of data. For ball bouncing each trial of 40s duration forms a set with approximately 50-70 bounces. For the following explanation of the decomposition method, two fictive data sets, Set A and Set B, are shown in the two-dimensional slice of the execution space of Figure 4B.

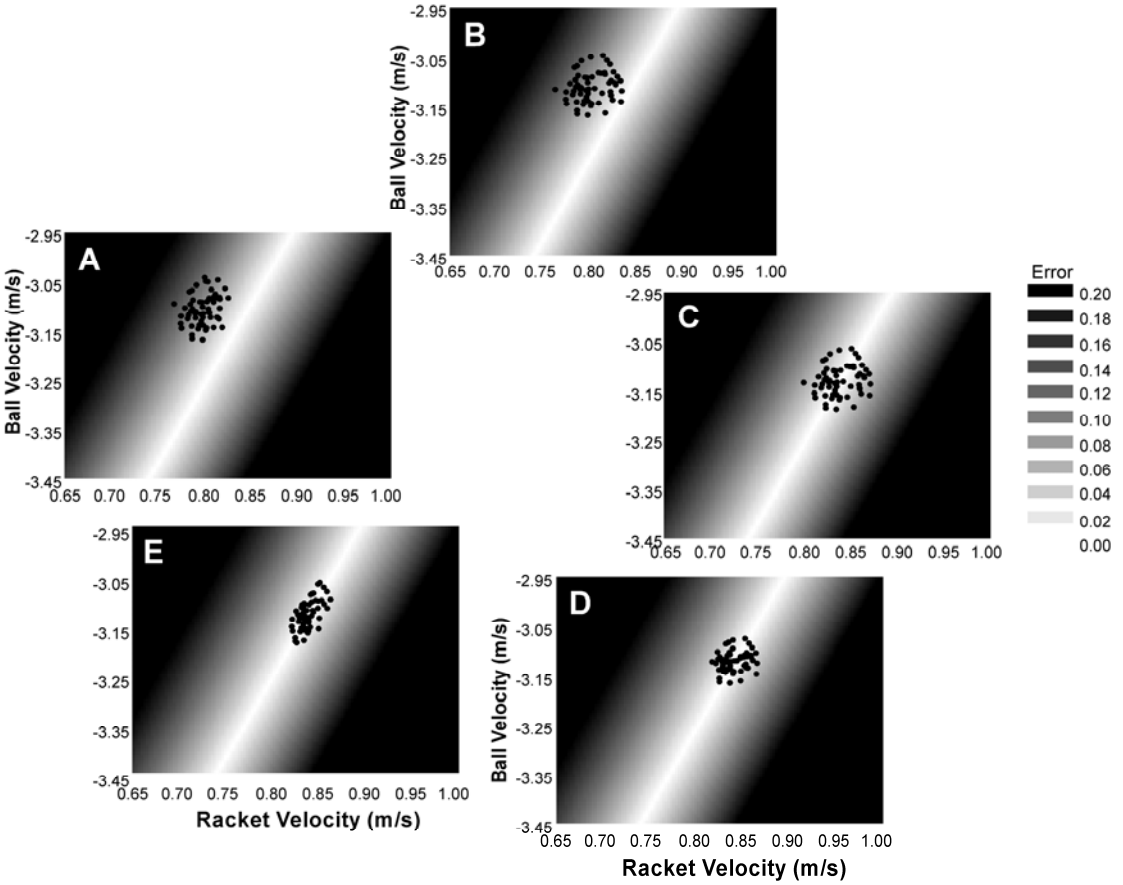

Figure 5: A schematic illustration of computational steps for TNC variability decomposition. Subplot A and E illustrate two fictive trials early and late in the learning process, respectively. See details in Appendix.

Panel A shows a data set of a single trial at the early phase of learning, while panel E shows data at the late phase of learning. Set E yields better performance due to three types of changes compared to Set A: Set E is at a location with smaller performance errors as indicated by the lighter shades of grey; Set E has less dispersion; Set E has an elliptic shape and aligns with the solution manifold. The performance improvement in *SD-AE* from Set A to Set E can be decomposed into three components to quantify these components: The performance improvement due the different locations in execution space between A to E is captured by *Tolerance, T*, visualized by the comparison of Set B and Set C. Note that before this shift in location,

Set A was permutated to eliminate influences from the orientation of the data distribution. The changes in performance before and after the permutation quantify the contribution of *Covariation C*. The third component quantifies how much a decrease in dispersion reduces the performance error, a component called *Noise N*. The component is visualized in the comparison between Set C and Set D. Note that Set D is obtained from transforming Set E by permutation. More calculation details for these three components are shown in the Appendix (see also Müller & Sternad, 2004)

To learn how *T* and *N* change during practice and transfer, comparisons were made between the first trial and each of the subsequent 43 trials of each day. *C* is calculated by taking the difference in performance between one data set and its permutated data set. Thus, *C* represents the covariation within a particular trial.

Results

Initial Performance Level. Before the specific analyses on the effects of practice and transfer were examined, it first had to be established that the two groups were of the same initial performance level. Second, it was of interest to test whether the preferred and non-preferred hand had different degrees of skill in the ball bouncing task. These questions were assessed by comparing the two dependent measures, the variability estimate standard deviations of absolute errors, *SD-AE*, and the racket acceleration at impact, *AC*, across hands and groups. For each participant the means of the two trials of the pre-test performed with the same hand were calculated and all participants' estimates were submitted to a 2 (group) x 2 (hand) ANOVA. Neither of the two dependent measures yielded significant effects. Hence, the two groups could be considered of comparable skill level before starting the experiment. Second, neither the degree of variability nor the values of racket acceleration were different for the dominant and non-dominant hand at the initial stage.

Performance Improvement during Practice. To evaluate the improvement in performance with practice on the two days the variability measures *SD-AE* across trials were fitted with exponential functions. Figure 6 shows the time course of the

group averages of *SD-AE* for each day, together with the fitted exponential curves for both groups on Day 1. Each data point is the average value over 10 participants of one group for each trial with the error bars denoting standard errors across the 10 participant means. Note the first two trials of the same hand from the pretest were included in the curve fits as they were regarded as part of the practice process. These exponential fits were only applied to the practice phases only, as the 10 trials in the transfer phase were not long enough to warrant the same curve fit. For Day 2 no exponential curves are shown because there was essentially no change in these performance measures and exponential fits were consequently not appropriate. The R^2-values of the shown fits for the group averages were .90 and .78, for Group 1 and 2 respectively. The learning rates were 7.84 and 8.35 for Group 1 and 2, respectively. The same fits were performed on each individual's data and the learning rates of these 10 curve fits were submitted to a one-way ANOVA comparing the two groups. No significant differences could be identified.

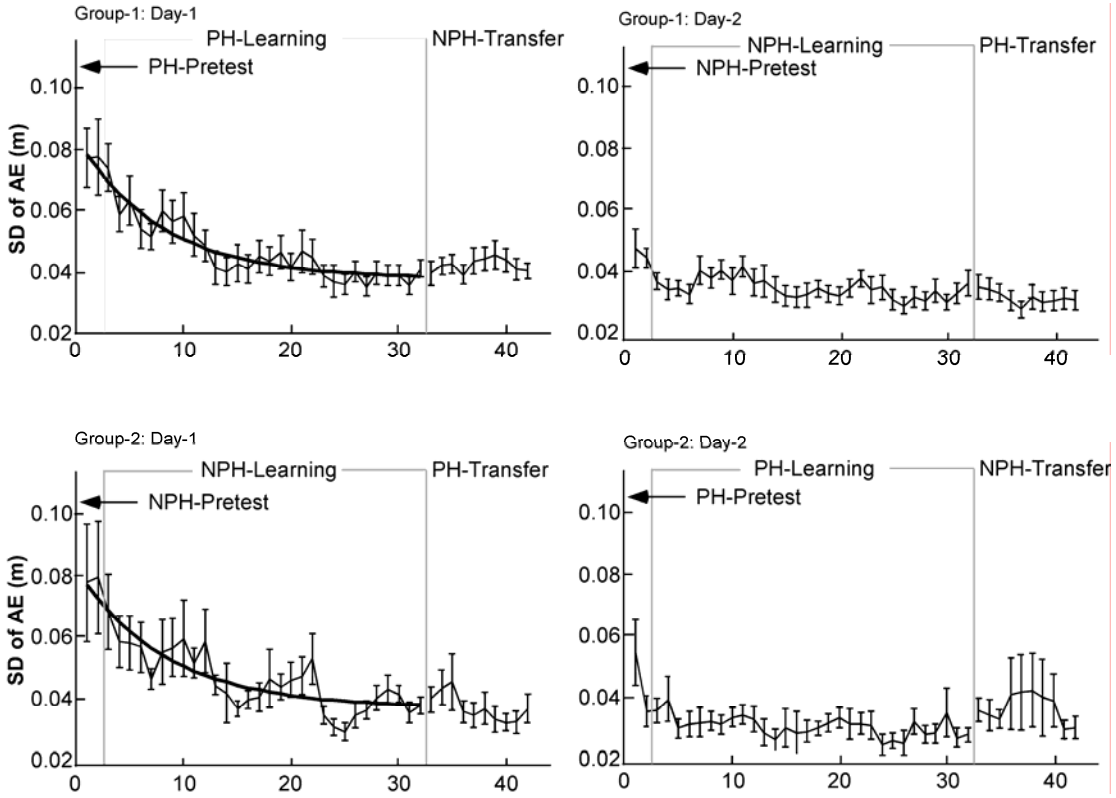

Figure 6: Standard deviations of the absolute error *SD-AE* are plotted as a function of trials for two groups and two days separately. The error bars stands for standard errors across participants. The bold lines are the fitted exponential curves.

To quantitatively assess the amount of improvement for the two groups on the two days, the difference in mean *SD-AE* in the two last trials of the practice phase (trials 33, 34) and the two pre-test trials of the same hand was calculated for each participant. This performance index, *PI-SD-AE,* was compared between days and groups with a 2 (group) x 2 (day) ANOVA. The analysis demonstrated that the improvement on Day 1 was larger than on Day 2 only a significant main effect for day, $F(1, 36) = 5.77$, $p < .05$ (Figure 7A). This showed that the change was comparable for dominant and non-dominant hand. As the exponential learning curves already illustrated above, the two groups showed similar learning rates. Given the relatively small performance improvement on the second day, additional *t*-tests evaluated whether the improvements on Day 2 were different from zero. One-sample *t*-tests against zero showed that the improvement was still significant for both groups on the second day (*p*-values < 0.05). However, as the time series in Figure 6 illustrates this difference in performance is probably only due to the warm-up decrement in the first few trials as the data appear to show no change during practice. To test this interpretation, the same performance index was calculated between the first trials of the practice phase (trial 5 and 6) and the last trials (trials 33 and 34). With this adjustment the difference became non-significant ($p = 0.40$). This suggests that in the practice phase on Day 2 there was no more noticeable improvement. A last close inspection of the plots suggests that on Day 2, the second group, which practiced with the preferred hand, lowered their variability score further than Group 1. This small difference, although not significant, may be helpful in interpreting a significant effect in the transfer phase below.

