

Advanced Robotics



ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/tadr20

Human control of complex objects: towards more dexterous robots

Salah Bazzi & Dagmar Sternad

To cite this article: Salah Bazzi & Dagmar Sternad (2020): Human control of complex objects: towards more dexterous robots, Advanced Robotics, DOI: 10.1080/01691864.2020.1777198

To link to this article: https://doi.org/10.1080/01691864.2020.1777198



Published online: 16 Jun 2020.



Submit your article to this journal 🖙



View related articles



🕖 View Crossmark data 🗹

FULL PAPER

Taylor & Francis

Check for updates

Human control of complex objects: towards more dexterous robots

Salah Bazzi ^{(Da,b} and Dagmar Sternad^{a,b,c}

^aDepartment of Biology, Northeastern University, Boston, MA, USA; ^bDepartment of Electrical and Computer Engineering, Northeastern University, Boston, MA, USA; ^cDepartment of Physics, Northeastern University, Boston, MA, USA

ABSTRACT

Manipulation of objects with underactuated dynamics remains a challenge for robots. In contrast, humans excel at 'tool use' and more insight into human control strategies may inform robotic control architectures. We examined human control of objects that exhibit complex - underactuated, nonlinear, and potentially chaotic dynamics, such as transporting a cup of coffee. Simple control strategies appropriate for unconstrained movements, such as maximizing smoothness, fail as interaction forces have to be compensated or preempted. However, predictive control based on internal models appears daunting when the objects have nonlinear and unpredictable dynamics. We hypothesized that humans learn strategies that make these interactions predictable. Using a virtual environment subjects interacted with a virtual cup and rolling ball using a robotic visual and haptic interface. Two different metrics quantified predictability: stability or contraction, and mutual information between controller and object. In point-to-point displacements subjects exploited the contracting regions of the object dynamics to safely navigate perturbations. Control contraction metrics showed that subjects used a controller that exponentially stabilized trajectories. During continuous cup-andball displacements subjects developed predictable solutions sacrificing smoothness and energy efficiency. These results may stimulate control strategies for dexterous robotic manipulators and human-robot interaction.

ARTICLE HISTORY

Received 23 January 2020 Revised 8 May 2020 Accepted 27 May 2020

KEYWORDS

Complex object manipulation; human motor control; underactuated; chaos; stability

1. Introduction

Physical interaction with objects and tools is a hallmark of human behavior as highlighted by the central role of increasingly sophisticated tools in the evolution of homo sapiens. Despite rapid development in robotic manipulation, humans still easily outperform robots in dexterous and adaptive behavior, especially when interacting with objects and using tools. This is remarkable as the human sensorimotor system appears inferior to robots in many ways: the information transmission speed is extremely slow, the actuator bandwidth is significantly smaller, and there is pervasive noise at all levels of the human sensorimotor system. This disparity between the biological and robotic systems raises the question of how humans achieve their remarkable dexterity. Human manipulation skills become particularly fascinating when the objects have internal degrees of freedom that add complex dynamics to the interactions. For example, it is part of daily life to handle objects that contain fluids such as a cup filled with coffee [1] or carry a tray to serve food. It is fair to say that humans even seek such complex challenges as demonstrated in many sports and circus acts, such as gymnastics with a ribbon or cracking a whip, both objects with infinitely many degrees of freedom [2]. Better understanding of the human sensorimotor control system may lead to advances in robotic control, specifically in manipulation.

Several existing approaches to robot control have already drawn inspiration from biological motor control. Most prominently, robotic locomotion has been inspired by central pattern generators in animals, which are semiautonomous rhythm generators located in the spinal cord of vertebrates [3,4]. The control of posture and balance in bipedal robots has included feedback loops with different time delays simulating those of humans [5]. Considerably fewer bridges to biological system have been seen in the area of manipulation. One exception is the framework of Dynamic Movement Primitives (DMP) that has been inspired by nonlinear dynamics and human research on discrete and rhythmic behaviors [6-8]. This framework has been applied to a variety of behaviors with visible success [9,10]. A deeper understanding of human motor control has also provided insights to improve physical interaction and collaboration between humans and

CONTACT Salah Bazzi S.bazzi@northeastern.edu Department of Biology, Northeastern University, Boston, MA 02115, USA; Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, USA

robots. For example, a recent study showed that when following a robot trajectory that shared the same velocity-curvature relation as human movements, the human actor exerted less force on the robot [11].

Most studies on object manipulation in human motor control have been limited to examine grip forces when holding or transporting solid objects [12-16]. The relatively small number of studies on the manipulation of more complex objects has focused on two tasks: balancing a pole, i.e. an unstable object, and manipulating a linear mass-spring system. Several studies on human pole balancing proposed a range of modeling concepts: intermittent or continuous control, predictive control with forward or inverse internal models [17-19]. For the manipulation of a linear mass spring, the modeling of human performance has been confined to optimizationbased models of the controller, with objective functions such as 'crackle', i.e. minimizing the fifth derivative of a trajectory [20], optimizing the trade-off between accuracy and effort [21], minimizing the acceleration of the center of mass [22], and minimizing a dynamically constrained jerk [23].

While these studies have provided interesting results, most of these approaches assumed that the human has, or eventually learns, a precise internal model of the manipulated object. This model is then the reference for online and predictive control [24]. While plausible and rational, it is difficult to imagine that this approach can extend to dynamically more complex objects, i.e. objects with nonlinear, underactuated, and potentially chaotic dynamics that are inherently unpredictable. Furthermore, fully relying on feedback to correct for inaccurate predictions is not sufficient as the feedback delays in the human neuromotor system are astonishingly long: compared to milliseconds in robots, human transcortial feedback loops are in the order of at least 200 ms [25]. The more complex the information that is processed, the longer the loop times become. Given the instantaneous nature of interactions, such delayed corrective feedback is unlikely to be successful. We therefore argue that learning accurate and precise internal models to serve for predictive control may not be the primary strategy for humans. Rather, we hypothesize that humans learn control strategies that modify the interactions with the object to be more predictable. Predictability is a concept that is core in a large range of scientific disciplines. In the context of human physical interaction with complex objects, we defined predictability as the degree to which the dynamical behavior of the object can be predicted. This implies that the uncertainty about the object's future behavior is low, which makes it easier for humans to predict the object motion, at least in the short term. We argue that

predictability can be achieved by simplifying the dynamics of the interaction and by seeking stable regimes. Predictability may then afford simpler internal models.

To test this hypothesis, we examined the control strategies that humans employ when physically interacting with a dynamically complex object. The experimental task was motivated by the everyday action of carrying a cup of coffee: A cup filled with coffee exemplifies a dynamically complex object that is underactuated, nonlinear, and potentially even chaotic. We conducted four studies that tested the hypothesis that humans seek to increase predictability of the object dynamics when physically interacting with such a dynamically complex object. To examine human control strategies, we have adopted a 'task-dynamic' approach: rather than positing a hypothetical control strategy, we first examine the task with its dynamics and constraints and derive the solutions that the task affords [26]. This allows for minimal assumptions, if any, about the human neuromotor controller. By analyzing what is known, namely the physical task, and the possible solutions, we can determine the variables that are under control of the performer and determine the space of all task results. Exact quantitative hypotheses can then be formulated for select executions to identify which criteria were of primary concern or, in some cases, which costs were minimized by the performer.

This paper reviews and integrates the results of four different studies on human manipulation of complex objects with the goal to present a human-based perspective on object control to the robotics community. The specific objective is to present how the umbrella concept of predictability may inform about human control strategies, in the hope that this may stimulate new control strategies for dexterous robotic manipulation. The paper is organized as follows. Section 2 presents the experimental task and the model of the complex object with which the subjects interacted, namely the cup-and-ball system. Section 3 illustrates the challenges presented by complex object manipulation and we show that current models of human control are insufficient emphasizing the need for new control principles in human manipulation. The complex properties and demands for the control of the cup-and-ball system are illustrated in Section 4. Section 5 reviews two experiments involving discrete movements, where predictability is quantified as stability. Rhythmic movements are examined in Section 6 where predictability is quantified by the information-theoretic concept of mutual information. In Section 7 the results and their implications are discussed, and we conclude the paper in Section 8 with a discussion of how these insights may inform robotics.



