Visual dynamics cues in learning complex physical interactions

Christopher J. Hasson and Paneed F. Jalili

Neuromotor Systems Laboratory; Northeastern University, Boston MA
Departments of Physical Therapy, Movement and Rehabilitation Sciences, Bioengineering, and Biology

ABSTRACT
This study investigated the role of visual dynamics cues in learning to interact with a complex physical system. Manual gait training was used an exemplary case, as it requires therapists to control the non-trivial locomotor dynamics of patients. A virtual analog of this activity was developed that allowed naïve subjects to manipulate the leg of a virtual stroke survivor (a virtual patient; VP) walking on a treadmill using a small robotic manipulandum. The task was to make the VP's leg pass through early- mid-, and late-swing gait targets. One group of subjects (n = 17) started practice seeing the VP's affected thigh and shank (visual dynamics cues); a second group (n = 16) only saw the point-of-contact (VP ankle). The results showed no learning benefits from visual dynamics cues. Subjects who began practice seeing the VP's leg were unable to fully adapt when the visual cues were removed in a generalization task. Full adaptation, defined as low errors on all gait targets, only occurred after subjects gained experience practicing without visual cues. This suggests that learning with visual dynamics cues may reinforce a dynamics representation associated with those cues, which may not generalize well if the visual context is changed.

INTRODUCTION
The performance of daily life tasks requires humans to interact with both familiar and unfamiliar objects with a range of dynamic complexities. Such activities include opening a door, drinking tea, driving a car, or interacting with another person. It is thought that humans accomplish this in part by internally modeling the dynamics of the interactions. An internal dynamics model (or representation) relates motor commands to object behavior as detected by the body's sensors. These models can be used to predict how an object (or another person) will respond to a motor input, and what motor commands are needed to elicit a desired interaction[1].

Proprioception has an established and critical role in learning novel dynamics[2-4]. However, visual cues can provide information not available to proprioception, which may aid in the development and recall of internal models. The visual presentation of object features (e.g., size, shape, or configuration) can inform expectations about static and dynamic object behavior. Seeing the geometry of familiar tools allows humans to correctly brace against reaction forces[5] and can improve generalization[6]. Sighted object features can assist with estimation of object stability characteristics, mass distribution, and material properties[7-9], and may influence how dynamics are represented in the nervous system[10]. Although prior studies have shown that visual dynamics cues can benefit learning[5-10], they have primarily used objects with relatively simple dynamics (e.g., a hammer, spring, or block). It is otherwise unclear whether the benefits would be similar for more complex physical interactions.

An archetypal complex physical interaction is with another person. As an exemplary case, manual gait training requires therapists to control the complex locomotor dynamics of patients. Following along with a patient's swinging leg and applying forces to achieve a desired kinematic trajectory is not trivial. A therapist must compensate for numerous forces, including those related to linked-segmental dynamics (e.g., gravitational, inertial, Coriolis, and centrifugal forces), as well as those from a patient's muscular contractions. A recent study used a virtual analog of manual gait training to show that healthy adults are capable of forming internal representations of pathologic locomotor dynamics[11]. However, this study did not manipulate the visual information provided to subjects, i.e., subjects could always see the simulated patient's leg during manipulation. Thus, it is unclear whether the visual dynamics cues had a significant impact on the learning outcomes.

Therefore, this study aimed to investigate the role of visual dynamics cues in learning to interact with a complex physical system in the form of a locomoting virtual patient (VP). Naïve subjects were asked to manipulate the leg of a virtual stroke survivor walking on a treadmill, which was modeled as an end-point impedance that subjects interacted with using a small robotic manipulandum[11]. Subjects were asked to apply forces to the VP's affected leg to make the VP's ankle pass through three gait targets, which would normalize the gait pattern with respect to the unaffected leg. One group of subjects (n = 17) started practice seeing the VP's affected thigh and shank; a second group (n = 16) only saw the point-of-contact (the VP's ankle). It was hypothesized that subjects who saw the VP's leg during practice would show improved task performance and improved generalization relative to the control group.

Address for Correspondence
Christopher J. Hasson
Northeastern University | 360 Huntington Ave | 301 Robinson Hall | Boston, MA 02115
E-mail: c.hasson@northeastern.edu

METHODS

Virtual Patient (VP) Model

The VP was modeled as an end-point impedance that subjects interacted with using a small robotic manipulandum (GeoMagic Touch, 3D Systems, Andover MA). Full details of the model, which was based on previously collected data on stroke survivors[12], are reported in Hasson and Goodman[11]. Using this stroke survivor data (the affected leg), a reference VP kinematic trajectory was specified (Figure 1; red trajectory). The step pattern associated with this reference trajectory had a smaller step height and width compared to the unaffected leg of the stroke survivors. When subjects applied a force that made the manipulandum deviate from the reference trajectory, they experienced an impedance force F_{IMP} as

$$F_{IMP} = a(F_D + F_M + F_T) \qquad (1)$$

where the quantities F_D, F_M, and F_T are vectors of forces that represent the anterior-poster (x) and vertical (y) components of impedance forces due to rigid body dynamics (F_D), neuromuscular contributions (F_M), and the treadmill interface (F_T). The scaling factor a was determined through experimentation ($a = 1/12$), and was chosen so that the forces required to move the VP into a "healthy" trajectory were large enough to be significant but not fatiguing (1-2 N).