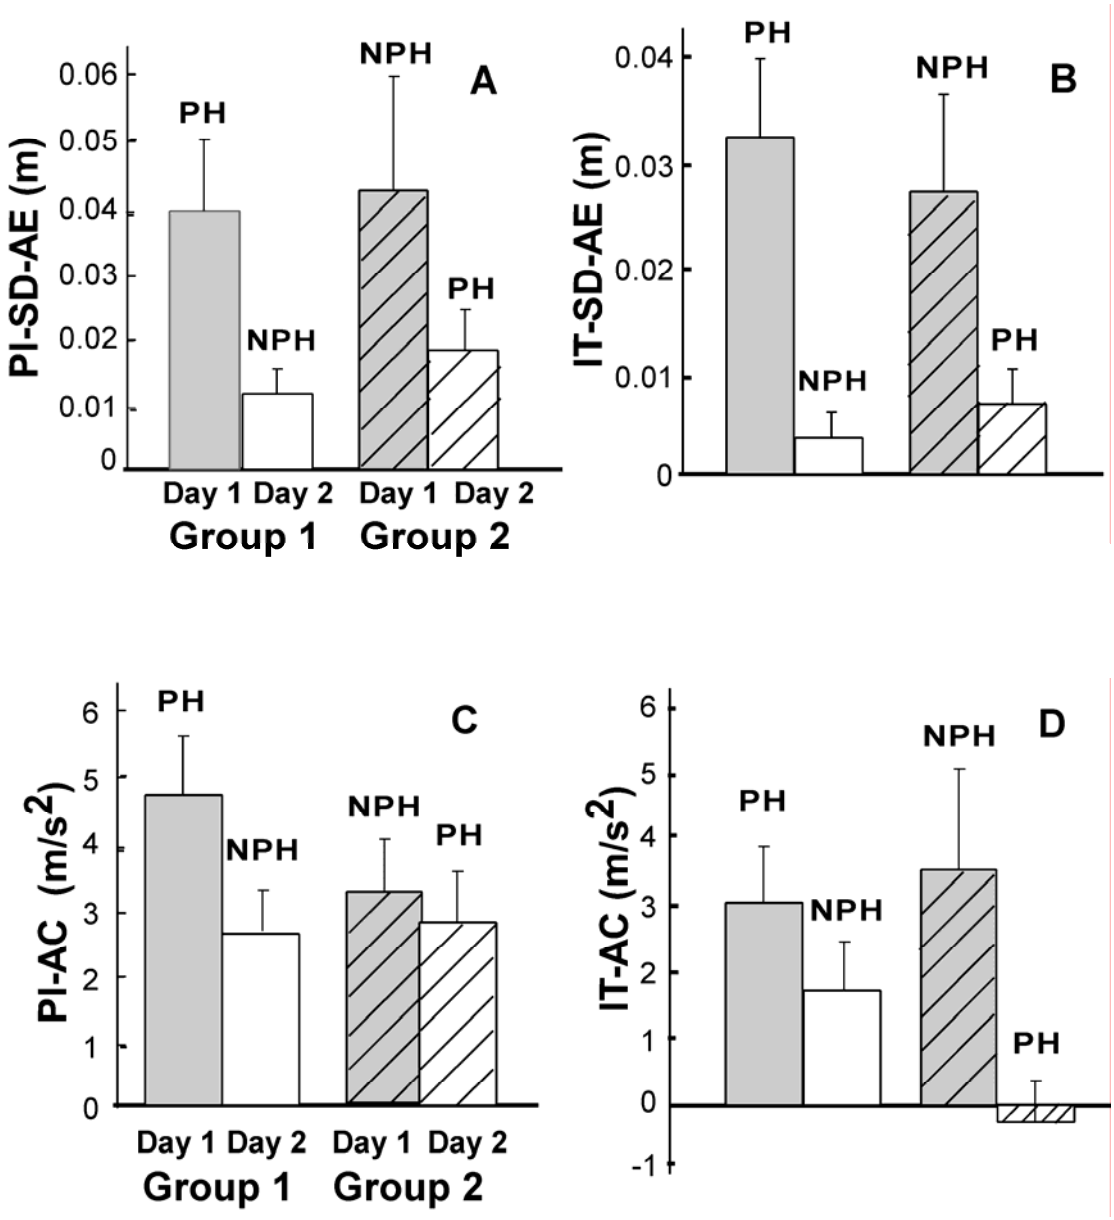

Figure 7: **A**, **B**: Performance index and index of transfer for the standard deviations of absolute error, *PI-SD-AE* and *IT-SD-AE*, respectively, for both days and groups (for definition see text). **C**, **D**: Performance index and index of transfer for the racket accelerations at impact, *PI-AC* and *IT-AC*, respectively, for both days and groups (for definition see text). The error bars denote the standard error across the participant means.

Transfer of Performance. To assess whether there was skill transferred from practice with one hand to the other hand, the index of transfer was calculated as the difference between the means of the first two trials of the transfer phase (trial 35, 36)

and the means of the two trials of the pre-test performed with the same hand (*IT-SD-AE*). *t*-tests compared these *IT-SD-AE* values against zero, separately for both groups and days. The *t*-tests for Day 1 yielded significant differences for both groups, attesting to the fact that there was positive transfer: Group 1: $t(9) = 4.32$, $p < 0.005$; Group 2: $t(9) = 2.80$, $p < 0.05$. In contrast, on the second day the differences were no longer significant: Group 1: $t(9) = 1.42$, $p = 0.19$; Group 2: $t(9) = 1.55$, $p = 0.16$. This absence of transfer on the second day was probably due to the fact that also the performance did not change very much any more. To assess whether this positive transfer was different for the two groups, *IT-SD-AE* was submitted to a two-way ANOVA. Results only confirmed the effect of day seen that was already in the *t*-test above: the main effect of day was significant, $F(1, 36) = 14.38$, $p < .005$. Figure 7B summarizes this result.

Effects of Practice in Racket Acceleration. Figure 8 displays the time course of the participant averages of *AC* per trial including the standard errors across the participant means. As above, the two pre-test trials that were performed with the same hand as the practice session are included in the figure. Both groups started with positive *AC* values on Day 1 that gradually decreased to values close to zero across the 30 practice trials. This trend persisted in the transfer phase. On Day 2 the average trial values were positive in the pre-test but then decreased towards negative values. To highlight the time course exponential functions were fitted to the data which are shown by the solid lines in the practice phases of both days. Note that in contrast to *SD-AE* also the second day of both groups shows an approximately exponential change in *AC*. The R^2-values and the learning rates for the mean fits were: Group 1, Day 1: 0.60, 2.44; Group 1, Day 2: 0.72, 7.27; Group 2, Day 1: 0.55, 1.60; Group 2, Day 2: 0.62, 9.49. The ANOVA on the individual learning rates did not reveal any significant differences between groups and days.

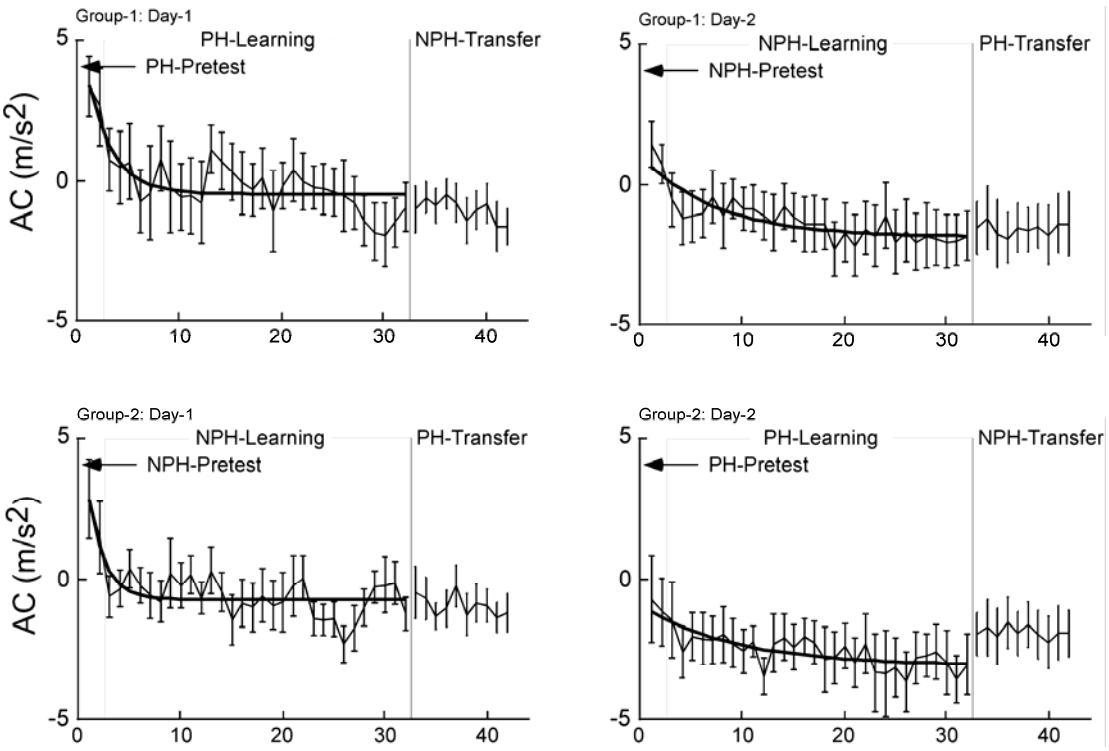

Figure 8: Racket acceleration at impact, *AC*, for both groups and days with both preferred and non-preferred hands, plotted for both groups across practice and transfer phases. The data show the average values across 10 participants; the error bars denoting standard errors across the 10 participant means. The solid lines for the practice phases on Day 1 and Day 2 are the exponential fits to the data.

To evaluate the change in *AC* due to practice, the same performance index as above was calculated and analyzed by a two-way ANOVA. Neither the two main effects, nor the interaction showed significant differences (Figure 7C). Contrary to the results on *SD-AE* (Figure 7A), *PI-AC* did not show significant differences between Day 1 and Day 2, indicating that there was still a comparable amount of change in *AC* on Day 2. This highlighted that while the performance variables *SD-AE* reached a plateau on Day 2, changes in *AC* still continued. To establish that the changes were indeed positive and different from zero, four *t*-tests were performed on *PI-AC*. All four *t*-tests verified a significant difference from zero with two *p*-values smaller than 0.005 and two greater than 0.05.

Further inspection of the data on Day 2 revealed that, similar to *SD-AE*, there were some differences between groups in the *AC* values. Although small, the *AC*-values that Group 2 reached at the end of practice on Day 2 were lower than in Group 1. This non-significant trend is in accord with the slightly lower variability observed in the performance results and is of interest when interpreting the change in the transfer phase observed after practice with the preferred hand on Day 2.

Transfer in Racket Acceleration. To test whether there was evidence for positive transfer in the *AC* values, the index of transfer, *IT-AC*, were calculated for both days and groups. Four *t*-tests compared the *IT-AC*-values against zero. Three comparisons established significant differences indicating positive transfer (*p*-values were smaller than 0.01, 0.05, and 0.05, for Group 1, Day 1 and 2, and Group 2, Day 1, respectively). The exception was the second day of Group 2 with a *p*-value of 0.706. Viewing the time profiles in Figure 8, this absence of a transfer effect can be understood by the increase of *AC*-values from the practice to the transfer phase of Day 2. Note, that this change in performance corresponds to a change observed in *SD-AE*, which, however, did not achieve significant results (see Figure 6). When *IT-AC* was submitted to a 2 (group) x 2 (day) ANOVA, the main effect of day was significant, $F(1, 36) = 6.21$, $p < .05$. The overall smaller transfer on Day 2 is partly attributable to the same observation, even though the interaction did not reach significance. In sum, after continued practice with the preferred hand the same improvement in performance is no longer transferred to the non-dominant hand.

TNC Decomposition. To reveal more aspects of these changes across the practice and transfer phases the changes in the variability measure *SD-AE* were decomposed into the three components Tolerance, *T*, noise, *N*, and covariation, *C*. Tolerance was calculated as a difference measure that quantifies the performance gain due to change in the location in execution space between the initial trial and trial *n*, where *n* refers to the trial number across practice and transfer. Performance gain was measured by the decrease in *SD-AE*. Analogously, *N* expresses the difference in performance between the initial trial in the pre-test and trial *n*. Covariation is derived from comparison

between the original data of trial *n* and the surrogate data set. Figures 9 and 10 shows the changes across practice and transfer on the two days, separately for the two groups. Inspecting Group 1 on Day 1, both *N* and *T* gradually increased throughout the practice phase while *C* remained almost constant. This lack of change in *C* indicates that the contribution from covarying execution variables remains unchanged throughout practice. While the highest contribution to performance improvement came from *N*, *T* showed a similarly growing contribution to the observed performance improvement.

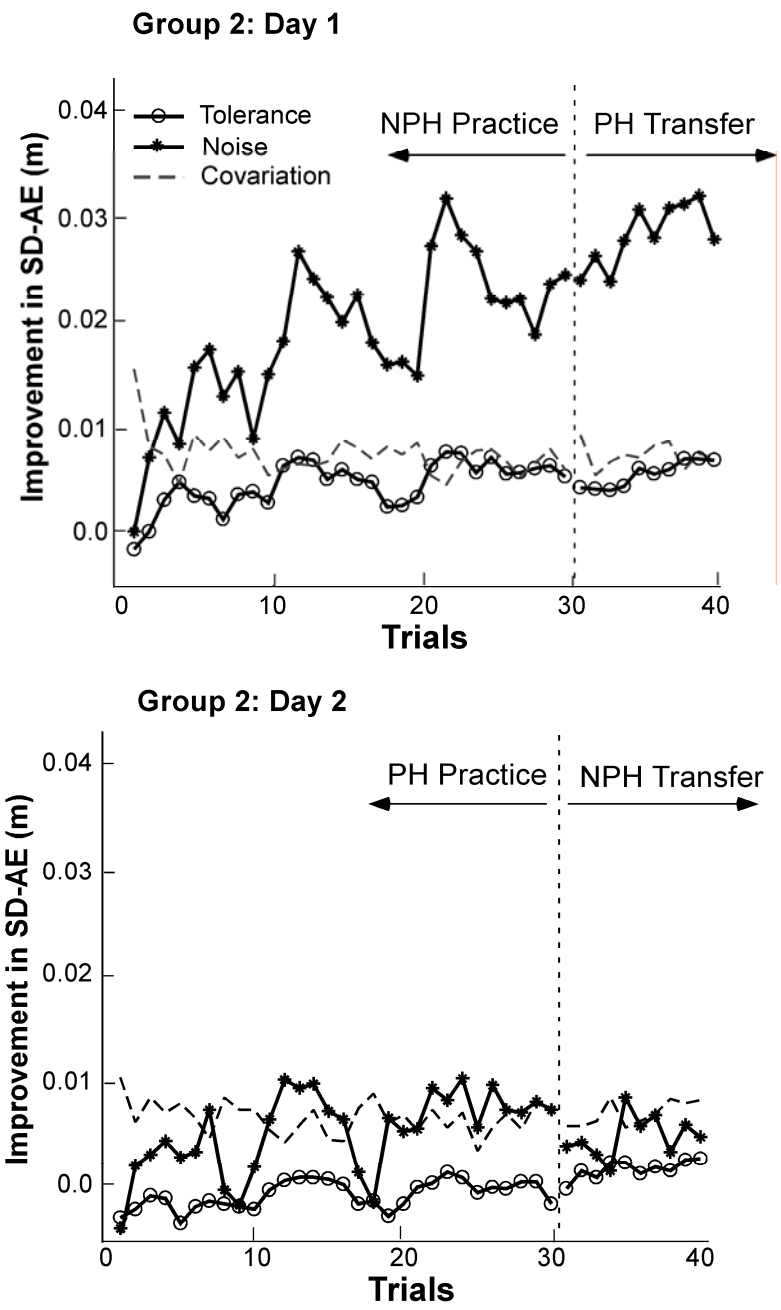

Figure 9: Time course of Tolerance, Noise, and Covariation and their relative contribution to performance improvement during acquisition and transfer in Group 1. The three components are plotted as a function of trials for two days.