Figure 1. (A) Real and simplified task. (B) Mechanical model. (C) Virtual environment with a subject operating the HapticMaster robot to manipulate the virtual cup. (D) Screen display.

2. The task: moving a cup of coffee

Transporting a cup filled with sloshing coffee is an example of a physical interaction with a dynamically complex object that has served as the testbed for our theoretical and experimental studies. However, modeling a realistic 3-dimensional cup with moving fluid requires solving nonlinear partial differential equations that immediately shows the computational challenge [1,27]. To turn this task into an experimental paradigm that can provide insights into control, the complexity of the object was significantly reduced, yet the essential elements of this dynamical system were retained: underactuation and nonlinearity, where the latter could induce chaotic behavior. The cup was simplified to a 2-dimensional semicircular arc with a ball rolling inside [28,29], representing the sloshing coffee (Figure 1(A)). The ball's motion was modeled by a pendulum suspended from a cart; the arc of the cup corresponded to the ball's semicircular path (Figure 1(B)). The motion of this 2D cup was limited to the horizontal axis only. To eliminate the complexity of the high-dimensional grasping forces the hand interfaced with this dynamic system only via a single port (Figure 1(B)). To measure human control of this dynamic object, this model system was implemented in a virtual environment with a visual and haptic interface using a robotic manipulandum (Figure 1(C)). For all experiments, a projection screen displayed the cup-and-ball system moving between a start box (Box A) and a target box on a horizontal line (Box B) (Figure 1(D)). The size and distance between the two boxes varied between the experiments, but were between 8 and 45 cm.

Mechanically speaking, the simplified system is equivalent to the widely studied undamped cart-and-pendulum system. The cart represents the cup that moves horizontally, although the cart is not shown; the ball is the pendulum bob attached to the cart via a massless rod. The arc of the pendulum is shown visually to represent the cup. Subjects controlled the ball indirectly by applying forces to the cart/cup. The ball could be lost when its angle exceeded the rim angle, equivalent to 'spilling' the coffee. The equations of motion of this cup-and-ball system were

$$(m+M)\ddot{x} = ml\left(\dot{\phi}^2\sin(\phi) - \ddot{\phi}\cos(\phi)\right) + F,\qquad(1)$$

$$l\ddot{\phi} = -g\sin(\phi) - \ddot{x}\cos(\phi), \qquad (2)$$

where x denoted the position of the cart, ϕ denoted the pendulum angle with a counter-clockwise positive convention, m was the mass of the pendulum bob, M was the mass of the cart, l was the length of the massless pendulum rod, and g denoted gravitational acceleration. The force exerted by the human subject on the handle of the robotic manipulandum was F.

Subjects manipulated the virtual cup-and-ball system via an admittance-controlled robotic manipulandum, which also exerted forces from the virtual object back onto the hand (HapticMaster, Motekforce, NL; [30]). The HapticMaster had 3 controllable degrees of freedom, but was constrained to move on a horizontal line for the experiments. The pendulum's ϕ and $\dot{\phi}$ were computed using a 4th-order Runge-Kutta integrator. The force that the ball imparted on the cup, F_{ball}, was computed based on (1) $F_{ball} = ml(\dot{\phi}^2 \sin(\phi) - \ddot{\phi} \cos(\phi))$ and presented as haptic feedback to the subject. This force, combined with any forces exerted by the subject F, accelerated the virtual mass (m + M). The robot motors moved the manipulandum according to \ddot{x} and the visual display was updated with a negligibly small delay. The force applied by the participants to the manipulandum and the kinematics of the ball and the cup were all recorded at 120 Hz.

3. Minimum jerk trajectory of the hand is not sufficient

In human motor control, a host of studies have shown or assumed that smoothness is a primary criterion that humans optimize [31–33]. Specifically, minimization of jerk has been the most common description for hand trajectories. Note though that the majority of studies have focused on simple unconstrained reaching movements, typically confined to the horizontal plane [34,35]. A few studies have generalized this notion and proposed that minimum jerk of the hand is also a reasonable description of how humans manipulate linear mass-spring systems [23,36]. However, this model has not yet been tested when the interaction included more complex objects. To demonstrate the scope and limitations of this control model for movements involving interactions, we conducted a simple simulation.

Subjects (n = 4) interacted with the virtual cup-andball system in two different ways: To simulate unconstrained reaching, the ball was fixed to the bottom of the cup. In this condition, the subject essentially interacted with a rigid object with a mass equal to the summed mass of cup and ball. In the interactive condition, the ball was free to move inside the cup subject to the force applied to the cup. In both cases, subjects were instructed to move the cup-and-ball system 45 cm from Box A to Box B at a comfortable pace (Figure 1(D)). The experimental session was comprised of 4 blocks of 30 trials (120 total). The conditions in the blocks alternated between the rigid system and the dynamic cup-and-ball system.

Figure 2 displays the velocity profiles of the cup (equivalent to the hand profiles) and the ball for one representative subject. The figure shows the last 10 trials of each condition; one representative human profile is highlighted by the bold blue line. This profile is compared with the simulated profile using the minimum jerk criterion. Figure 2(A) shows the rigid object condition where the subject interacted with the object when the pendulum was fixed: as demonstrated before, the hand trajectories show a repeatable bell-shaped profile with some symmetry in the ascending and descending branch. This profile was well approximated by the minimum jerk profile shown in Figure 2(A), bottom panel. However, this coincidence disappeared when the subject interacted with the dynamic cup-and-ball system (Figure 2(B,C)). The velocity profiles of the cup were no longer bell-shaped, but contained additional 'shoulders' and 'humps', caused by the moving ball acting back upon the cup. The ball trajectories were relatively variable, especially in the later segment, but converged to a profile with one undulation before coming to rest. Figure 2(B) (bottom panel) directly compares one representative human cup velocity profile with the simulated cup profile generated by an input F that would generate a minimum jerk cup trajectory in the rigid object case. When using this smooth cup profile to move the non-rigid system, Figure 2(C) (bottom panel) shows the resulting ball profile: the time course of the



Figure 2. (A) Experimental and simulated cup trajectories when the ball was fixed to the bottom of the cup. (B) Experimental and simulated cup trajectories when the ball was free to move. (C) Experimental and simulated ball trajectories when the ball was free to move.

ball trajectory differs significantly, with large ball angles that may lead to a loss of the ball. This illustration clearly suggests that the subject compensated for the ball force, with the final result that the ball has only low-amplitude oscillations at the target.

This simple demonstration highlights that while the minimum jerk principle is successful for generating unconstrained point-to-point movements, it is not appropriate for interacting with a complex object. More generally, these results demonstrated that insights gained from studying unconstrained reaching movements have limited value to complex interactions with objects. This highlights the need for other control principles that are ultimately able to generate the rich behavioral repertoire that humans exhibit - control principles that robots can mimic.

4. Chaos and unpredictability in the cup-and-ball system

Consider the task of oscillating a rigid object or a linear mass-spring system with the goal of attaining a specific periodic behavior. The resulting amplitude and frequency of oscillation are linearly related to the applied external force. In contrast, a nonlinear system such as the cup-and-ball system does not exhibit such a linear mapping: the same external force input can cause the system to oscillate at a range of different frequencies. For extended durations, the dynamics can become unpredictable, and even chaotic.