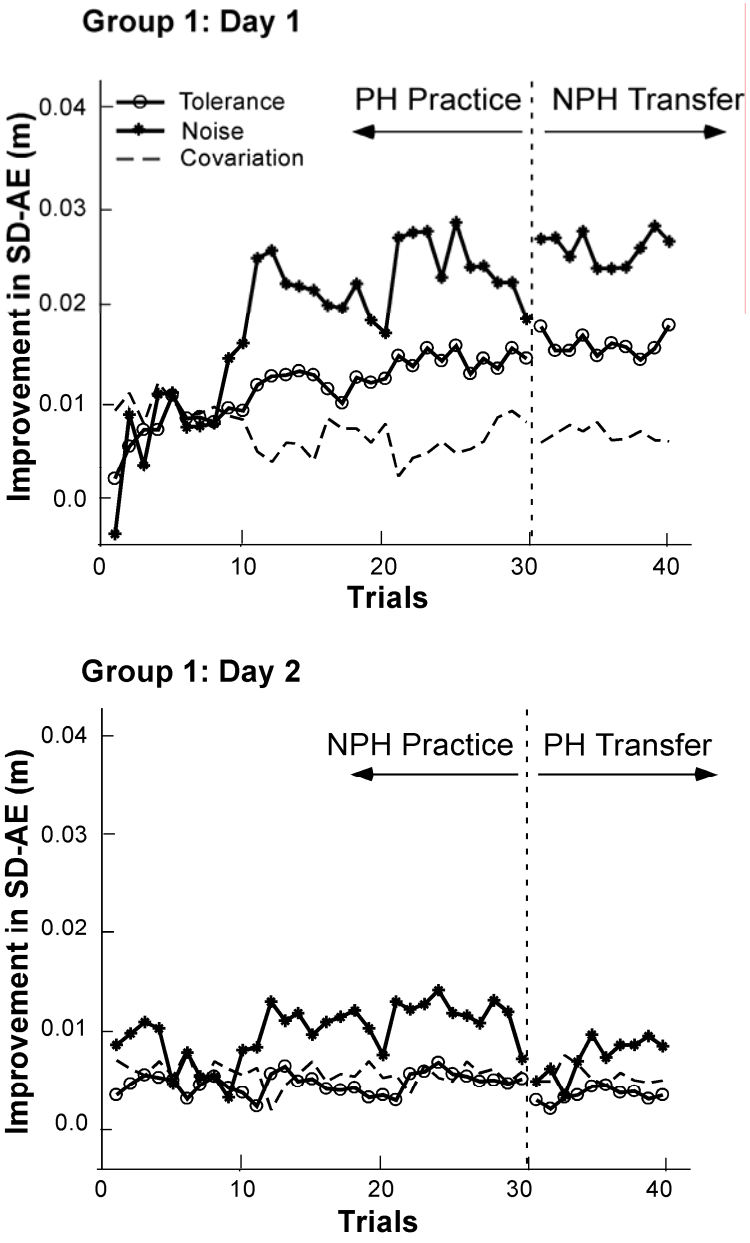

Figure 10: Time course of Tolerance, Noise, and Covariation and their relative contribution to performance improvement during acquisition and transfer in Group 2. The three components are plotted as a function of trials for two days.

When switching to the non-preferred hand in the transfer phase, the three variability components did not show any abrupt changes, indicating relatively complete transfer of all three components. On Day 2 for Group 1 the three components were small and close to zero, reflecting again the fact that there was not much change in the performance variable *SD-AE*. This result was accompanied by a weak contribution of the three components in the transfer phase. For Group 2 the variability components showed similar patterns of change across practice and transfer trials, except one noticeable difference: On Day 1 *T* contributed visibly less than *N* and *C*. Hence, practice with the non-preferred hand showed less changes in *T*, exploration of the execution space.

To better understand this result Figure 11 illustrates one participant with such a pattern of result. The dark symbols represent a subset of bounces from one early trial that shows visibly more scatter than the later trial represented by hollow symbols. The average location of this early trial changes and the data migrate to the apparently preferred location of the late trial. At this location the data are more clustered and align with the manifold leading to better results. The component tolerance quantifies this migration in execution space. As the summarized data show participants who practice with their dominant hand first tend to search the execution space more leading to better performance towards the end of the practice.

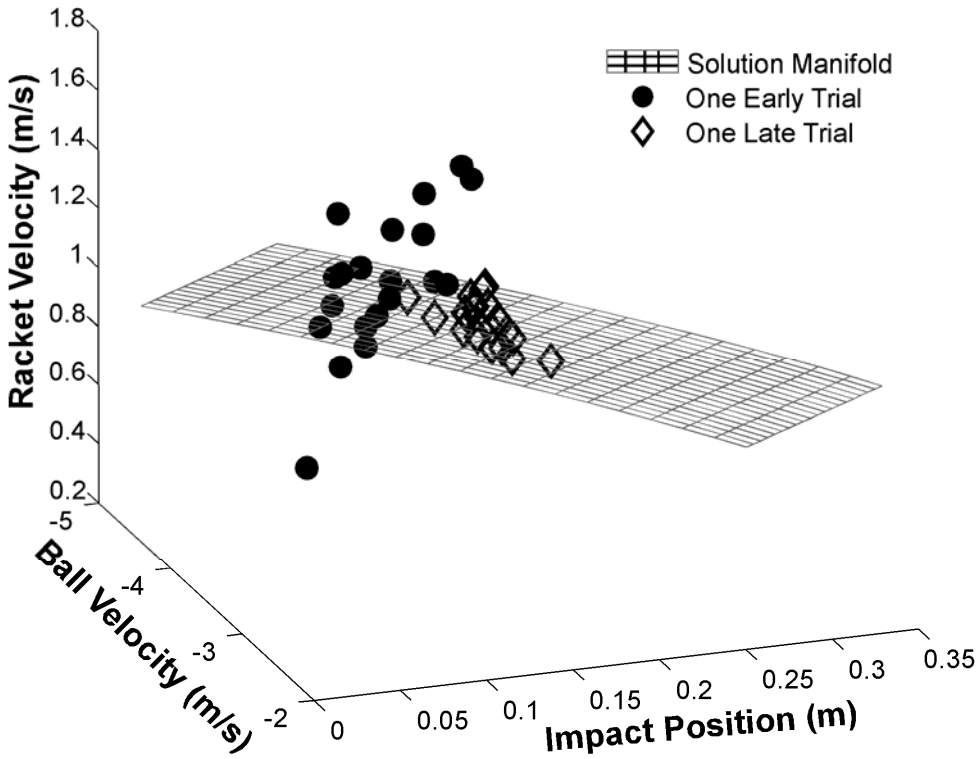

Figure 11: Execution space with a solution manifold. The dots represent individual bounces from an exemplary subject. The dark symbols represent a subset of bounces from one early trial and the hollow symbols represent the later trial.

Discussion

The objective of the present study was to examine learning and transfer of a perceptual-motor skill and to better understand *what* is learnt and *what* is transferred given the apparent improvements seen in the result variables. The task of rhythmic ball bouncing served as the experimental vehicle as it affords an independent investigation of stability and variability to assess the process of acquisition and interlimb transfer.

Results on the typical performance variables, standard deviation of error, clearly and not surprisingly demonstrated that participants improved their skill and performed with less variability. Regardless of hand, both groups showed evidence that they improved their performance with practice to the same degree. Practicing with the

preferred hand showed the same improvement as practice with the non-preferred hand. This has been evidenced by a host of previous research over many decades (see Schmidt, 2005). Further, this decrease in variability was also transferred without decrement to the other hand and this bilateral transfer was symmetrical to both the preferred and the non-preferred hand. On the second day this improved skill was maintained and no further change was observed. Inspection of the standard deviations of ball height showed a trend that the group that practiced with their preferred hand acquired a slightly lower level of variability. However, these results were not significant. Further, there were some indications that this level of skill was not completely transferred to the non-preferred hand. Again, these results were not significant due to high inter-individual variability.

Typically this reduction of variability is equated with an increase in stability, a characteristic in skilled performance. However, this simple equation is not always technically correct and is at best an operationalization of the concept. By using ball bouncing as the experimental task, we aim to examine the two aspects separately, with the goal to obtain more insight into learning and transfer. As previous studies already showed ball bouncing lends itself to a separate analysis of stability independent of variability of performance measure. Based on linear stability analyses of a mathematical model of the task criteria for dynamical stability were derived (Dijkstra et al., 2004; Schaal et al., 1996; Sternad et al., 2001). The racket acceleration at the moment of contact was shown to be an indicator of performance stability. Specifically, negative racket acceleration at the moment of impact indicates stable performance.

Using this independent measure of stability, the data showed that participants significantly changed their performance strategy. Starting with positive racket acceleration on Day 1, both groups showed an exponential change in this quantity towards more negative values, characteristic for passively stable solutions. This parallel change in variability and racket acceleration was already shown in few subjects (Sternad et al., 2000; Dijkstra et al. 2004). However, the previous study was performed in a real experimental set-up without the virtual display. Hence, it is important to see that the same change occurred when the task was performed in a

virtual set-up. It should be emphasized that this solution is not trivial and participants can also perform the task with positive accelerations. In fact, from a biomechanical point of view, solutions with appositive or zero acceleration would be expected, as at this moment of maximum velocity the maximum ball amplitude would be obtained. The parallel learning curves are a strong support for the interpretation that the measure acceleration is indicative of a strategy. The strategy was discussed in more detail in other investigations where perturbations were applied to test for the predicted passive (Wei, Dijkstra, & Sternad, submitted-a, submitted-b).

Comparing the exponential changes of both variability of the error score and acceleration with practice it is noteworthy that variability reached a plateau at the end of Day 1 and there was no more change seen on Day 2. This fast asymptotic learning was observed for both hands. In contrast, the exponential change in racket acceleration continued on the second day. This continued reorganization was seen in both groups, indicating that this effect was not dependent on the performing hand. Argued differently, the absence of hand effects were probably due to the presence of high interindividual variability. Hence, even if this variability may suppress such laterality effects, it did not prevent large effects in these variables. The continued change in acceleration on day 2 and absence of change in variability strongly indicates a reorganization in the underlying strategy that was no longer reflected in the observable measure.

Performance Differences between Hands. It is commonly acknowledged that dominant and non-dominant hands differ in their degree of adeptness, although the underlying causes are still little understood. The current results show surprisingly little evidence for different skill in performance. The comparison between both groups at the pre-test did not reveal differences between hands, neither in variability nor acceleration. This could be either because interindividual variability was too large or this task was novel to all participants. Closer inspection of individual data did not reveal any further hints at hand differences. What did we do? Is there any evidence that hand dominance becomes more pronounced with more practice? Conversely, the task is special because control is effectively confined to one moment when the racket

hits the ball. However, playing tennis clearly reveal hand differences.

The most likely explanation is that at the initial stage hand differences are not pronounced. One more piece of evidence points to this explanation. The data of group 2 that practiced with their dominant hand on the second day reached slightly better performance than group 1 practicing with their non-dominant hand. Hence, it appears that the dominant hand continues to improve. Further, the change back to the non-dominant hand in the transfer phase reveals that this change was indeed specific to the dominant hand as it could no longer be transferred to the non-dominant hand. This finding is visible in both dependent measures and contrast with the results of group 1.