To illustrate that the relatively simple, yet nonlinear cup-and-ball system can exhibit chaotic behavior, at least when the interactions are extended in time, inverse



Figure 3. (A) Required force profile for maintaining a sinusoidal cup motion when $\phi_0 \approx -28.4$ deg. The resulting cup and ball trajectories are also displayed. (B) Required force profile for maintaining the same sinusoidal cup motion in (A) when $\phi_0 \approx 17$ deg. The resulting cup and ball trajectories are also displayed. (C) Diagram of strobed forces to summarize the complexity of the required input forces for different initial conditions: For every initial ball angle, at every maxima of cup displacement, the value of *F* was determined, as displayed by the red points at each peak in (A) and (B). The resulting strobed force value was plotted vertically above the initial ball angle to obtain the marginal distribution of the strobed force profile. These distributions were obtained for each simulation with different ϕ_0 .

dynamics simulations were used to obtain the force F required to sustain a specific oscillatory motion of the cup. Starting from different initial ball angles ϕ_0 with the initial angular velocity set to 0, forces F were generated that produced a cup trajectory with a given amplitude A and frequency f. Figure 3 displays two simulations that started from different initial ball angles, but produced the same oscillatory movement of the cup, albeit with very different ball angles. Inverse dynamics calculations reveal that the force profiles required to produce this cup movement are qualitatively different in the two cases. In Figure 3(A), the force is periodic and predictable, whereas in Figure 3(B) it is highly irregular and potentially chaotic. To characterize the patterns of force profiles in relation to the cup dynamics, the system was simulated for different initial ball angles and the force profiles were generated. To summarize the (ir-)regularity of the force profiles for the different initial ball angles, the force profiles were strobed, i.e. sampled at the time of each peak in the cup displacement. The resulting force values were projected to obtain the marginal distributions for each simulation. These marginal distributions of strobed force values were plotted as a function of initial ball angle ϕ_0 in Figure 3(C). This figure shows essentially input-output relations, i.e. force input and ball output. The figure reveals a pattern similar to the period-doubling behavior of chaotic systems [37], implying that for some initial conditions, the system is chaotic. (Note though, that this inverse dynamics does not consider any feedback processes that may occur in real manipulation, but only characterizes the behavior of the cup-and-ball system.) Nevertheless, this figure has important implications for the control of this dynamic system: small changes in initial states can dramatically change the behavior of the system and render it essentially unpredictable in the long term. Such small perturbations readily arise from the fact that human movements are always variable due to their intrinsic noise.

5. Predictability and stability

We hypothesized that humans seek solutions with predictable object behavior. While plausible and intuitive at first sight, predictability is an umbrella concept with multiple features that can be mathematically defined in several ways. A first proposition to operationalize predictability is in terms of stability: A dynamic system that is stable rejects small perturbations and returns to its attractor. This attractor is an invariant set, present throughout the evolving dynamics, and is therefore predictable. This obviates the need for explicit error corrections and extensive computations based on an accurate and precise model of the system's nonlinear dynamics. When the system is at a stable attractor, model-based closed-loop control becomes less critical. Stability also provides robustness to the ubiquitous noise present at all levels of the sensorimotor system. We therefore proceeded to evaluate stability of the human trajectories when moving the cup-and-ball system to a target.

5.1. Quantifying predictability as stability or contraction

To assess stability, existing methods require and assume the behavior to be at steady-state, i.e. close to or at a fixed-point attractor or a limit cycle [38]. However, for the task of manipulating the cup-and-ball system, these assumptions do not hold: performing a point-to-point movement does not reside on either of these two attractor types. Even if there is an attractor, it is not known. We propose that contraction analysis, a differential form of stability analysis for nonlinear dynamical systems, is an appropriate tool for this scenario [39]. Contraction analysis quantifies the convergence or divergence of neighboring trajectories in state space. A favorable property of contraction analysis is that it does not require knowledge of the stable solution/attractor. This is useful for studying dynamically complex physical interactions as in the model task studied here. When transporting the cupand-ball system from a starting to a target position the system is not at steady state, but rather in a transient state.

Consider a nonlinear dynamical system

$$\dot{x} = f(x, t), \tag{3}$$

with two neighboring trajectories x(t) and $x^*(t)$, which are two solutions of (3) with initial conditions x_0 and x_0^* , respectively. At any given fixed time, the infinitesimal distance between these trajectories is denoted as δx and referred to as the virtual displacement (Figure 4(A)). The temporal evolution of this virtual displacement is governed by the exact differential relation

$$\delta \dot{x} = \frac{\partial f}{\partial x}(x, t)\delta x,\tag{4}$$

from which we can derive the equation for the squared distance

$$\frac{d}{dt}(\delta x^T \delta x) = 2\delta x^T \frac{\partial f}{\partial x} \delta x.$$
(5)

We can see that if the Jacobian $\partial f/\partial x$ is uniformly negative definite, then the virtual displacement $\|\delta x\|$ converges exponentially to zero. A contraction region is any region in the state space in which this negative definiteness property holds. Mathematically, the Jacobian is said to be uniformly negative definite if all eigenvalues of its symmetric part are uniformly negative definite.



Figure 4. (A) Two solutions of the same dynamical system starting from different initial conditions, x_0 and x_0^* , displaying contracting behavior. The virtual displacement δ_x between the trajectories is shown. (B) Two solutions of the same dynamical system starting from different initial conditions, z_0 and z_0^* , displaying diverging behavior.

Contraction is only a sufficient condition for exponential convergence. A both necessary and sufficient condition for exponential convergence can be formulated by considering a continuously differentiable coordinate transformation of the system

$$\delta z = \Theta(x, t) \delta x, \tag{6}$$

where $\Theta(x, t)$ is a square matrix that satisfies $\Theta^T \Theta > 0$. Θ is referred to as the contraction metric and the generalized Jacobian is now

$$K = \left(\dot{\Theta} + \Theta \frac{\partial f}{\partial x}\right) \Theta^{-1}.$$
 (7)

In this new frame, a contraction region is one where the generalized Jacobian is uniformly negative definite.

To summarize, a region of the state space is said to be contracting if there exists a metric $\Theta(x, t)$ with $\Theta^T \Theta > 0$ such that $1/2(K + K^T) < 0$. Any trajectory entering a ball of constant radius around another trajectory within a contraction region remains in that ball and converges exponentially to that trajectory.

This theoretical framework was applied to the pointto-point movements of the cup-and-ball system collected in human experiments.

5.2. Study 1: navigating perturbations in a point-to-point transport task

Human subjects (n = 7) were instructed to transport the cup-and-ball system from Box A to Box B (Figure 1(D)), without losing the ball [38,40,41]. To create an additional challenge and make stability a useful feature for success, a perturbation of magnitude 40 N, duration 20 ms, was presented at 60% of the travel distance. The position of the perturbation was displayed as a small bump, but the virtual cup moved through the bump and remained on the horizontal line. This perturbation acted either in the

direction of motion of the cup, which is referred to as assistive, or it acted against it, referred to as resistive. Importantly, the perturbations were visible and present in every trial so that subjects could learn how to navigate the perturbation. Subjects were instructed to perform the discrete movement as fast as possible (the distance between the boxes was 45 cm). The experiment consisted of 4 blocks: Block 1 comprised 60 trials without perturbation to allow subjects to familiarize themselves with the task. Blocks 2 and 4 comprised 60 trials each and involved a series of either assistive or resistive perturbations, respectively. Block 3 presented 10 unperturbed trials to separate the two perturbation conditions. The hypothesis was that subjects would learn to exploit contraction regions to accommodate for these perturbations. If so, these strategies should be distinct for the two types of perturbations.

To make the task sufficiently challenging, the simulated dynamics incorporated a positive gain for the cup acceleration; this rendered the ball movements more sensitive to the force applied to the cup. Further, to ensure the existence of contraction regions, the system had to include energy dissipation. This was achieved by adding damping to model the coupling between the hand and the object, i.e. including hand impedance. The modified dynamical model of the system was

$$(m+M)\ddot{x} = ml\left(\dot{\phi}^2\sin(\phi) - \ddot{\phi}\cos(\phi)\right) + F - B\dot{x}, \quad (8)$$

$$l\ddot{\phi} = -g\sin(\phi) - G\ddot{x}\cos(\phi), \qquad (9)$$

where B is the damping coefficient and G is the cup acceleration gain factor.

In the system's canonical form presented in (8)–(9), the Jacobian was not uniformly negative definite in any region of the state space. However, this did not rule out the existence of contraction regions since negativity of the Jacobian is only a sufficient condition. Therefore, the next step was to find a contraction metric that would reveal the contraction regions of the system. This was achieved by solving the following differential equation, adopted from [39]

$$\frac{\partial\Theta}{\partial X}f + \Theta J = -\Theta, \tag{10}$$

where f is the vector field describing the dynamics of the system and J is the Jacobian of the system in its canonical form (8)–(9). This partial differential equation was solved numerically to obtain the contraction metric, which in turn enabled the computation of the generalized Jacobian K. The contraction regions of the system were then computed by checking the points in the state space where K was uniformly negative definite. Note that the force F



Figure 5. Human trajectories when negotiating the assistive perturbation. (A) Early trial: the subject did not exploit any contraction region. (B) Late trial: the subject encountered the perturbation entirely inside a contraction region. P^- denotes the instant just before the perturbation while P^+ denotes the instant after the perturbation. The black arrowhead indicates the starting point of the trajectory. The exact ellipse indicates that the cup has reached the target box at rest, while the pendulum was still oscillating.