The variability decomposition provides further consistent results. The most striking result was that during the first day of practice noise and tolerance steadily contribute in the group that performs with their preferred hand. In contrast, practicing first with the non-preferred hand reveal that only noise contributes to performance, not tolerance. As explained above, tolerance is an operationalization of migration and exploration. Therefore, it can be concluded that the dominant hand appears to explore execution space more, potentially to find the best solution and strategy. Hence, the slightly increasing performance for the dominant hand may be attributed to having fond a better location in execution space.

The seamless continuation of the three components without any interruption when changing from one hand to the other is one further strong indication that the transfer of skill across limbs was very smooth. This holds for both the first and second day of both groups. Covariation did not contribute very much because the task had a very planar manifold with the iso-error bands symmetrical and unchanging across locations.

Acknowledgements

This research was supported by grants from the National Science Foundation BCS-0450218, the National Institutes of Health R01 HD045639, and the Office of Naval Research N00014-05-1-0844.

Appendix

For a trial with n bounces, each bounce i (i = 1, .. n) is described by the vector $\boldsymbol{e_i}$ = (X_{IP}, V_B, V_R), where X_{IP} is the impact position, V_B the ball velocity immediately before impact, and V_R is the racket velocity at impact. Thus, all bounces from a trial produce a matrix of execution variables $\boldsymbol{E}$ with a dimension of $n \times 3$. Each row of this matrix corresponds to one bounce and one data point in execution space. Calculating the resulting ball amplitude and the resulting error from the target line according to Equation 4, $\boldsymbol{E}$ produces a vector with n values of *AE*. As variability of *AE* is regarded as the primary measure of performance, the standard deviations of *AE* within one trial are calculated.

Using two fictive trials, the five panels in Figure 5 illustrate the computational steps for the TNC-decomposition. Panel A represents one trial at the early stage of learning, panel E at the late stage of learning. The five data sets in panels A-E are labeled as $\boldsymbol{E_A}$, $\boldsymbol{E_B}$, $\boldsymbol{E_C}$, $\boldsymbol{E_D}$ and $\boldsymbol{E_E}$ for clarity. The average result of *SD-AE* of $\boldsymbol{E_A}$ is worse than $\boldsymbol{E_E}$ as it occurs earlier in the learning process.

Step 1: Permutate each column of $\boldsymbol{E_A}$ and $\boldsymbol{E_E}$ separately. This permutation rearranges the data in the three execution variables such that a different result is obtained for each bounce. From these new results *SD-AE* are calculated. The permutations of Sets A and E are shown by $\boldsymbol{E}_B$ and $\boldsymbol{E}_D$, respectively.

Step 2: The mean locations of $\boldsymbol{E_B}$ and $\boldsymbol{E_D}$ are calculated by taking the means of each execution variable, separately. Move $\boldsymbol{E_B}$ to the mean location of $\boldsymbol{E_D}$ by adding the distance between $\boldsymbol{E_B}$ and $\boldsymbol{E_D}$ for each execution variable to $\boldsymbol{E_B}$. This step creates the new data set $\boldsymbol{E_C}$.

Step 3: Calculate *SD-AE* for each set by applying Equation 4.

Step 4: Calculate the following differences in *SD-AE* to obtain the three components:

- The difference between $\boldsymbol{E_E}$ and $\boldsymbol{E_D}$ is called *Covariation*, *C*, which quantifies the performance improvement due to co-varying execution variables in a

goal-relevant way. Covariation can also be calculated between $\boldsymbol{E_A}$ and $\boldsymbol{E_B}$.

- The performance difference between $\boldsymbol{E_B}$ and $\boldsymbol{E_C}$ is called *Tolerance*, *T*, which quantifies the performance improvement obtained from a better location that is more tolerant to errors or noise. The only difference between $\boldsymbol{E_B}$ and $\boldsymbol{E_C}$ is the location in execution space.
- The performance difference between $\boldsymbol{E_C}$ and $\boldsymbol{E_D}$ is termed *Noise*, *N*, which quantifies performance improvement due to decreased scatter in the data. Note that both $\boldsymbol{E_C}$ and $\boldsymbol{E_D}$ are at the same location in execution space but have no covariation.

References

de Rugy, A., Wei, K., Müller, H., & Sternad, D. (2003). Actively tracking "passive" stability. *Brain Research, 982*(1), 64-78.

Dijkstra, T. M. H., Katsumata, H., de Rugy, A., & Sternad, D. (2004). The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies, 11*(3), 319-345.

Imamizu, H., Uno, Y., & Kawato, M. (1995). Internal representations of the motor apparatus: Implications from generalization in visuomotor learning. *Journal of Experimental Psychology, 21*(5), 1174-1198.

Katsumata, H., Zatsiorsky, V., & Sternad, D. (2003). Control of ball-racket interactions in the rhythmic propulsion of elastic and non-elastic balls. *Experimental Brain Research, 149*, 17-29.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinions in Neurobiology, 9*(6), 718-727.

Kelso, J. A. S., & Zanone, P. G. (2002). Coordination dynamics of learning and transfer across different effector systems. *Journal of Experimental Psychology: Human Perception and Performance, 28*(4), 776-797.

Müller, H., Frank, T., & Sternad, D. (2007). Variability, covariation and invariance with respect to coordinate systems in motor control. *Journal of Experimental Psychology: Human Perception and Performance, 33*(1), 250-255.

Müller, H., & Sternad, D. (2003). A randomization method for the calculation of covariation in multiple nonlinear relations: Illustrated at the example of goal-directed movements. *Biological Cybernetics, 39*, 22-33.

Müller, H., & Sternad, D. (2004). Decomposition of variability in the execution of goal-oriented tasks – Three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance, 30*(1), 212-233.

Riccio, G. E. (1993). Information in movement variability about the quantitative dynamics of posture and orientation. In K. M. Newell & D. M. Corcos (Eds.),

Variability and motor control (pp. 317-357). Champaign, IL: Human Kinetics.

Schaal, S., Sternad, D., & Atkeson, C. G. (1996). One-handed juggling: A dynamical approach to a rhythmic movement task. *Journal of Motor Behavior, 28*(2), 165-183.

Schmidt, R. A. (1975). A schema theory of discrete motor skill learning. *Psychological Review, 82*(4), 225-260.

Schmidt, R.A., & Lee, T.D. (2005). *Motor Control and Learning: a Behavioral Emphasis*. Champaign, IL: Human Kinetics.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2000). Dynamics of a bouncing ball in human performance. *Physical Review E, 63*, 011902-011901 - 011902-011908.

Sternad, D., Duarte, M., Katsumata, H., & Schaal, S. (2001). Bouncing a ball: Tuning into dynamic stability. *Journal of Experimental Psychology: Human Perception and Performance, 27*(5), 1163-1184.

Wei, K., Dijkstra, T. M. H., & Sternad, D. (submitted-a). Fluctuations and passive stability in a ball bouncing task. *Journal of Neurophysiology*.

Wei, K., Dijkstra, T. M. H., & Sternad, D. (submitted-b). Passive stability and active control in a rhythmic task. *Journal of Neurophysiology*.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science, 269*, 1880-1882.

Zanone, P. G., & Kelso, J. A. S. (1992). Evolution of behavioral attractors with learning: Nonequilibrium phase transitions. *Journal of Experimental Psychology: Human Perception and Performance, 18*(2), 403-421.

Zanone, P. G., & Kelso, J. A. S. (1997). Coordination dynamics of learning and transfer: Collective and component levels. *Journal of Experimental Psychology: Human Perception and Performance, 23*(5), 1454-1480.

CHAPTER 7

General Discussion

7.1. Hard and Soft Constraints to Human Performance

At first sight human behavior appears to have no bounds: the number of variations and the complexity of patterns that humans can exhibit appear limitless. Even simple tasks such as opening a door or climbing stairs are performed differently every time. Dancers and musicians repeatedly show us that seemingly impossible movements are indeed possible. And yet, despite this rich repertoire of observed behaviors there are numerous constraints on human and, in fact, all biological movements. Such constraints arise from the physical conditions of the environment and the mechanics of the human body, from the specific limits of the nervous system, and finally from the task that the actor is engaged in. Some of these constraints are hard such as physical laws that cannot be violated. Other constraints are soft and do not preclude a certain behavior but impose disadvantages or costs.

Examples for hard constraints set by the physical conditions of the environment and the body are many. Moving in the gravitational field imposes limits that cannot be overcome. The fact that a normal sized humans will never run 100 meters in 5 seconds and will never walk faster than 3 m/s is due to physical constraints set by the muscular-skeletal properties of the human body and gravity. For example only as long as the upward acceleration during a step is smaller than the gravitational pull humans can walk otherwise they fly off the ground, as when walking on the moon (Alexander, 1992). On the other hand, walking speed is also subject to soft constraints: when walking casually, humans will most likely adopt a stepping frequency that is close to the eigenfrequency of the legs as determined by their inertial properties (Holt, Hamill, & Andres 1990, 1991). Although humans can certainly walk faster and slower than this preferred speed, there will be a cost such as increased metabolic rate, higher variability, or inferior accuracy if the task involves accuracy.

Both hard and soft constraints are also set by the nervous system. The delays in information transmission due to the limited transmission speed of nervous signals are one obvious example. A similar hard constraint is the asymmetry of the human body. The left and right cerebral hemispheres have a different morphology and, concomitantly, different functional specializations. In human behavior this is seen in hand-dominance that leads to different preferences and levels of skill in the dominant and non-dominant hands. People will use their preferred hand to write or to brush their teeth unless they are required to do otherwise. More subtle differences in functional specialization have been revealed in the kinematics of dominant and non-dominant performance in visually guided reaching tasks (Sainburg, 2005). Relatively hard constraints are also seen in bimanual coordination where it has been demonstrated that the two hands have a tendency for synchronization in both rhythmic but also non-rhythmic behaviors (Cohen, 1970; Kelso, Southard, & Goodman, 1979). Well-known examples for intrinsically preferred modes include the rhythmic coordination in inphase and antiphase between two oscillating limbs (Yamanishi, Kawato, & Suzuki, 1980; Kelso, 1984).

A less obvious example for a soft constraints from the nervous system and information processing capabilities is the finding that when reaching from one target to another humans move in an almost straight line with a bell-shaped velocity of the endeffector trajectory (Morasso, 1981). Imposing visuo-motor transformations from joint angle trajectories to the perceived path Flanagan and Rao (1995) showed that visually guided reaching movements remain to be generated to produce a straight-line path in perceived space, even if the generated movement itself becomes nonlinear (Wolpert, Ghahramani & Jordan, 1995). The movement path obviously need not be straight but in the absence of other requirements humans tend to choose so.

Another source of soft constraints that has been less recognized is presented by the stability properties of the task. Following Warren (2006) a task is defined over the agent and its environment that constitute a pair of dynamical systems that are coupled mechanically and informationally. Their interactions give rise to the task dynamics, “a

vector field with attractors that correspond to stable task solutions, repellers that correspond to avoided states, and bifurcations that correspond to behavioral transitions" (Warren, 2006; Saltzman & Kelso, 1987). Research on robot locomotion has highlighted the central role of such stability properties by demonstrating that, when properly tuned, the rigid-body system can walk in a dynamically stable manner even without actuators (McGeer, 1990; Collins et al., 2005). The human-like behavior of such simple purely passive machines has strongly suggested that similar principles underlie human walking. A similar strategy was taken by Jin and Zacksenhouse (2003) in their approach to the task of playing yo-yo. As the task was analyzed to be open-loop unstable, an oscillatory controller with a closed-loop strategy was developed to perform the task in a stable manner. These few examples demonstrate the important role of constraints and stability as defined by the task for understanding movement. Yet, to our knowledge there are not many studies that have attempted to understand how task constraints shape human motor control.

7.1.1. Constraints and Stability in the Ball Bouncing Task

In this wake, the objective of this thesis was to develop a quantitative approach to shed light on the control and coordination of humans in the context of a task. The strategy was to first model the dynamics of the task and then derive predictions about stability as the preferred solutions for humans. The focus was not on the limits of behavior but on the soft constraints as presented by the stability of the task. The hypothesis was that the task dynamics is a fundamental determinant that shapes the movement pattern and action strategies that are available to the actor. This proposition is consistent with the claim of Newell (1986) that coordinated movements are determined by three kinds of constraints emanating from the environment, the organism, and the task. Note that this proposition does not mean that movement solutions are "specified" but rather that some are less advantageous and less efficient.

With this objective the first step was to develop a mathematical model of the task. In earlier work Sternad and colleagues had derived a simple model which captured the mechanical interactions between the ball and the oscillating surface (Sternad et al.,

2001; Dijkstra et al., 2003). It is important to point out that the bouncing ball model does not include properties of the actor and the nervous system. In the model the racket movements follow an invariant sinusoidal waveform. However, it is assumed that the actor becomes part of the system by moving the racket and phasing it with respect to the ball trajectory. The question is then how the actor chooses to perform these racket movements when being part of this task system. Based on a stability analysis of the model the conditions for stable behavior were derived. The central hypothesis was that actors prefer modes of behavior that satisfy these conditions. In experiments that simulated the simplified assumptions of the model human behavior was compared with these predictions. In sum, the strategy in this line of research was to examine the extent to which the physics of the task shapes and constrains the selection of the actor's behavior.

The physical model of the bouncing ball is based on few assumptions: ballistic flight of the ball, instantaneous collision with the racket with loss of energy defined by the coefficient of restitution, sinusoidal movements of the racket, a mass of the racket that is significantly greater than the mass of the ball such that there is no rebound of the racket de to the impact (Sternad et al., 2001; Dijkstra et al., 2004). There is no coupling from the ball trajectory to the racket movement, nor any other neuromuscular aspects of how the racket movements are generated. Hence, the model is entirely open-loop or "passive". However, this "lifeless" system exhibits a stable period-1 attractor in which the ball is bounced with invariant ball amplitudes, i.e. rhythmically. This stable attractor is obtained if the racket trajectory impacts the ball at a moment where the racket is in its upward decelerating phase. The stability analyses have derived quantitative predictions about the range of the racket accelerations for which the system is stable. Stable behavior in accordance with passive dynamics means that small perturbations of the ball trajectory converge back to the steady state without explicit error corrections of the racket. This stability was termed "passive stability" as no active control of the racket trajectory was included. The importance of the presence of such a stable attractor is that small fluctuations in

performance dissipate without having to explicitly correct for these errors by changing the racket movements (Dijkstra et al., 2004; Schaal et al., 1996; Sternad et al., 2000, 2001). This strategy is advantageous in terms of control effort as theoretically there is no need to correct for small errors and an invariant racket movement suffices for producing stable ball amplitudes.

7.1.2. Evidence of Sensitivity to Passive Dynamics of the Task

A series of empirical studies on ball bouncing verified that experienced performers are indeed sensitive to the stable solutions and bounce the ball with negative racket accelerations. The series of four experiments in this thesis continued to provide more fine-grained support for the central role of passive stability in this simple model task. Previously, the experimental support was obtained using a real racket and ball in three different variations in 1D and 3D. Given that the four experiments were conducted in a newly designed virtual set-up it was important to first replicate that actors performed the task with negative acceleration, i.e., in line with the stability predictions from the model.

All experiments provided such replication for a range of different coefficients of restitution. In Experiment 1 when no perturbations were applied, experienced participants consistently bounced the ball with negative racket accelerations. More fine-grained support was seen in the data when the ball was perturbed to either overshoot or undershoot the target: the racket accelerations at successive impacts remained largely unchanged at negative values. This was in contrast to the significant changes in other kinematic variables of the racket that deviated from the stationary values over bounces immediately following the perturbations. This indicated that actors changed their racket trajectory to accommodate for the perturbations but simultaneously maintained the ball-racket impacts at negative accelerations to provide conditions for stability and thereby facilitate return to steady state performance.

In Experiment 2, the stability characteristics of the task were changed parametrically by setting the coefficients of restitution α to seven different values (from 0.3 to 0.9). While the conditions for the period-1 stable attractor remained the

same - racket accelerations were negative -, its basin of attraction changed across α conditions. Irrespective of α, participants hit the ball at a phase with negative acceleration. Evidence that also more subtle stability characteristics mattered was seen in the fact that performance variability changed. Concomitantly, covariation among execution variables also changed. The less stable the solution (as indicated by higher values of the largest engenvalue), the more covariation - or control - was exerted. These results demonstrated that participants were sensitive to very subtle aspects of the task dynamics.

Experiment 3 provided support for the passive stability hypothesis by showing that during practice novice performers changed their strategy from an unstable solution characterized by positive accelerations to a stable one with consistently more negative accelerations. This exponential change was accompanied by a similar decrease in the variability of performance, confirming that this change in strategy was consistent with improvements in performance, although causality cannot be attributed to this correlation. It may be worth pointing out that actors were not aware of these changes in their strategy and only strove for decreasing the deviations in ball height and thereby variability in their performance. Another interesting observation was that after introducing the distortion to the execution space the impact characteristics changed in a systematic manner. The majority of actors exhibited large negative accelerations that gradually returned to the smaller pre-distortion values.

Experiment 4 strengthened these observations. Using a similar scenario with a practice session of 40 trials, a group of different actors replicated the findings of Experiment 3. Further, this behavior was observed regardless of whether the dominant or non-dominant hand performed the task. The central result of this experiment was that this strategy was not learnt and specific to the performing limb but rather was more abstract in nature as it transferred to the other hand, almost seamlessly.

Taken together, the findings of the four experiments consolidated previous support for the central role of passive stability in a task. It can be safely stated now that actors are sensitive to stability properties and organize their movements

accordingly to utilize the passive stability offered by the task dynamics.

7.1.3. Essential and Non-Essential Variables

In a broader and more historical context, this thesis provides experimental contributions that conform with the requirements for a theory of skilled actions as laid out by Fowler and Turvey as early as 1978. In this programmatic article the authors integrated propositions by Gel'fand and Tsetlin (1962, 1971), Bernstein (1967), Greene (1969, 1972), and Gibson (1966, 1977), and provided a framework which has since influenced, implicitly or explicitly, a wide spectrum of experimental research. The fundamental conceptual claims for a theory of acquisition and performance of skilled activity are threefold: 1) The minimal system of analysis should be an event, encompassing both the actor and the environment as the necessary support for movements. 2) The level of description should be coarse-grained and compatible with the actor's self-description and the environment. 3) In forming a controllable system, or a coordinative structure, the actor identifies an "organizational invariant" by which the many degrees of freedom of the task are constrained. Gel'fand and Tsetlin (1971) proposed that, as such, the problem of coordination is "well-organized": the variables indigenous to the specific task can be partitioned into essential and non-essential ones. Actors need to first identify essential variables as they determine the topology of the solution and then tune the non-essential variables as they parameterize the space. With a view to the three requirements, we argue that the critical variable racket acceleration has the status of an essential variable in the sense of Gel'fand and Tsetlin (see also Sternad et al., 2001). 1) The task is extremely simple and centered around the event of racket-ball interactions. 2) The level of description is course-grained at spatio-temporal scale that matter for the agent. As stated before, more microscopic neurophysiological detailed are not included in this analysis. 3) The actors have exhibited a strategy that seems to concentrate on an organizational invariant.

7.1.4. Evidence of Passive Dynamics Combined Active control

Importantly, the passive stability arising from the task dynamics does neither prescribe behavior nor does it describe all aspects of behavior. As introduced above,

passive stability poses a soft constraint such that other solutions are realizable but will have a cost. Active control on the basis of perceptual information is evidently possible for human actors. Taking a look at the literature on dynamic walking again, a rich array of studies have extended the original pioneering work of McGeer. These include passive walkers in two and three dimensions without, but also with additional powering (Coleman & Ruina, 1998; Collins, Wisse & Ruina, 2001; Collins et al., 2005; Tedrake et al., 2004). While these walkers differ much in their morphology and actuation, they all exploit passive stability and, probably as a consequence, exhibit stable human-like walking pattern. For instance, the passive dynamic walking models can take on a variety of physical configurations and still maintain the feature of being passively stable. But actuators have added advantages: passive walkers can walk on level ground as energy loss is offset by simple means of energy input, and three-dimensional walkers can be built with lateral stability using sensory feedback to correct for destabilizing deviations (Bauby & Kuo, 2000; Donelan et al., 2004).

In a similar fashion it can be expected that performance with passive stability does not fully account for the bouncing action. Not only does passive stability not afford a lot of flexible behaviors, it is also unlikely that actors do not utilize visual and haptic information available to them. Evidence for the presence of active control was supplied in Experiments 1 and 2. In Experiment 1, the quantitative predictions about the response to perturbations were not matched. Actors reestablished steady-state performance generally much faster than the model predicted. This held for both large perturbations but also for very small perturbations. In principle if only open-loop stability were present the large perturbations should have destabilized performance and actors should have lost the pattern or established solutions other than period-1. This, however, was never observed and spoke to the fact that subjects perceived the unexpected deviations of the ball trajectory and modulated their racket movements accordingly. Importantly, analysis of the racket trajectories following even small perturbations also invoked some immediate changes in the racket trajectory. Theoretically, such small perturbations within the basin of attraction could have

converged to steady state due to passive stability without requiring changes in the racket movement. However, this would have taken many more bounces than was observed.

Analysis of the racket trajectory showed that the periods and amplitudes scaled with the perturbation magnitudes. One important finding worth highlighting is that the active adjustments of the racket kinematics also assisted the re-establishing of the necessary condition for exploiting passive stability. The adjustments in racket period and racket amplitude effectively changed the racket waveform. However, the impact accelerations remained unchanged at the same negative values as in the bounces leading up to the perturbations. In other words, the actor actively corrected the perturbation effect by making necessary changes and simultaneously keeps attuned to the passively stable regime. These findings suggest that instead of solely relying on the passive stability to automatically correct for errors the actor perceived the disturbance to the system and made quick adaptations that scaled to the magnitude and direction of the disturbance. This shows that the actor is not "locked" onto the attractor of the task dynamics but he/she can make quick adaptations according to task-relevant information.

Experiment 2 pursued a different strategy to scrutinize behavior for the signatures of passive stability and active control. The rationale for the experiment is that even when no explicit perturbations are applied and the system is at steady state, both the ball and the racket trajectory exhibit a certain level of fluctuations arising from either the intrinsic performance or externally from the environment. These small fluctuations are like very small perturbations and, assuming purely open-loop behavior, they die out due to the passive stability of the task system. To test whether these small fluctuations die out in a manner consistent with the model, the ball bouncing map was extended with a stochastic component of small amplitude. The time course of these small perturbations at multiple time scales was then examined in simulations and quantified in the covariance structure. In addition, stability analysis of the model revealed that the largest eigenvalues are different for the different

coefficients of restitution, with larger α having less stability than smaller α. Predictions were tested by comparing the simulated or analytically determined covariance functions with the data. Results showed that for small coefficients of restitutions a good match was obtained; larger α values, on the other hand, led to significant deviations from these model predictions. As these mismatches generally implied that the fluctuations equilibrated faster, it was inferred that this behavior could only be brought about by active control. This finding suggested that also during steady state behavior, for higher α conditions, the small fluctuations were actively compensated for.

To further evaluate the presence of active control for higher α conditions, the measure of covariation was introduced. It has long been recognized that for a task with redundant degrees of freedom variations between variables can compensate for each other to yield relative constancy in the result (Bernstein, 1967; Müller, 2001; Stimpel, 1933; Turvey & Carello, 1995). This compensation was captured by quantifying covariation between the variables in the execution space. Covariation between the three execution variables was calculated by the permutation method introduced by Müller and Sternad (2003), a method that could also capture potential nonlinear relations. In a previous study it was shown that covariation increased with improving skill level and it was argued that this measure is a signature of movement control (Müller & Sternad, 2003; Scholz, Schöner, & Latash, 2000). Based on this interpretation the present study showed that covariation indeed increased for higher α conditions where task stability decreased. These changes across α conditions were complemented by a declining variability in the outcome variables. This result supports the interpretation that the racket movement was controlled more consistently and more effectively when the stability of the task was reduced by increasing α. Taken together, these adaptations, shown as increase in covariation and reduction in variability of movements, demonstrate the flexibility of the perceptual-motor system under changing task conditions.

7.1.5. The Strategy: Modeling Hard Constraints to Reveal Soft Constraints for

Control

Studies on passive dynamics in the task have been rare in the research on human motor control. Instead, the common focus has been on human performance specifically the biomechanics and neurophysiology of the moving limbs without much consideration of the task constraints within which the human is operating in. The present thesis aimed to demonstrate the potential of this task-focused research strategy to provide insights into the fundamental problems of human behavior.

While few, there do exist other studies that can be viewed in this light. As already referenced several times, research on passive dynamic walking has attracted some attention, mostly because these simple walking robots appear so human-like. Note though that the passive biped robot that walks downhill without any actuators or feedback controllers is dependent on the meticulous parameterization of the physical linkages. This is similar to the parameterization of the ball bouncing model, which requires a periodic racket waveform with certain amplitudes and impact accelerations within a certain negative range. Another interesting parallel in both tasks is that both the walking robot and the bouncing ball system involve intermittent contacts that are the critical moments for controlling the movement pattern: the racket-ball contact is the only moment the actor can exert influence on the ballistic flight of the ball; the foot-ground contact is essential for adjustments between steps since the movement of the leg in the swing phase is largely passive as a free pendulum (Mochon & McMahon, 1980).