Figure 6. Human trajectories for the condition with resistive perturbations. (A) Early trial: in this trial, the subject dropped the ball from the cup, hence the trajectory ends just after P^+ . (B) Late trial: the subsequent trajectory after the perturbation entered a contraction region. P^- denotes the instant just before the perturbation while P^+ denotes the instant after the perturbation. The black arrowhead indicates the starting point of the trajectory. The exact ellipse indicates pendulum oscillations for zero cup velocity at the target box.

does not contribute to the Jacobian, and hence this analysis only revealed contraction regions of the open-loop, unforced system.

With this theoretical framework, the human data could be evaluated. To test the hypothesis that subjects exploited the contraction regions, the experimental trajectories of the cup-and-ball were plotted in 3-dimensional state space. The calculated contraction regions were then overlaid and displayed as volumes. Figure 5(A) illustrates an early trial of one representative subject when facing an assistive perturbation, while Figure 5(B) shows a late trial of the same subject navigating the same perturbation. At the beginning of practice, the trajectory did not pass through any of the contraction regions. However, with more practice, the subject learned to enter the contraction region just before the onset of the perturbation, causing the perturbation

to occur within the contraction region. This mitigated the destabilizing effect of the perturbation. For resistive perturbations, a different strategy emerged. Similar to the assistive perturbations, early trials did not approach the contraction regions (Figure 6(A)). However, with practice subjects shaped their trajectories to pass through a contraction region directly after the perturbation. This attenuated the destabilizing effects of the perturbation (Figure 6(B)). This pattern was consistent across all subjects.

These observations supported our hypothesis that subjects exploited contraction regions to preempt destabilizing and unpredictable perturbations that would require computationally expensive corrections. In more general terms, we propose that these findings provide first support for the hypothesis that humans seek predictability via stability in complex object manipulation.

5.3. Study 2: control contraction metrics for assessing stability

So far, the contraction analysis only identified the contraction regions of the free unforced system; the human control input was not included in the analysis. As a next step we extended the question to 'Does the human control policy actively seek to stabilize the cup trajectories and make them contracting?' To answer this question, our approach was to use existing control design methods to develop a stabilizing controller and compare its output to the human force applied. Several control design tools exist for the stabilization of nonlinear systems: Lyapunov-based [42], passivity-based [43], and contraction-based [44]. For the same reasons we mentioned before, a contraction-based method was chosen: the concept of control contraction metrics (CCMs) was leveraged [44]. A CCM is a contraction metric that guarantees the existence of an exponentially stabilizing (contracting) controller for any dynamically feasible trajectory of a nonlinear system. First, the system has to be expressed in its control-affine form:

$$\dot{x} = f(x) + B(x)F. \tag{11}$$

Then, the CCM is mathematically defined as one that satisfies the following conditions [44]:

$$\partial_{b_i} M(x) + M(x) \frac{\partial b_i}{\partial x} = 0, \quad i = 1, \dots, m$$
 (12)

$$\delta_x^T \left(\partial_f M(x) + \widehat{M(x)} \frac{\partial f}{\partial x} \right) \delta_x < -2\lambda \delta_x^T M(x) \delta_x, \quad (13)$$

where *M* is the CCM, b_i is the *i*th column of *B*, δx is an infinitesimal displacement, and $\lambda > 0$ is the contraction rate. The first condition (12) ensures that, along the actuated directions in the vector field, distances are preserved and trajectories do not diverge. Condition (13) ensures that, in the uncontrolled directions orthogonal to the controlled ones, the system is naturally contracting. In principle, CCMs are a generalization of control Lyapunov functions (CLFs), where the energy of the geodesics between neighboring trajectories is the CLF [45].

The search for a CCM can be written as a convex optimization problem by considering the dual of (12)-(13):

$$\partial_{b_i} W(x) - \frac{\partial \widehat{b_i}}{\partial x} W(x) = 0, \quad i = 1, \dots, m$$
 (14)

$$B_{\perp}^{T}\left(-\partial_{f}W(x) + A(\widehat{x,F})W(x) + 2\lambda W(x)\right)B_{\perp} \prec 0,$$
(15)

where *W* is the dual metric $W(x) := M(x)^{-1}$, $A(x, F) := (\partial f / \partial x) + \sum_{i}^{m} (\partial b_i / \partial x) F_i$, and B_{\perp} satisfies $B_{\perp}^T B = 0$. A

computationally tractable finite-dimensional approximation of this CCM feasibility problem can be obtained by casting it as a Sums-of-Squares (SOS) program; the metric W is parameterized as a matrix of polynomials, and the Linear-Matrix-Inequality (LMI) (15) is relaxed and expressed as a SOS constraint [46].

Once a CCM is found, its corresponding exponentially stabilizing controller can be defined. Importantly and consistent with our task-dynamic approach, the analysis made no assumptions about the structure of the human controller. Rather, the experimental measurements of force from the subjects performing the task were evaluated at each time instant and compared to the values that an exponentially stabilizing controller would generate. On that basis, each data point of a given human trajectory was classified as either contracting or not. In each experimental trial, the degree of contraction was quantified by the number of data samples that showed contraction divided by the total number of samples of the trial, which we referred to as the contraction ratio.

Figure 7(A) shows the mean contraction ratios for every trial across all subjects. The increase across trials indicated that with practice subjects learned to stabilize more or longer segments of their trajectories. Correlating the contraction ratio to trial duration, a strong and statistically significant correlation was present (Figure 7(B,C)). Note that trial duration was one metric for the goodness of task performance as subjects were instructed to navigate the perturbation as fast as possible. This suggested that subjects achieved better and faster performance by a strategy that contracted their trajectories.

For a subsequent analysis, successful trials for each perturbation type were split into two categories: fast and slow, based on the average trial duration. When averaging the trials in each category and summing the contracting regions, fast trials showed more contracting segments than slow trials. Figure 8 shows the average cup velocity over the horizontal position; the color gradient expresses how many trials at each instant were contracting. In both assistive and resistive conditions, the fast trials had more contracting segments directly prior to and immediately after the perturbation onset than the slow trials. This indicated that subjects utilized contraction to attenuate the effect of the perturbation.

Integrating the results of these two studies, we deduce the following. While both studies showed that humans sought contraction at the onset of the perturbation, whether by guiding the system to a contraction region or by rendering the trajectories contracting through control effort, the second study revealed insight beyond this. An important element of the subjects' strategy was preempting the perturbation by contracting the trajectory before impact, and then following through by contracting



Figure 7. (A) Mean contraction ratios across all subjects \pm one standard error. (B) A negative correlation between trial duration and contraction ratio for the assistive perturbations. (C) A negative correlation between trial duration and contraction ratio for the resistive perturbations.



Figure 8. Average fast and slow trajectories and their contracting segments across all subjects in the respective perturbation blocks. The color code represents the percentage of trials that were contracting at a given position.

the trajectory after the perturbation as well. Together, these two different analyses are complementary since they reveal to what degree the subjects relied on the objects passive stability properties versus actively stabilizing the system through control effort.

These two studies are the first that applied contraction analysis to human data showing that contraction analysis is a suitable tool to provide insights into human movements. Importantly, the results were consistent with our hypothesis that humans sought to increase contraction to preempt the effect of perturbations that cause unpredictable after-effects. These findings support that humans seek stability as one way to make the complex interactions predictable and potentially avoid error corrections. We propose that this human strategy may also inspire roboticists to develop control algorithms that generate contracting trajectories in robots. In fact a recent study has initiated efforts in that direction by applying control contraction metrics for generating robust motion plans for quadrotors [47].

6. Predictability and mutual information

The two previous studies pursued stability or contraction as a way to quantify predictable interactions in a point-to-point movement. In these short movements that started with defined initial conditions, the dynamics remained relatively simple. To allow for more complex dynamics in this simple model task, two complementary studies investigated extended interactions with the cup-and-ball system that revealed more of the challenges arising from the nonlinearity of the complex object. To test the hypothesis that humans seek predictable interactions, these studies used an alternative measure, mutual information, to quantify predictability.