The passive dynamic walking studies also provide insights into the human control of locomotion. For instance, the original 2D passive dynamic walking biped developed by McGeer (1990) has rigid legs only and the common understanding was that an additional upper body will only make the control more complicated. Wisse and his colleagues added an upper body to investigate passive walking in 3D. Surprisingly it was shown that the upper body actually increased the fore-aft stability (Wisse, Schwab, & van der Helm, 2004). Similarly when adding two arms swinging in anti-phase the destabilizing yaw motion around the vertical axis of the body was

reduced (Collins et al., 2001). These studies show that the physical properties of the human body can actually facilitate the control of locomotion and the whole body dynamics has to be taken into consideration in order to understand the control of human walking.

Another task similar to the ball bouncing task that involves a dynamic model of the task is cascade juggling. The juggling of three or more balls is a highly constrained task where the cyclic movements of the hands and the objects are tightly coupled. It requires the juggler to throw and catch N balls rhythmically while the hands' movements are coordinated to the flight cycles of the ball with some constrained phase relationship. Based on the work by Claude Shannon (Horgan, 1990; Raibert, 1986), Beek and colleagues proposed that the ratio of the duration of the ball loaded on the hand divided by the complete hand movement cycle (k) was a key variable for temporal organization of juggling (Beek, 1988; Beek & Turvey, 1992; Beek & Santvoord, 1992, 1996; van Santvoord & Beek, 1996). This variable can be freely manipulated by the juggler to vary between 0 and 1. Following the principle of frequency locking and "tiling of time", the authors argued that $k = 0.75$ characterizes a stable solution for juggling (Beek, 1988, 1989; Beek and Turvey, 1992). This stable regime is offered by the task dynamics of juggling and is independent from the number of hands and balls involved. In real tasks, however, the opportunity to choose k for increasing N might decline due to the timing and spatial constraints. But k remains to converge to 0.75 regardless of the juggling frequency.

These predictions were confirmed by empirical studies. In a learning experiment novice participants showed a trend towards $k = 0.75$ over successive training sessions (Beek & Santvoord, 1992). However, for skilled jugglers, k was found to range between 0.54 and 0.83 with a mean of 0.71 (Beek, 1989). This result was interpreted as evidence for multiple possible mode locks other than $k = 0.75$ in the work space of juggling. As $k = 0.75$ characterizes the most stable mode, the ability to operate at values other than 0.75 was interpreted as a high degree of flexibility. Indeed, when severe task demands were imposed in five- and seven-ball juggling, even skilled

jugglers demonstrated the predominance of $k = 0.75$ (Beek & Turvey 1992).

Finally, also some studies on the control of posture have taken an analysis of the task as their entry into understanding control. Anticipatory activity of postural muscles just prior and during voluntary movement has been studied for several decades since Belenkii and colleagues (1967) reported that activation of postural muscles preceded the activation of muscles for a voluntary arm movement. However, the functional role of this anticipatory muscle action was unclear with two possible mechanisms under debate: postural muscle activation stabilizes the body's center of mass (COM) or, alternatively, it stabilizes the individual joints affected by the interaction torques arising from the incoming voluntary movement. Patla, Ishac, & Winter (2002) investigated voluntary arm movements during a standing posture. To assess the functional role of anticipatory activity, they based their reasoning on modeling upright posture as an inverted pendulum and taking the joint reactive moments as input and the COM trajectory as the output. They found that the simulated and the measured COM profiles were identical until about 200 ms after the arm movement's onset. This indicated that a mechanical model without muscle control could account for COM migration 200 ms after the voluntary arm movement was initiated. Thus, the anticipatory trunk muscle activation did not stabilize the COM. Therefore, they concluded that the COM excursion at the early stage of the voluntary movement was passive in nature and the anticipatory muscle activations at the joints were mainly for joint stabilization and not for the control of the COM.

The approach of studying the passive dynamics of the task opens a window to reveal the fundamental aspects in movement control. The rationale common to all these studies is that task dynamics is a central determinant that shapes the movement pattern and the strategy that are available to the actor. As elaborated above, the complex and ordered behavior of actors is constrained by the dynamics, both through its hard and soft constraints. Our study on the ball bouncing task has modeled the hard constraints, the physics of the racket-ball system, stability analyses have revealed the soft constraints that the nervous can or cannot obey. This proposition is similar in

spirit to Newell's tenet that coordinated movements are determined by the confluence of constraints from the environment, the organism, and the task. It is also consistent with the concept of synergy as a task-specific device assembled out of the many degrees of freedom of the human action system (Fowler & Turvey, 1978).

The strategy of passive dynamics has taken these tenets into a research strategy that quantifies the physics of the task or the event (Fowler & Turvey, 1978). While simplicity is warranted for practical reasons, the event should be modeled including the hard constraints from the basic morphology of the actor and the object manipulated, the environment the actor interacts with, and the essential physical interactions. Neural control involving properties of the neuromuscular substrate, feedback control mechanism and perception, all essential to biological systems, lives within these constraints.

7.2. Learning and Transfer of the Passive Strategy

The major goal of Experiments 3 and 4 was to study the acquisition, adaptation, and transfer of the skill of bouncing a ball on the background of the findings on passive stability and control in Experiments 1 and 2. Special emphasis was given to the process of tuning to the passively stable regime with practice by analyzing the essential variable racket acceleration at ball impact. In addition, changes in variability during practice, adaptation, and transfer were evaluated to complement and further understanding of the underlying processes. Importantly, while it has been common to evaluate stability as the inverse of variability, our approach presents two independent routes of evaluation of stability and variability. Changes in variability were scrutinized by decomposing it into three distinct functional components by using the TNC method (Müller & Sternad, 2003, 2004a, b; Müller, Frank, & Sternad, 2007). This decomposition relies on a parsing of the task into the redundant space of execution variables and the associated result variables. It thereby creates a novel space or reference frame in which performance of the task can be evaluated.

7.2.1. Evidence for Attunement to Passive Stability with Practice

Both Experiments 3 and 4 found that novice participants started their performance with positive impact accelerations that gradually decreased to negative values, indicating that they gradually became sensitive to passive stability. At the same time, primary measures of performance variability, standard deviations in ball height and the absolute error from the target line, decreased along with this tuning process. Hence, this correlation supports the interpretation that exploiting passive stability is beneficial and indeed facilitates stable rhythmic bouncing. Further, the fact that all participants showed positive acceleration values at the initial stage of practice suggests that utilizing negative impact acceleration is an acquired and not immediately intuitive strategy through learning. These results were consistent with previous findings although this experiment was conducted in a real set-up where participants manipulated a real ball instead of a virtual ball as in the current virtual setup (Dijkstra et al., 2004).

Experiment 3 went further and introduced a perturbation such that all bounces at the previously preferred region in the execution space were perturbed to examine the adaptive behavior of the actor. Using the virtual set-up it was possible to create systematic deviations to the previously preferred solutions of each individual such that the originally approximately planar solution manifold became distorted. After the expected initial deterioration in performance, all participants moved out of the distorted region and re-established stability in performance. More importantly, preference for continually utilizing passive stability was found in the abruptly changed task space. Four out of seven participants showed an abrupt drop in racket acceleration to more negative values upon the introduction of perturbation. Along with this adaptation they managed to make racket acceleration return to the pre-perturbation values that were in the optimal range as predicted by the model. The remaining three participants did not exhibit abrupt changes in their impact values, instead, their racket acceleration values remained in the negative range but with large variations. These results provide additional evidence that racket acceleration is indeed

an essential variable that the actor tries to control in a certain optimal range in order to exploit the passive stability. Applied perturbations can force the actor to abandon the originally preferred location in the task space. However, the actor re-organized the movement to displace to a new location in the task space while simultaneously re-establishing optimal impact characteristics.

Experiment 4 required participants to practice first with one hand and then switch to the opposite hand. No matter with which hand they started their practice a similar tuning to the passively stable regime was found. Participants started off with positive values of racket acceleration which then gradually decreased to negative values with practice. The learning rate, the amount of change during one practice session, and the final value of racket acceleration were all similar in both hands. More importantly, once actors learnt to exploit passive stability with one hand, they continued to use this strategy when they switched to the other hand. The racket acceleration values right before and after switching hands did not show any significant difference. This result suggests that utilizing dynamic stability of the task is an acquired abstract skill and does not depend on which effectors execute the motor task.

A second important finding in this experiment is that during the second day of practice, the performance measures did not show any further significant improvements. However, racket acceleration continued to show a decrease to more negative values. This observation documented that racket acceleration is a very sensitive variable, more sensitive to change than the traditionally used performance measures. This observation also highlights that with the continued practice the actor's strategy continued to change even though overt statistical measures no longer indicated improvement and signaled that actors have reached an asymptote. This issue of performance plateau and continued reorganization of underlying processes has been a topic for discussion since the studies of Bryan and Harter (1897, 1899). These early studies on telegraph language reported plateaus in their learning curves and interpreted them as a time for reorganization of the underlying processes. While subsequent studies had problems in replicating these results, the phenomenon

remained alive in the discussion of skill learning (Adams, 1987). The present results add one piece of evidence to this discussion.

7.2.2. Passive Stability and Automaticity

Recalling the original argument that better attunement to the passively stable regime will alleviate the necessity for active corrections of small errors in performance, the evident conclusion is that the skill can be performed with less attention, or more automatically. Automaticity of movement has long been recognized as a signature of skilled performance characterizing later stages of the learning process (James, 1890; Fitts, 1964). Fitts has introduced his much discussed three stages of learning: The first cognitive stage denotes a phase where considerable cognitive activity is required to determine the appropriate strategy for the task at hand. This is followed by the associative stage, where the individual starts to make more subtle adjustments. The final stage is labeled autonomous, where performance has become automatic and is less subject to interference from other simultaneous activity and attention can be directed to secondary issues. The latter stage evidently matches with the definition of performance with passive stability, i.e., the strategy where attention is no longer required to focus on error correction. Even though Experiments 3 and 4 did not investigate the automaticity of performance directly, the increasing tuning to the passively stable regime is consistent with the characterization of the third stage in Fitts' theory of learning.

7.2.3. Variability Decomposition Analysis in Execution Space

Besides the analysis of passive stability Experiments 3 and 4 also employed a novel method to analyze variability to shed light on the process of learning, adaptation, and transfer. The TNC analysis focused on the distributional properties of a set of data in execution space and how it changes with practice. Three conceptual components of learning that have been discussed in the literature for many years are quantified by a novel method that parses the overall distributional properties into three components. The first component tolerance captures the essence of what researchers have described as the exploratory stage in learning. In many contexts exploration has been

discussed as one part of the observed variability. For example, in postural control the function of body sway has been interpreted as partly exploratory (Riccio, 1993). In developmental studies, the seemingly random flailing of infants' limb movement has been interpreted as an exploration of their capabilities and work space. The amount of exploration is represented in the migration in execution space with respect to the solution manifold. The component tolerance quantifies exactly this. The second component covariation has already been introduced as the compensatory relationship between variables. The component C evaluates the effect of such covariation in terms of the result variable. Note covariation is not covariance and does not use the linear statistical method of calculation. Lastly, the decrease of stochastic noise is a commonly acknowledged factor in motor learning. Within the TNC method noise is calculated as the third component that account for the changes in the result.

Novel to the TNC method is that the variability of a set of data points is evaluated in terms of the difference in the result variable. Therefore, unlike the UCM method for example, which has the same objective, the method can be applied to an execution space spanned by variables of different dimensions. Note that in ball bouncing the execution variables have position and velocity as dimensions rendering the space non-homogenous and distances in different directions cannot be evaluated. Further, it is not the covariance of the execution variables that is analyzed with the tools of linear algebra, but rather the distributional properties and their effect on the result variable. Hence, the components tolerance, covariation, and noise can be quantified regardless of whether the solution manifold is linear or nonlinear. This is not very important in the ball bouncing task as the solution manifold is an approximately planar surface; however this is not the case in other tasks, for example skittles or darts (Müller & Sternad, 2004; Müller & Loosch, 1999).

7.2.4. Stages of Learning

During practice in the first sessions of the two experiments, respectively, the decrease in noise N was the leading contributor to the improvement in performance. Following an initial significant increase N reached a plateau towards the end of the

practice phase. The contributions from both tolerance *T* and covariation *C* were modest. The observed leveling off of *N* may be interpreted as indicating that a first stage in the learning process has been completed. This interpretation is further supported by the fact that the performance measures, standard deviation and mean of the absolute error, exhibit the same qualitative time course. Hence, the measure *N* indeed corresponds to the stochastic noise processes that are also quantified by the standard variability measures. While not providing much additional information, it is informative that the other two components showed so little change. Previous research on different tasks, skittles and dart throwing, had identified more significant contributions of *T* and *C* (Müller & Sternad, 2004; Müller & Loosch, 1999). This difference is most likely attributable to the fact that the solution manifold of the ball bouncing task was approximately planar with similar sensitivity at all locations. With a view to the previous central result on passive stability, it is noteworthy that racket acceleration had a comparable approximately exponential time course as *N* with a similar time scale.