Mutual information is a nonlinear correlation measure defined between two probability density distributions of two random variables; it quantifies the information shared between the two signals. For the present purpose, mutual information (MI) was calculated between the applied force F and the motion of the object, captured by the phase of the ball ϕ_{ph} :

$$MI(\phi_{ph}, F) = \iint P(\phi_{ph}, F) ln \left[\frac{P(\phi_{ph}, F)}{P(\phi_{ph})P(F)} \right] d\phi_{ph} dF,$$
(16)

where p(.) denotes a probability density function. In this context, MI quantifies the degree to which the long-term evolution of the applied force can be predicted if the object's trajectory was known [48,49]. MI is a scalar measure that summarizes the predictability of the performer's strategy for a given set of four variables: cup amplitude A, cup frequency f, initial ball angle ϕ_0 , and initial ball velocity $\dot{\phi}_0$. Natural logarithm was used in (16) and hence we use nat as the unit of MI.

6.1. Study 3: rhythmic interactions with prescribed frequency of oscillation

Human subjects (n = 8) were instructed to oscillate the cup-and-ball system for 45 s between two large target boxes (see Figure 2). They were paced by a metronome to complete one back-and-forth movement in 1 s, i.e. at a frequency of 1 Hz. As the target boxes were relatively wide, subjects were free to choose their preferred movement amplitude between 8 and 44 cm; there were no specific instructions about the ball movements [49]. As in the previous studies, the robot manipulandum gave online haptic feedback about the ball forces acting on the hand. Each subject completed 50 trials of 45 s each in one experimental session.

Without any assumptions about the controller, the model task was first analyzed to distinguish between execution variables, i.e. variables that are at the control of the subject, and result variables, i.e. those that resulted from the dynamics of the object given the subject's input. In this case, the amplitude A, movement frequency f, initial ball phase ϕ_0 , and initial phase velocity $\dot{\phi}_0$ were the execution variables as they could be controlled by the subjects. Note that cup oscillation frequency f was fixed to 1 Hz and initial ball angular velocity $\dot{\phi}_0$ was set to 0. Task execution could be simulated for the range of all execution variables to quantify the predictability of the resulting dynamics in these different strategies. Forward simulations of Equations (1) and (2) were performed for different initial ball angles ϕ_0 and cup amplitudes A to generate cup displacement x at a frequency 1 Hz. Initial ball velocity was set to zero. For each value of A and ϕ_0 , the mutual information between the applied force and the resulting ball dynamics was quantified for the different execution strategies. Figure 9(A) shows the result space for mutual information as a function of A and ϕ_0 ; lighter shadings indicate higher mutual information and, hence, higher predictability. The large point marks the strategy

with the highest mutual information. This result space presents the reference and subject's data can be plotted into the same space as the execution variables can be experimentally measured. With this task-dynamic analysis, the hypothesis can be tested that subjects seek the point of maximum mutual information, i.e. predictability of the object's dynamics.

To compare predictability with alternative objective functions, exerted force and movement smoothness of the resulting trajectories, widely accepted in the human control literature, were also computed from the simulated trajectories. To this end, the expended force was calculated by squaring and then averaging the force profile F(t)over the course of a trial and denoted as mean squared force MSF:

$$MSF = \frac{1}{kT} \int_{0}^{kT} F(t)^{2} \,\mathrm{d}t,$$
 (17)

where *k* denoted the number of cycles and T = 1/f the period of each cycle. The result space for mean squared exerted force is shown in Figure 9(B); lighter shades refer to strategies requiring less force with the large point at the smallest amplitude highlighting the minimum force solution.

The second alternative hypothesis was that subjects would maximize smoothness. As discussed above, this objective function is widely regarded adequate for unconstrained free movements. To quantify smoothness, the normalized mean absolute jerk of the ball trajectory was calculated:

$$Jerk = \frac{1}{T(\ddot{\phi}_{\max} - \ddot{\phi}_{\min})} \int_0^T |\ddot{\phi}| \, dt.$$
(18)

Smoothness was evaluated in the ball trajectory since the cup trajectory was assumed to be sinusoidal following task instructions and hence smooth by definition. Figure 9(C) shows the result space for the smoothness of the ball trajectory for a given strategy. Lighter shades denote higher smoothness and the large point represents the strategy with maximum smoothness of the ball trajectory. Importantly, the three maxima lie in different locations of the execution space, clearly differentiating the alternative hypotheses. Therefore, by extracting the amplitudes and relative phases from each trial, the three hypotheses could be evaluated.

To permit comparison between model simulations and experimental data, equivalent measures had to be derived from the experimental data, specifically for the four execution variables A, f, ϕ_0 , and $\dot{\phi}_0$. To pay tribute to the fact that human trajectories were not fully determined by the initial states of the 45-s-long trial as corrective online changes were likely, the estimates of



Figure 9. Result spaces spanned by the execution variables cup amplitude *A*, cup frequency *f* (fixed at 1 Hz), initial ball angle ϕ_0 , and initial ball velocity $\dot{\phi}_0$ (set to 0). (A) Mutual information for each (A, ϕ_0) ; lighter colors denote higher *MI*. (B) Mean squared force (log transformed) for each (A, ϕ_0) ; lighter colors denote lower mean squared force. (C) Normalized mean absolute jerk of the ball trajectory for each (A, ϕ_0) ; lighter colors denote lower jerk. The large points in each graph denote the optimal locations for the respective measure.



Figure 10. Main results in the result space for mutual information. The green contours represent different values of mean squared force superimposed onto the same space. (A) Data points for all subjects where each point represents one trial. Darker red denotes earlier trials. Data shows that the later trials (lighter red) were concentrated in the areas of high mutual information. (B) Average data points for each subject. Arrows represent one subject, with the tail coinciding with the mean of the first 5 trials and head is located at the mean of the last 5 trials. Most of the subjects moved towards more predictable regions in the result space.

the corresponding execution variables were extracted for each cycle k, treating each cycle as a separate initial conditions. To obtain equivalent measures from the subjects' data, the estimated execution variables from each cycle kwere then averaged across cycles to obtain \overline{A} , \overline{f} , $\overline{\phi_0}$, and $\overline{\phi_0}$.

The experimental results are summarized in Figure 10. The figure shows how subjects' strategies started in areas with low mutual information, but, with practice, they steered towards areas with higher MI. The left panel shows the strategies for each trial pooled for all subjects; each point represents a trial with darker red indicating earlier trials and lighter red indicating later trials. The same data is displayed in the right panel but separated by subject. The red arrows mark how each subject's strategy changed with practice from the mean of the first 5 trials to the mean of the last 5 trials. The arrows show that most subjects change from higher mutual information towards the area of lower mutual information in a gradientdescent-like manner, regardless of the different starting points.

The green contours represent isocontours of mean values of mean square force MSF, with the minimum MSF for the lowest amplitude. Clearly, the data do not steer towards the point of minimum force. In fact by increasing mutual information, the expended force increased, not decreased, as might have been expected. Further, not directly shown here, is that the data are clearly distinct from the point of minimum smoothness (Figure 9(C)). This result resonates with the initial demonstrations that a minimum jerk hand trajectory does not achieve successful behavior. Note though, that these simulations evaluated smoothness of the ball trajectory, not the hand trajectory, as in the earlier simulations at the beginning. Overall, the results rejected the alternative explanations of minimizing effort and maximizing smoothness, and supported the hypothesis that subjects favored predictable solutions.

6.2. Study 4: rhythmic interactions with prescribed amplitude of oscillation

To examine these results further, a complementary study investigated subjects' behavior (n = 10) when instructed to continuously oscillate the cup-and-ball system at frequency of their own choice [50]. The cup amplitude was specified by two smaller target boxes on the screen that specified an amplitude of 45 cm for the cup movements (see Figure 1). As before, the experimental session was comprised of 50 trials, each trial lasting 45 s. Again, the main hypothesis was that subjects seek to increase the predictability of the manipulated object dynamics. Specifically, this experiment tested whether subjects selected movement frequencies that enhanced the predictability of the cup-and-ball system. Specifically, the hypothesis was that subjects would exploit the resonant frequencies of the system. The alternative hypothesis was as above, i.e. subjects minimize exerted force.