With a view to extant formulations of stages in learning this profile may be correlated with the intermediate stage in the terminology of Turvey, Mitra, and colleagues (1997). In their proposition the intermediate stage is characterized by a reduction in active degrees of freedom as deduced from phase space embedding of an observed quantity. In this space the dimensionality of the attractor is estimated. Hence, somewhat akin to the TNC approach, the determination of active degrees of freedom also relies on an unpacking a space spanned by its embedded dimensions. The TNC approach begins with the overt result variables and spans the space of execution variables to examine control processes in a different representation. What is different between the two approaches is that Turvey's approach aims to estimate the number of dimensions of the underlying attractor, whereas our more mechanically grounded representation simply spans another space to obtain underlying variability. Yet, the results may be correlated as both representations aim to reveal the complexity and its changes during learning. It may hence be conjectured that our observation of

gradually reducing noise processes may be related to the postulated decrease in dimensionality at the intermediate stage.

7.2.5. Exploration in Adaptation and Transfer

Another central finding in Experiment 3 is that the actor showed sensitivity to the variability of the execution variables. This was seen when a distortion in execution space was applied by systematically changing the physical relations between execution and result in the previously preferred region. As performance suffered noticeably actor's execution migrated to a new location in execution space that was unperturbed and followed the original correct physical relations. This migration showed an interesting feature: the ellipsoid that included 95% of the variance within a given set of bounces moved completely out of the distorted region, but only until the ellipsoid no loner overlapped with the distorted ellipsoid. This illustrates that actors are aware of their variance. Further, instead of following random routes to move away from the perturbation ellipsoid, all participants showed a similar adaptation path. This constrained process of adaptation suggests that adaptation might be subject to task constraints such that performance within the novel solution manifold is still governed by the previously learned task dynamics. Additionally, the fact that actors just barely moved out of the perturbed region and not further, suggests that they not only change their performance (many other locations in execution space would have been possible), but they are also attracted by the previously preferred location in the task space. This resonates with the proposition that learning a new coordination pattern is governed by the interaction between initial coordination dynamics and the to-be-learned coordination pattern (Zanone & Kelso, 1992). The final location outside of the perturbation ellipsoid indicates the adaptation is an outcome of the compromise between the initial and the to-be-adapted patterns.

The adaptation and the transfer phases of the two experiments also provided more "opportunity" for the component tolerance to play a major role. In order to move outside the perturbed region, T increased substantially and became the most significant component. This is in accordance with the functional role of tolerance as

indicator of performance improvement due to moving to a better and more "tolerant" area in execution space (Müller, 2001; Müller & Sternad, 2004). While not unintuitive, the analysis highlighted this conception of exploration well acknowledged in the motor control literature. In Experiment 4, the component tolerance revealed a less obvious result: in the first session participants who performed the task with their dominant hand showed noticeably more changes in tolerance T than the group that used their non-dominant hand. Without any speculation this result signifies exploration. While many differences have been ascribed to dominant and non-dominant performance, the ability to explore has not yet been discussed.

Taken together, these results on three variability components suggest that variability decomposition is a valuable method for capturing the learning and adaptation in the task space.

7.3. Conclusions

The present dissertation examined a ball bouncing task in a virtual reality environment. Using this toy task as an experimental vehicle, questions about human control strategies, their acquisition and transfer were investigated. Control and coordination was found to be constrained by the task dynamics as the human actor exploits the stability properties of the task. This statement was based on the quantitative analysis and testing of a mechanical model that captured the essence of the task. At the same time, the actor is able to flexibly utilize the behavioral information when perturbations were applied to the system or when the system stability was altered by varying parameter settings. This blend of control, simultaneously conforming to passive stability but also applying active control guided by perceptual information renders the actors' movements stable and. Building a passive model without perceptual feedback has been proven to be a viable methodological window to unravel the contributions of passive dynamics and neural control. Experiments on learning and adaptation further supported that exploiting passive stability is indeed an acquired strategy. It is also suggested to be abstract

knowledge of the system dynamics as it is symmetrically transferable between two limb systems. Compared to the mere examination of result or outcome measures, variability decomposition method has been shown to provide more insights about processes of learning, adaptation, and transfer.

REFERENCES

Adams, J. A. (1987). Historical review and appraisal of research on the learning, retention, and transfer of human motor skills. *Psychological Bulletin, 101*, *1*, 41-74.

Bauby, C. E., & Kuo, A. D. (2000). Active control of lateral balance in human walking. *Journal of Biomechanics, 33*, *11*, 1433-1440.

Beek, P.J., Peper, C.E., & Stegeman, D.F. (1995). Dynamical models of movement coordination. *Human Movement Science*, *14, 4*, 573-608.

Belenkii, V.Y., Gurfinkel, V.S., & Paltsev, Y.I. (1967). Elements of control of voluntary movements. *Biophysics*, *12,* 154-161.

Bernstein, N.A. (1967). *The coordination and regulation of movements.* London: Pergamon Press.

Bernstein, N.A. (1996). On dexterity and its development. In M.L. Latash & M.T. Turvey (Eds.), *Dexterity and its development* (pp.3-244). Mahwah, NJ: Lawrence Erlbaum.

Bryan, W. L., & Harter, N. (1897). Studies in the physiology and psychology of the telegraphic language. *Psychological Review, 4*, 27-53.

Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review, 6*, 345-375.

Bühler, M., Koditschek, D.E., & Kindlmann, P.J. (1994). Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research, 13*, 101-118.

Cohen, L. (1970). Interaction between limbs during bimanual voluntary activity. *Brain, 93*, 253-272.

Coleman, M. J., & Ruina, A. (1998). An uncontrolled walking toy that cannot stand still. *Physical Review Letters, 80*, 3658-3661.

Collins, S.H., Ruina, A., Wisse, M., & Tedrake, R. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science*, *307*, 1082-1085.

Collins, S.H., Wisse, M., & Ruina, A. (2001). A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research, 20, 7*, 607-615.

Corbetta, D., & Thelen, E. (1996). The developmental origins of bimanual coordination: a dynamic perspective. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 502–522.

Darling, W.G., & Cooke, J.D. (1987). Changes in variability of movement trajectories with practice. *Journal of Motor Behavior*, *19, 3*, 291-309.

Donelan, J. M., Shipman, D. W., Kram, R., & Kuo, A. D. (2004). Mechanical and metabolic requirements for active lateral stabilization in human walking. *Journal of Biomechanics, 37,6*, 827-835.

de Rugy, A., Wei, K., Müller, H. & Sternad, D. (2003). Actively tracking 'passive' stability in a ball bouncing task. *Brain Research, 982,* 64-78.

Dijkstra, T.M.H., Katsumata, H., de Rugy, A. & Sternad, D. (2004). The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies*, *11, 3*, 319-345.

Dijkstra, T.M.H., Schöner, G., Gielen, C.C.A.M. (1994a). Temporal stability of the action-perception cycle for postural control in a moving visual environment. *Experimental Brain Research, 97*, 477–486.

Dijkstra, T.M.H., Schöner, G., Giese M.A., Gielen, C.C.A.M. (1994b). Frequency

dependence of the action-perception cycle for postural control in a moving visual environment: relative phase dynamics. *Biological Cybernetics, 71*, 489–501.

Dingwell, J.B., Mah, C.D., & Mussa-Ivaldi, F.A. (2004). Experimentally confirmed mathematical model for human control of a non-rigid object. *Journal of Neurophysiology, 91*, 1158-1170.

Faugloire, E., Bardy, B.G., Stoffregen, T. A., (2006). Dynamics of learning new postural patterns: Influence on preexisting spontaneous behaviors. *Journal of Motor Behavior*, *38, 4*, 299-312.

Fitts. P. (1964). Perceptual-motor skill learning. In: Melton, A.W. (Ed.). *Categories of Human Learning* (pp. 244-285). New York: Academic Press.

Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience, 5, 7*, 1688-1703.

Flanagan, J. R., & Rao, A. K. (1995). Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space. *Journal of Neurophysiology, 74*, *5*, 2174-2178.

Fowler, C., & Turvey, M.T. (1978). Skill acquisition: An event approach with special reference to searching for the optimum of a function of several variables. In: Stelmach, G. (Ed.). *Information Processing in Motor Control and Learning* (pp. 1-40). New York: Academic Press.

Gelfand, I. M., & Tsetlin, M. L. (1962). Some methods of control for complex sytems. *Russian Mathematical Surveys*, *36*, 307-319.

Gibson, J.J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J.J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.). *Perceiving, acting and knowing: Toward an ecological psychology* (pp.67-82). Hillsdale, NJ: Erlbaum.

Goldfield, E.C., Kay, B.A., Warren, W.H. (1993). Infant bouncing: the assembly and tuning of action systems. *Child Development*, *64*, 1128-1142.

Goodman, L., Riley, M. A., Mitra, S., & Turvey, M. T. (2000). Advantages of rhythmic movements at resonance: Minimal active degrees of freedom, minimal noise, and maximal predictability. *Journal of Motor Behavior, 32*, 1, 3-9.

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

Haken, H. (1983). *Synergetics: An introduction*. Berlin, Germany: Springer-Verlag.

Haken, H. (1996). *Principles of brain functioning: A synergetic approach to brain activity, behavior and cognition.* Berlin, Germany: Springer-Verlag.

Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transition in human hand movements. *Biological Cybernetics, 51*, 347-356.

Hasan, Z. (2005). The human motor control system's response to mechanical perturbation: Should it, can it and does it ensure stability? *Journal of Motor Behavior*, *37, 6*, 484-493.

Holt, K. G., Hamill, J., & Andres, R. O. (1990). The force-driven harmonic oscillator as a model for human locomotion. *Human Movement Science, 9*, 55-68.

Holt, K. G., Hamill, J., & Andres, R. O. (1991). Predicting the minimal energy costs of human walking. *Medicine and Science in Sports and Exercise*, *23, 44*, 491-498.

Hoyt, D. E., & Taylor, G. R. (1981). Gait and the energetics of locomotion in horses.

Nature, 292, 239-240.

Higgins, J. R., & Spaeth, R. K. (1972). Relationship between consistency of movement and environmental condition. *Quest, 17,* 61–69.

Imamizu, H., Miyauchi, S., Tamada T., Sasaki, Y., Takino, R., Putz, B., Yoshioka, T., Kawato, M., 2000. Human cerebellar activity reflecting an acquired internal moddel of a new tool. *Nature*, *403*, 153-4.

James, W. (1890). *The principles of psychology* (vol. 1). New York: Holt.

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

Jin, H.-L., & Zacksenhouse, M. (2003). Oscillatory neural networks for robotic yo-yo control. *IEEE Transactions on Neural Networks 14*, *2*, 317-325.

Jordan, M.I., & Rumelhart, D.E. (1992). Forward models: supervised learning with a distal teacher. *Cognitive Science*, *16, 3*, 307-354.

Jordan, J.W. & Smith, P. (1987). *Nonlinear ordinary differential equations* (2nd ed). New York: Oxford.

Kawato, M. (1990). Computational schemes and neural network models for formation and control of multijoint arm trajectory. In W. T. Miller Iii, R. S. Sutton & P. J. Werbos (Eds.), *Neural networks for control* (pp. 197-228). Cambridge, MA: MIT Press.

Keele, S.W. (1968). Movement control in skilled motor performance. *Psychological Bulletin*, *70*, 387-403.

Kelso, J.A.S. (1984). Phase transitions and critical behavior in human bimanual coordination. *American Journal of Physiology: Regulatory, Integrative and*

Comparative Physiology, 15, R1000-R1004.

Kelso, J.A.S. (1994a). Elementary coordination dynamics. In: S. Swinnen, H. Heuer, J. Massion and P. Casear (Eds.), *Interlimb coordination: Neural, dynamical and cognitive constraints* (pp: 301-318). New York: Academic Press.

Kelso, J. A. S., & Jeka, J. J. (1992). Symmetry breaking dynamics of human multilimb coordination. *Journal of Experimental Psychology: Human Perception and Performance, 18, 3*, 645-668.

Kelso, J.A.S., & Schöner, G. (1988). Self-organization of coordinative movement patterns. *Human Movement Science, 7*, 27-46.

Kelso, J.A.S., Southard, D.L., & Goodman, D. (1979). On the coordination of two-handed movements. *Journal of Experimental Psychology: Human Perception and Performance, 5, 2,* 229-238.

Kelso, J. A. S., & Zanone, P. G. (2002). Coordination dynamics of learning and transfer across different effector systems. *Journal of Experimental Psychology: Human Perception and Performance, 28*, *4*, 776-797.

Körding, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*, 244-247.

Kugler, P. N. (1986). A morphological perspective on the origin and evolution of movement patterns. In M. G. Wade & H. T. A. Whiting (Eds.). *Motor development in children: Aspects of coordination and control* (pp. 459-525). Dordrecht, The Netherlands: Martinus Nijhoff.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement*. Hillsdale, NJ: Erlbaum.

Kuo, A. (2002). The relative roles of feedforward and feedback in the control of

rhythmic movements. *Motor Control*, *6*, 129-145;

Latash, M.L. (1993). *Control of human movement.* Chicago, IL: Human Kinetics.

Liu, Y.-T., Mayer-Kress, G., & Newell, K. M. (2003). Beyond curve fitting: A dynamical systems account of exponential learning in a discrete timing task. *Journal of Motor Behavior, 35*, *2*, 197-207.

Liu, Y-T, Mayer-Kress, G., & Newell, K.M. (2006). Qualitative and quantitative change in dynamics of motor learning. *Journal of Experimental Psychology: Human Perception and Performance*. *32, 2,* 380-393.

Magill, R.A. (2001). *Motor learning and its applications*. (6th Eds.). Boston: MacGraw Hill.

McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research 9, 2*, 62–82.

Mitra, S., Amazeen, P.G., & Turvey, M.T. (1998). Intermediate motor learning as decreasing active (dynamical) degree of freedom. *Human Movement Science*, *17, 1*, 17-65.

Mochon, S., & McMahon, T.A. (1980). Ballistic walking: an improved model. *Mathematical Biosciences*. *52, 3-4*, 241-260.

Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research, 42*, 223-227.

Müller, H., Frank, T.D., & Sternad, D. (2004). Variability, covariation, and invariance with respect to coordinate systems in motor control: Reply to Smeets and Louw (2007). *Journal of Experimental Psychology: Human Perception and Performance. 33, 1*, 250-255.

Müller, H., & Loosch, E. (1999). Functional variability and an equifinal path of

movement during targeted throwing. *Journal of Human Movement Studies, 36*, 103-126.

Müller, H., & Sternad, D. (2003). A randomized method for the calculation of covariation in multiple nonlinear relations: Illustrated with the example of goal-directed movements. *Biological Cybernetics, 89*, 22-33.

Müller, H., & Sternad, D. (2004). Decomposition of variability in the execution of goal-oriented tasks: Three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance, 30*, 212-233.

Newell, K. (1986). Constraints on the development of coordination. In: Wade, M.G., Whitting, H.T.A. (Eds.). *Motor Development in Children: Aspects of Coordination and Control* (pp. 341-360). Martinus Nijhoff, Dordrecht.

Newell, K.M., Broderick, M.P., Deutsch, K.M., & Slifkin, A.B. (2003). Task goals and change in dynamical degrees of freedom with motor learning. *Journal of Experimental Psychology: Human Perception and Performance*. *29, 2*, 379-387.

Newell, K. M., Liu, Y.-T., & Mayer-Kress, G. (2001). Time scales in motor learning and development. *Psychological Review, 108*, *1*, 57-82.

Newell, A., & Rosenbloom, P.S. (1981). Mechanisms of skill acquisition and the law of practice. In J.R.Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1-55). Hillsdale, NJ: Erlbaum.

Patla, A.E., Ishac, M.G., & Winter, D.A. (2002). Anticipatory control of center of mass and joint stability during voluntary arm movement from a standing posture: interplay between active and passive control. *Experimental Brain Research. 143*, 318-327.

Port, R.F., & van Gelder, T. (Eds). (1995). *Mind as motion: Explorations in the*

dynamics of cognition. Cambridge, MA: MIT press.

Riccio, G. E. (1993). Information in movement variability about the quantitative dynamics of posture and orientation. In K. M. Newell & D. M. Corcos (Eds.), *Variability and motor control* (pp. 317-357). Champaign, IL: Human Kinetics.

Richardson, M. J., Marsh, K. L., & Schmidt, R. C. (2005). Effects of visual and verbal interaction on unintentional interpersonal coordination. *Journal of Experimental Psychology: Human Perception and Performance, 31*, *1*, 62-79.

Rose, D.J. (1997). *A multilevel approach to the study of motor control and learning*. Boston: Allyn and Bacon.

Rosenblum, L.D., & Turvey. M.T. (1988). Maintenance tendency in coordinated rhythmic movements: relative fluctuations and phase. *Neuroscience*, *27*, 289-300.

Rosenbaum, L.D. (1992). *Human motor control*. Cambridge: MIT Press.

Sainburg, R. L. (2005). Handedness: differential specialization for control of trajectory and position. *Exercise Sport Science Reviews, 33*, *4*, 358-370.

Saltzman, E., Kelso, J.A.S. (1987). Skilled actions: A task dynamic approach. *Psychological Review, 94*, 84-106.

Schaal, S., Atkeson, C.G. & Sternad, D. (1996). One-handed juggling: A dynamical approach to a rhythmic movement task. *Journal of Motor Behavior, 28*, 165-183.

Schmidt, R.A. (1975). A schema theory of discrete motor skill learning. *Psychological Review, 82*, *4*, 225-260.

Schmidt, R.A., & Lee, T.D. (2005). *Motor Control and Learning: a Behavioral Emphasis*. Champaign, IL: Human Kinetics.

Schmidt, R. C., Bienvenu, M., Fitzpatrick, P. A., & Amazeen, P. G. (1998). A

comparison of within- and between-person coordination: Coordination breakdowns and coupling strength. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 884-90

Schmidt, R. C., Carello, C., & Turvey, M. T. (1990). Phase transitions and critical fluctuations in the visual coordination of rhythmic movements between people. *Journal of Experimental Psychology: Human Perception and Performance, 16*, *2*, 227-247.

Schmidt, R.A., Zelaznik, H., Hawkins, B., Frank, J.S., & Quinn, J.T Jr. (1979). Motor-output variability: a theory for accuracy of rapid motor acts. *Psychological Review*, *47, 5*, 415-451.

Scholz, J.P. & Schöner, G. (1999). The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research, 126,* 289-306.

Scholz, J.P., Schöner, G. & Latash M.L. (2000). Identifying the control structure of multijoint coordination during pistol shooting. *Experimental Brain Research, 135, 382-404.*

Schroeder, M. (1991). *Fractals, chaos, power laws: Minutes from an infinite paradise.* New York: Freeman.

Shadmehr, R., & Wise, S. P. (2005). *Computational neurobiology of reaching and pointing: A foundation for motor learning*. Cambridge, MA: MIT Press.

Smeets, J. B. J. (2000). The relation between movement parameters and learning. *Experimental Brain Research, 132,* 550-552.

Spencer, J.P., Thelen, E., Spatially specific changes in infants' muscle coactivity as they learn to reach. *Infancy*, *1, 3*, 275-302.

Sternad, D., Duarte, M., Katsumata, H. & Schaal, S. (2001). Bouncing a ball: Tuning

into dynamical stability. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 1163-1184.

Stimpel, E. (1933). Der Wurf (On throwing). In F. Krüger & O. Klemm (Eds.), *Motorik (On motor control)* (pp. 109-138). München, Germany: Beck.

Tedrake, R., Zhang, T.W., Fong, M., Seung, H.S. (2004). Actuating a simple 3D passive dynamic walker. *IEEE International Conference on Robotics and Automation*, *5*, 4656-4661.

Thelen, E., Corbetta, D., Kamm, K., Spencer, J.P., Schneider, K., & Zernicke, R.F. (1993). The transition to reaching: mapping intention and intrinsic dynamics. *Child Development*, *64*, 1058–1098.

Thelen, E., Corbetta, D., & Spencer, J.P. (1996). Development of reaching during the first year: role of movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, *22, 5*, 1059–1076.

Thelen, E., & Smith, L.B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT press.

Tufillaro, N.B., Abbott, T., & Reilly,J. (1992). *An experimental approach to nonlinear dynamics and chaos*. Redwood City, CA: Addison-Wesley.

Turvey, M.T., Fitch, H.L., & Tuller, B. (1978). Part V: degrees of freedom, co-ordinative structures and tuning. In: Kelso, J.A.S. (Ed). *Human Motor Behavior : An introduction.* Erlbaum, New Jersey.

Turvey, M.M., Schmidt, R.C., Rosenblum, L.D., & Kugler, P.N. (1988). On the time allometry of coordinated rhythmic movement. *Journal of Theoretical Biology*, *130*, 285-325.

van Asten, W.N.J.C., Gielen, C.C.A.M., & Denier van der Gon, J.J. (1988). Postural

movements induced by rotations of visual scenes. *Journal of the Optical Society of America*, *5, 10*, 1781-1789.

Vorro, J. R. (1973). Stroboscopic study of motion changes that accompany modifications and improvements in throwing performance. *Research Quarterly, 44,* 216–226.

Warren, W.H. (2006). The dynamics of perception and action. *Psychological Review*, *113, 2*, 358-389.

Wisse, M., Schwab, A.L., van der Helm, F.C.T. (2004). Passive dynamic walking model with upper body. *Robotica*, *22*, 681-688.

Wolpert, D.M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, *11, 7-8*, 1317-1329.

Wolpert, D.M., Miall, R.C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*, 338-347.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). Are arm trajectories planned in kinematic or dynamic coordinates? *Experimental Brain Research, 103*, 460-470.

Worringham, C. J. (1991). Variability effects on the internal structure of rapid aiming movements. *Journal of Motor Behavior, 23,* 75–85.

Worringham, C. J. (1993). Predicting motor performance from variability measures. In K. M. Newell & D. M. Corcos (Eds.), *Variability and motor control* (pp. 53–63). Champaign, IL: Human Kinetics.

Yamanishi, J., Kawato, M., & Suzuki, R. (1980). Two coupled oscillators as a model for the coordinated finger tapping by both hands. *Biological Cybernetics, 37*, 219-225.

Yang, J. F., & Scholz, J. P. (2005). Learning a throwing task is associated with

differential changes in the use of motor abundance. *Experimental Brain Research, 163*, *2*, 137-158.

Yang, J. F., Scholz, J. P., & Latash, M. L. (2007). The role of kinematic redundancy in adaptation of reaching. *Experimental Brain Research, 176*, *1*, 54-69.

Zanone, P. G., & Kelso, J. A. S. (1992). Evolution of behavioral attractors with learning: Nonequilibrium phase transitions. *Journal of Experimental Psychology: Human Perception and Performance, 18*, *2*, 403-421.

Zanone, P. G., & Kelso, J. A. S. (1997). Coordination dynamics of learning and transfer: Collective and component levels. *Journal of Experimental Psychology: Human Perception and Performance, 23*, *5*, 1454-1480.

Vita: Kunlin Wei

EDUCATION

2000	**B.S in Biomechanics**, Beijing University of Physical Education, China
2002	**M.S in Kinesiology** (Motor Control), the Pennsylvania State University
2003	**M.S in Electrical Engineering** (Control Systems), the Pennsylvania State University
Anticipated May 2007	**Ph.D in Kinesiology**, the Pennsylvania State University

AWARDS

2006	Dissertation Award from the School of HHD at the Penn State Univ.
2004	Fellowship from the 8th International Conference on Cognitive and Neural Systems.
2004	Fellowship of the Santa Fe Institute for Complex Systems
2000-2002	University Fellowship of the Pennsylvania State University
2000	Fellowships from the Univ. of North Carolina, Univ. of Illinois at Urbana-Champaign, Univ. of Texas at Austin, and The Iowa State Univ. (all declined)

PUBLICATIONS

de Rugy, A., **Wei, K**., Müller, H., & Sternad, D. (2003). Actively tracking "passive" stability. *Brain Research*, *982, 1*, 64-78.

Wei, K., Wertman, G., & Sternad, D. (2003). Interactions between rhythmic and discrete components in a bimanual task. *Motor Control, 7,* 134-154.

Sternad, D., **Wei, K**., Ivry, R.B., & Diedrichsen, J. (2006). Rhythmic and discrete elements in bimanual coordination of healthy and split-brain subjects*. Experimental Brain Research*. Aug 18; [Epub ahead of print] PMID: 16917769.

Wei, K., Hogan, N., & Sternad, D. (in preparation). Effects of load on the interaction between discrete and rhythmic movements.

Boonstra, T.A., **Wei, K.**, & Sternad, D. (in preparation). Variability and stability during the acquisitionn and adaptation of a rhythmic task.

Wei, K., Dijkstra, T.M.H., & Sternad, D. (submitted). Passive stability and active control in a rhythmic task. *Journal of Neurophysiology*

Wei, K., Dijkstra, T.M.H., & Sternad, D. (in preparation). Passive stability and active control during steady-state performance.

Tlili, M., **Wei, K**., & Sternad, D. (submitted). Variability and stability in learning and transfer of a ball bouncing task.

Tlili, M., **Wei, K.,** & Sternad, D. (in preparation). Exploration strategies in a complex task with disjunct nonlinear solution manifolds.