A first inspection of the subjects' behavior focused on the frequencies of the cup oscillations that subjects preferred. The histogram in Figure 11 that plots the frequencies across all subjects shows two pronounced modes at two distinct frequencies separated by a 'dip' (Note that each back-and-forth movement is one data point to obtain a sufficiently large number of samples). For the low-frequency strategies, the cart and pendulum movements were in-phase (the maximum angle of the pendulum was synchronized with the maximum position of the cart in the same direction), whereas for the high-frequency strategies they were in anti-phase relation (ball angle excursion and cup maximum were in



Figure 11. Distribution of frequencies adopted by all subjects. The histogram shows the frequencies of every single cycle of all trials. Note that the *x*-axis is in log scale.

opposite direction). Subjects adopted either one or the other strategy throughout the experiment and did not switch between them. These two distinct modes in subjects' choices suggested that there were two resonances separated by an anti-resonance. For this resonance structure, the system needed to be 4th-order. Therefore, the model of the cup-and-ball system alone was not enough to fully capture the data. In order to guide any further analysis of experimental data, the model needed to be extended.

The simplest extension that could reproduce these two modes of behavior includes the hand and a controller coupled to the cup-and-ball system (Figure 12). The hand dynamics were assumed to be an ideal force generator (F_{input}) in parallel with a spring (stiffness, K) and a damper (damping coefficient, B). F_{input} was the force required to follow a desired trajectory ($x_{des}(t)$, $\dot{x}_{des}(t)$) whereas the spring and damper were a simplified model of hand impedance. This appended the equations of motion to:

$$(m+M)\ddot{x} = ml\left(\dot{\phi}^2\sin(\phi) - \ddot{\phi}\cos(\phi)\right) + F, \qquad (19)$$

$$l\ddot{\phi} = -g\sin(\phi) - \ddot{x}\cos(\phi), \qquad (20)$$

$$F = F_{input} - K(x - x_{des}) - B(\dot{x} - \dot{x}_{des}).$$
(21)

Given the task instructions, the desired trajectory was $x_{des}(t) = A \sin(2\pi ft + \pi/2)$. F_{input} was assumed to be the force required to manipulate a rigid object of equal mass $F_{input}(t) = (m + M)\ddot{x}_{des}(t)$, since humans can manipulate rigid objects very accurately suggesting that they have a good model of such a simple object. If the full dynamics of the task were perfectly anticipated, the subjects would be able to generate the exact



Figure 12. The extended model used in Study 4 to include coupling between the hand impedance and the cup.

 F_{input} required for the cart to precisely follow x_{des} . In reality, however, it was unlikely that subjects learned the exact model including the perturbing forces of the ball; hence, the hand impedance served to resist and restore any deviations from the desired cart trajectory.

To test the hypothesis that subjects selected movement frequencies that resided at the resonance frequencies of the system, we proceeded to conduct a frequency response analysis of the coupled model (with linearization). Evidently, the frequency response for the linearized system depended on the values of stiffness K and damping B. Hence, stiffness and damping values had to be estimated from the data. Even though K and B could not be measured directly from the experimental data, they were estimated for every trial using an optimization procedure. Using forward simulations with all feasible combinations of K and B, a root mean square optimization identified the K and B values that resulted in the best fit of the simulation with the human trial. This procedure was applied to all individual trials (for more details see Appendix B in [50]). Two clusters of K and B values emerged from these fits that were associated with the two separate frequencies.

Figure 13 shows the frequency response for these two sets of values of K and B. The left column shows the frequency response with K and B values representative of the low-frequency group, whereas the right column uses *K* and *B* values representative of the high-frequency group. The response of the cup displacement for both sets of values shows a sharp valley, indicating the antiresonance at 0.74 Hz between the two resonances (equivalent to the natural frequency of the pendulum). This anti-resonance frequency was systematically avoided by subjects. The phase plots in Figure 13 display the relative phase between the input force and the cart movement (red line) and the relative phase between the input force and the pendulum movement (blue line). These two curves highlight that for low frequencies the cart and pendulum are in phase, whereas for frequencies higher than the anti-resonance frequency, cart and pendulum movements are anti-phase.

This mathematical analysis then served to evaluate the experimental data. Comparing the resonant peaks of the models with the experimental data, for the low-frequency group, the peak in the distribution was very close to the resonance peak of the system. For the high- frequency group, participants showed a very broad distribution that matched with the smeared-out resonance peak of the system. These analyses confirmed the appropriateness of the control model and suggested that subjects chose or tuned their *K* and *B* values to move at one of the two resonance frequencies of the system.

Further analyses of the experimental data used this extended model. Based on this model, a strategy could be now defined by 4 different execution variables: amplitude *A* (which was prescribed), frequency *f*, stiffness *K*, and damping *B*. Note that initial ball angle and velocity ϕ_0 and ϕ_0 are essentially included in the choice of *f* since the movement frequency dictates the behavior of the ball (in-phase or anti-phase). Note that for the experimental analysis the execution variables have to be known. However, *K* and *B* could not be measured directly from the experimental data. Rather, they were estimated for every trial using an optimization procedure. The algorithm chose the *K* and *B* values for which the forward simulated trajectories best fitted the experimental trajectories (for more details see Appendix B in [50]).

In one next step, this hand-cup-and-ball model could serve as basis to evaluate whether these two frequency strategies corresponded to those with higher predictability, i.e. higher mutual information. Consistent with the logic of the task-dynamic analysis, execution variables had to be separated from result variables. Given the extended model, 4 execution variables fully defined a strategy: amplitude A (which was prescribed), frequency f, stiffness K, and damping B. Note that initial ball angle and velocity ϕ_0 and $\dot{\phi}_0$ are essentially included in the choice of f since the movement frequency dictates the behavior of the ball (in-phase or anti-phase).

Figure 14 displays the three-dimensional execution space spanned by the execution variables f, K, and B as simulated from the model equations (A was fixed). For each strategy (f, K, B), predictability was quantified by mutual information. To address the alternative hypothesis, minimizing exerted force, MSF was also calculated. In Figure 14(A), the areas of high values of MI are shown in pink, in Figure 14(B) the areas of low exerted force are displayed in green. The blue points represent the execution variables or strategies for each trial estimated for each subject. Note that the blue data are identical in both figures. Both figures clearly show that none of the subjects chose a strategy that lay in areas with low predictability. Conversely, very few trials overlapped with low MSF solutions. Comparison between Figures 13 and 14 reveals



Figure 13. Bode amplitude and phase plots of the linearized coupled model, for different values of hand impedance. (A) Stiffness (K) = 100 N/m and damping coefficient (B) = 10 N·s/m, typical for the low-frequency group. (B) K = 200 N/m and B = 15 N·s/m, typical for the high-frequency group. Phase plots of the cart and pendulum are superimposed to highlight the synchronization of their movements. For comparison, the gray histogram represents the distribution of frequencies adopted by participants in the experiment (identical to Figure 11).

that the two resonance frequencies of the system coincided with areas of high MI (pink shading). This suggests that the behavior of the system is easily predictable when oscillating at a resonance frequency. Conversely, the anti-resonance frequency coincides with a region of low MI; therefore, the behavior of the system is hard to predict when oscillating at or around the anti-resonance frequency. These results are consistent with both the primary hypothesis that subjects seek predictability, but also that this coincides with exploiting the resonances of the coupled system.

7. Discussion

Research in human motor neuroscience to date has focused on simple unconstrained movements, predominantly reaching or pointing. The evident rationale has been that one should study simple problems first and control for confounding variability. However, daily activities are full of much richer coordination demands that humans skillfully manage with great ease. To gain insight into biological control principles that can then inform robot control, such problems deserve attention.

Unfortunately, principles that have been identified for the control of unconstrained movements do not easily 'scale up' to more complex problems. This was highlighted in a straightforward simulation using a minimum jerk trajectory for moving the cup-and-ball system. Specifically, in movements involving interactions with objects new challenges arise: complex interaction forces are generated that present continuous 'perturbations' that the controller has to deal with. In fact, introducing extended contact with objects can quickly generate chaos (in the technical sense), as the inverse dynamics simulations demonstrated. It is these complex interactions that also present challenges for the control of robots. Given the simultaneous nature of interaction forces, feedback-based corrections are insufficient. In humans, the long delays and the high levels of noise in the neuromotor system imply that some form of predictive feedforward control is necessary; but control based on internal models of complex nonlinear and potentially chaotic dynamics seems daunting, both for humans and may be also for robots. Are there simpler and more elegant ways to achieve dexterity?

In order to gain more insight into the control architecture in humans, we adopted a task-based approach:



Figure 14. (A) 3-D plots of the mutual information in the space of the execution variables frequency *f*, stiffness *K*, and damping *B*. Pink shading represents areas of relatively high mutual information. (B) 3-D plots of the mean square exerted force in the space of the same three execution variables. Green shading represents areas of relatively low force.

Rather than starting with assumptions or hypotheses about the controller directly, the task and its possible solutions to a task goal were analyzed first. To facilitate this task-dynamic analysis, the object was simplified to the cart-and-pendulum model to afford insights and intuitions. Importantly, this model retained the essential challenges of an underactuated nonlinear object. All four studies adopted this approach, although in different implementations.

In the first study, the physical model of the cup-andball object alone was the basis for identifying the contraction regions; the controller was not included in the analysis. Similar to the third study, the calculated contraction regions were those of the free, unforced system. Based on this analysis, the subjects' trajectories were examined whether they intersected those regions of stability. The fact that subjects' trajectories indeed visited those subspaces after some practice suggested that subjects shaped their movements system to regions where it was naturally or passively contracting, without the need for external stabilizing control effort. Following these encouraging results, the next study used more recently developed tools in contraction analysis to include the controlled system into the analysis (object+human controller). The framework of control contraction metrics was adapted such that it could be used not only as a framework for the synthesis of a controller, but also as a tool for the analysis of stability in experimentally acquired human movements. This allowed to examine whether the human control policy actively sought to stabilize the trajectories and make them contracting. Note that this analysis still refrained from making specific assumptions about the human controller: The control input values were analyzed at each data point and evaluated for whether they satisfied a necessary criterion for contractivity. The results of that study further supported our hypothesis that humans seek stability, and thereby predictability, in these complex object manipulation tasks.

The two experiments involving extended rhythmic movements elicited more complex trajectories that further probed into the question whether subjects learned to make interactions predictable. The hypotheses were assessed by evaluating the human executions of this task with mutual information between the controller and the resulting object dynamics. Again, following the same mantra to not a priori assume a controller, the first step was to parse the dynamical system into execution and result variables. By mapping the execution variables into the result variables, result spaces were generated from all possible execution variables. Quantitative hypotheses on the chosen result variables were formulated and the human executions were directly evaluated in the result spaces generated for selected hypotheses. The human data conformed to strategies that pursued predictable interactions and rejected minimization of force and maximization of smoothness.

While the first rhythmic study only used inverse dynamics without any assumptions about the controller, the second study went beyond this minimal approach and coupled hand and object with a hand impedance. However, only two minimal assumptions were made about the controller: First, the input force (F_{input}) was assumed to be equal to the force required to move a rigid object of mass equal to the cup-and-ball system in sinusoidal fashion. Second, the hand impedance was assumed to be constant throughout any given trial. While these choices can be changed to better represent reality, such 'better models would become more complex and probably add little to our understanding of the core problem. After modifying the model to include the hand-object coupling, the same task-based analysis steps as the first study were followed. Using forward simulations, the results again favored strategies that achieved predictable interactions. Further, these predictable interactions were achieved by moving at the resonance frequencies of the coupled system. It may be concluded that humans tuned their hand impedance to match the resonance frequencies with those that they were able to perform.

Our research used two different metrics to quantify predictability: stability and mutual information. While these two metrics originate in two different theoretical frameworks, dynamical systems theory and information theory, the latter has developed a range of computational tools that have been successfully applied to dynamical systems to quantify predictability [51,52]. Intuitively, if a system is evolving along its stable attractor, then the trajectory will be confined to a restricted range in state space and the same states will be revisited over and over. Therefore, we expect to see shared information between any two observed variables of the system, rendering high mutual information: knowing one variable reduces the uncertainty in the other. Alternatively, when at a chaotic attractor, there is little shared information between states now and at a future state. We opted to relate the mutual information between hand control and the resulting ball trajectory to quantify predictability of the object dynamics.

Taken together, these four studies converged on a main result: the human controller seeks predictability. However, these results do not yet provide any indication about how this 'high-level' objective is learned and realized by the human controller. Moreover, it is still not clear how the subjects sense what is predictable and how they can distinguish predictable and unpredictable dynamics. Despite these open questions, we posit that the presented results are a first step towards addressing these issues. The findings advance our understanding of how humans control physical interactions with complex objects.

We would like to suggest that these findings may also inspire developments for the control of robots. For example, the results obtained from contraction analysis can be used to design contraction-based controllers for robotic manipulation to generate robust motion plans online. In fact, a recent study had applied control contraction metrics for the control of quadrotors in the presence of cross-wind disturbances [47]. For biomimetic applications, the parameters of the controller, for example the contraction rate, can be set to match those of humans. Moreover, similar to human behavior, the robot can be programmed to seek contraction only when a perturbation arises, in order to save control effort. The results from the study where humans can select their preferred frequency indicate that humans may change their muscle impedance (stiffness and damping parameters) depending on the desired behavior. This 'parametric-control'

could inspire control strategies for robots with elastic components, especially soft robots. Considering the central role of stability, one can envision robotic systems that tune the parameters of their elastic components to make the desired behavior a stable attractor of the system.

Please also note that our simple cart-and-pendulum system inspired by the problem of transporting a nonrigid object is also a simplified version of a system that carries a payload. Imagine a quadrotor transporting a package attached via a pendular string. Even though there is no human hand guiding the quadrotor, the control challenges that are faced in quadrotors with payloads are surprisingly close to those of our simple manual task.

The insights into human behavior are also relevant for settings where robots and humans must physically interact. Developing bio-inspired robot control strategies based that include or respect these principles in human behavior can make robotic movements more legible and predictable to human co-workers in cooperative, collaborative, and comanipulation tasks. However, the exchange of knowledge in humans and machines also works both ways: more advances in robot control theory may also provide useful tools to further our understanding of human movement control and generation. Human movement neuroscience and robotics form a natural synergy and more research studies that exploit this interplay are needed to push the frontiers in both fields.

8. Conclusions

This work presented the task-dynamic approach to understand how humans physically interact with an underactuated nonlinear dynamic object. Given the unpredictable nature of the object dynamics and the interactions, the hypothesis was that humans seek to make their interactions predictable. Using two different metrics for predictability, mutual information and stability, experimental results supported the hypothesis that indeed predictability of the object's dynamics was a primary objective of human control.

Disclosure statement

No potential conflict of interest was reported by the authors.

Notes on contributors

Salah Bazzi received his B.Eng. and Ph.D. in Mechanical Engineering from the American University of Beirut in 2012 and 2017. His is currently a Postdoctoral Research Associate in the Action Lab at Northeastern University. His central research interest is to understand the acquisition and control of skilled human sensorimotor behavior, with the goal of using that knowledge to help endow robots with dexterous physical interaction and manipulation capabilities.

Dagmar Sternad received her B.S. in Movement Science and Linguistics from the Technical University and Ludwig Maximilians University of Munich and her Ph.D. in Experimental Psychology from the University of Connecticut. From 1995 until 2008, she was Assistant, Associate, and Full Professor at the Pennsylvania State University. Since 2008, she holds an interdisciplinary appointment as full professor in the departments of Biology, Electrical and Computer Engineering, and Physics at Northeastern University in Boston. She was recently appointed as University Distinguished Professor. Her research is documented in over 170 peer-reviewed publications, book chapters, and several books. She has had editorial appointments in several scientific journals and was regular member of a NIH study section. Her research has been continuously supported by the National Institutes of Health, National Science Foundation, American Heart Association, Office of Naval Research, and others.

ORCID

Salah Bazzi D http://orcid.org/0000-0002-8631-0426

References

- Mayer HC, Krechetnikov R. Walking with coffee: why does it spill? Phys Rev E. 2012;85(4):046117.
- [2] Goriely A, McMillen T. Shape of a cracking whip. Phys Rev Lett. 2002;88(24):244301.
- [3] Ijspeert AJ. Central pattern generators for locomotion control in animals and robots: a review. Neural Netw. 2008;21(4):642–653.
- [4] Liu GL, Habib MK, Watanabe K, et al. Central pattern generators based on matsuoka oscillators for the locomotion of biped robots. Artif Life Rob. 2008;12(1–2):264–269.
- [5] Ott C, Henze B, Hettich G, et al. Good posture, good balance: comparison of bioinspired and model-based approaches for posture control of humanoid robots. IEEE Rob Autom Mag (RAM). 2016;23(1):22–33.
- [6] Schaal S, Sternad D. Programmable pattern generators. 3rd international conference on Computational Intelligence in Neuroscience; Research Triangle Park (NC). 1998. p. 48–51.
- [7] Sternad D, Dean WJ, Schaal S. Interaction of rhythmic and discrete pattern generators in single-joint movements. Hum Mov Sci. 2000;19(4):627–664.
- [8] Schaal S, Mohajerian P, Ijspeert A. Dynamics systems vs. optimal control a unifying view. Prog Brain Res. 2007;165:425-445.
- [9] Gams A, Nemec B, Ijspeert AJ, et al. Coupling movement primitives: interaction with the environment and bimanual tasks. IEEE Trans Rob (T-RO). 2014;30(4):816–830.
- [10] Travers MJ, Whitman J, Schiebel PE, et al. Shape-based compliance in locomotion. Robotics: Science and Systems; 2016 Jun 18–22; Ann Arbor (MI). 2016.
- [11] Maurice P, Huber ME, Hogan N, et al. Velocity-curvature patterns limit human-robot physical interaction. IEEE Rob Autom Lett (RA-L). 2017;3(1):249–256.
- [12] Flanagan JR, Tresilian J, Wing AM. Coupling of grip force and load force during arm movements with grasped objects. Neurosci Lett. 1993;152(1–2):53–56.
- [13] Flanagan JR, Wing AM. The role of internal models in motion planning and control: evidence from grip force

adjustments during movements of hand-held loads. J Neurosci. 1997;17(4):1519–1528.

- [14] Santello M, Flanders M, Soechting JF. Postural hand synergies for tool use. J Neurosci. 1998;18(23):10105– 10115.
- [15] Fu Q, Santello M. Coordination between digit forces and positions: interactions between anticipatory and feedback control. J Neurophysiol. 2014;111(7):1519–1528.
- [16] Slota GP, Latash ML, Zatsiorsky VM. Grip forces during object manipulation: experiment, mathematical model, and validation. Exp Brain Res. 2011;213(1):125–139.
- [17] Mehta B, Schaal S. Forward models in visuomotor control. J Neurophysiol. 2002;88(2):942–953.
- [18] Gawthrop P, Lee KY, Halaki M, et al. Human stick balancing: an intermittent control explanation. Biol Cybern. 2013;107(6):637–652.
- [19] Insperger T, Milton J, Stépán G. Acceleration feedback improves balancing against reflex delay. J R Soc Interface. 2013;10(79):20120763.
- [20] Dingwell JB, Mah CD, Mussa-Ivaldi FA. Experimentally confirmed mathematical model for human control of a non-rigid object. J Neurophysiol. 2004;91(3):1158– 1170.
- [21] Nagengast AJ, Braun DA, Wolpert DM. Optimal control predicts human performance on objects with internal degrees of freedom. PLoS Comput Biol. 2009;5(6): e1000419.
- [22] Leib R, Karniel A. Minimum acceleration with constraints of center of mass: a unified model for arm movements and object manipulation. J Neurophysiol. 2012;108(6):1646–1655.
- [23] Svinin M, Goncharenko I, Kryssanov V, et al. Motion planning strategies in human control of non-rigid objects with internal degrees of freedom. Hum Mov Sci. 2019;63:209–230.
- [24] Kawato M, Wolpert D. Internal models for motor control. Novartis Found Symp. 1998;218:291–303.
- [25] Kandel ER, Schwartz JH, Jessell TM, et al. Principles of neural science. Vol. 5. New York (NY): McGraw-Hill; 2012, of Biochemistry D.
- [26] Sternad D. Control of intermittent and continuous objects. In: Laumond JP, Mansard N, Lasserre JB, editors. Geometric and Numerical Foundations of Movement. New York (NY): Springer; 2017. p. 301–335.
- [27] Han J. A study on the coffee spilling phenomena in the low impulse regime. Achievements Life Sci. 2016;10(1): 87–101.
- [28] Hasson CJ, Shen T, Sternad D. Energy margins in dynamic object manipulation. J Neurophysiol. 2012;108(5): 1349–1365.
- [29] Sternad D, Hasson CJ. Predictability and robustness in the manipulation of dynamically complex objects. Progress in motor control. 2016. p. 55–77.
- [30] van der Linde RQ, Lammertse P. Hapticmaster–a generic force controlled robot for human interaction. Ind Rob: Int J. 2003;30:515–524.
- [31] Flash T, Hogan N. The coordination of arm movements: an experimentally confirmed mathematical model. J Neurosci. 1985;5(7):1688–1703.
- [32] Uno Y, Kawato M, Suzuki R. Formation and control of optimal trajectory in human multijoint arm movement. Biol Cybern. 1989;61(2):89–101.

- [33] Hogan N, Sternad D. Sensitivity of smoothness measures to movement duration, amplitude, and arrests. J Mot Behav. 2009;41(6):529–534.
- [34] Wolpert DM, Ghahramani Z, Jordan MI. Are arm trajectories planned in kinematic or dynamic coordinates? an adaptation study. Exp Brain Res. 1995;103(3): 460–470.
- [35] Rohrer B, Fasoli S, Krebs HI, et al. Movement smoothness changes during stroke recovery. J Neurosci. 2002;22(18): 8297–8304.
- [36] Svinin M, Goncharenko I, Luo ZW, et al. Reaching movements in dynamic environments: how do we move flexible objects? IEEE Trans Rob (T-RO). 2006;22(4):724–739.
- [37] Strogatz SH. Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering. Boca Raton (FL): CRC Press; 2018.
- [38] Bazzi S, Ebert J, Hogan N, et al. Stability and predictability in dynamically complex physical interactions. IEEE International Conference on Robotics and Automation (ICRA); Brisbane, Australia. IEEE; 2018.
- [39] Lohmiller W, Slotine JJE. On contraction analysis for nonlinear systems. Automatica. 1998;34(6):683–696.
- [40] Bazzi S, Ebert J, Hogan N, et al. Stability and predictability in human control of complex objects. Chaos: Int J Nonlinear Sci. 2018;28(10):103103.
- [41] Bazzi S, Sternad D. Robustness in human manipulation of dynamically-complex objects through control contraction metrics. IEEE Rob Autom Lett (RA-L). 2020;5:2578–2585.
- [42] Sontag ED. Control-lyapunov functions. In: Open problems in mathematical systems and control theory. London: Springer; 1999. p. 211–216.

- [43] Ortega R, Perez JAL, Nicklasson PJ, et al. Passivity-based control of euler-lagrange systems: mechanical, electrical and electromechanical applications. London: Springer Science & Business Media; 2013.
- [44] Manchester IR, Slotine JJE. Control contraction metrics: convex and intrinsic criteria for nonlinear feedback design. IEEE Trans Automat Contr. 2017;62(6):3046– 3053.
- [45] Sontag ED. A Lyapunov-like characterization of asymptotic controllability. SIAM J Control Optim. 1983;21(3): 462–471.
- [46] Parrilo PA. Semidefinite programming relaxations for semialgebraic problems. Math Program. 2003;96(2): 293-320.
- [47] Singh S, Majumdar A, Slotine JJ, et al. Robust online motion planning via contraction theory and convex optimization. IEEE International Conference on Robotics and Automation (ICRA); Singapore. IEEE; 2017. p. 5883–5890.
- [48] Cover TM, Thomas JA. Elements of information theory. Hoboken (NJ): John Wiley & Sons; 2012.
- [49] Nasseroleslami B, Hasson CJ, Sternad D. Rhythmic manipulation of objects with complex dynamics: predictability over chaos. PLoS Comput Biol. 2014;10(10): e1003900.
- [50] Maurice P, Hogan N, Sternad D. Predictability, force and (anti-) resonance in complex object control. J Neurophysiol. 2018;120(2):765–780.
- [51] Shaw R. Strange attractors, chaotic behavior, and information flow. Z Naturforsch A. 1981;36(1):80–112.
- [52] Shaw R. The dripping faucet as a model chaotic system. Santa Cruz (CA): Aerial Press; 1984.