Model Details

The rigid body forces (F_D) that subjects felt when manipulating the VP's leg were specified by

$$F_D = J^{-T}[I(\theta)\ddot{\theta}_{err} + V(\theta,\dot{\theta})\dot{\theta}_{err} + G(\theta)\theta_{err}] \quad (2)$$

where $I(\theta)$, $V(\theta,\dot{\theta})$ and $G(\theta)$ are matrices that specify geometry-dependent inertial torques, centrifugal and Coriolis torques, and gravitational torques, respectively; the quantities θ_{err}, $\dot{\theta}_{err}$, and $\ddot{\theta}_{err}$ are difference vectors comparing the segmental kinematics associated with the stroke reference trajectory (x_R, $\dot{x}_R$, and $\ddot{x}_R$ in Figure 1) and those associated with the actual state of the VP (i.e., manipulandum position; x_M, $\dot{x}_M$, and $\ddot{x}_M$ in Figure 1), and J^{-T} is the inverse transpose of the Jacobian, which transforms joint torques to end-point (ankle) forces (details for I, V, G and J^{-T} matrices are in Hasson and Goodman[11]).

Neuromuscular impedance forces (F_M) depended on three quantities: a muscular spring stiffness K_{SPR}, a geometric stiffness K_{GEO}, and a damping factor B. The quantity K_{SPR} modeled the spring-like action of human musculotendon units, K_{GEO} modeled how the limb endpoint stiffness due to muscular action is affected by skeletal geometry in the presence of a contact force[13], and damping B simulated the viscosity-like effects of

Figure 1. Schematic of virtual patient model. When a subject moved the manipulandum in a way that caused the virtual patient kinematics (x_M, $\dot{x}_M$, and $\ddot{x}_M$, i.e., the virtual patient's ankle position, velocity, and acceleration, respectively) to deviate from the reference stroke model kinematics (x_R, $\dot{x}_R$, and $\ddot{x}_R$; red ankle trajectory), the subject felt an impedance force that depended on programmed rigid body, muscular, and treadmill dynamics. The impedance force was equal and opposite from the assistance force provided by the subject (F_S).

velocity-dependent muscular properties[14]. These viscoelastic effects were expressed as

$$F_M = (K_{SPR} + K_{GEO})p_{err} + B\dot{p}_{err}. \qquad (3)$$

where p_{err} and $\dot{p}_{err}$ are difference vectors between the stroke reference and actual VP ankle (manipulandum) positions and velocities, respectively (details for K_{SPR}, K_{GEO}, and B matrices in Hasson and Goodman[11]).

Finally, a treadmill normal force prevented the VP's ankle from penetrating the treadmill. A virtual spring (stiffness $k_N = 2.0$ N/mm) slid along a semicircular arc (the treadmill) beneath the VP's ankle position. A semicircular treadmill was needed because the VP's leg was short one degree of freedom (the foot). An additional anti-slip spring ($k_S = 0.7$ N/mm) prevented slipping due to frictional forces, and simulated the effect of the treadmill in dragging th e foot rearwards during the stance phase. The net treadmill impedance force (F_T) was given by

$$F_T = (k_N\hat{p}_{err}) + (k_S p_{err}) \qquad (4)$$

where $\hat{p}_{err}$ is the position of the VP's ankle projected onto the surface of the treadmill.

Subjects

Thirty-three young adults were randomly assigned to either a visual dynamics cues (VDC) group (n=17; 10 female, 7 male; age: 22 ± 1.7 yrs [mean ± SD]; height: 1.7 ± 0.1 m; weight: 67 ± 11 kg;) or a control group (n=16; 11 female, 5 male; age: 24 ± 4 yrs [mean ± SD]; height: 1.7 ± 0.1 m; weight: 67 ± 10 kg). The VDC group was able to see the VP's leg during practice and the control group never saw the VP's leg. All subjects were healthy and free from neurological and musculoskeletal impairments affecting upper extremity control and gave informed consent. The study was approved by the Northeastern University Institutional Review Board and all methods were performed in accordance with the relevant guidelines and regulations.

Task Instructions & Goal

Subjects were told that they were interacting with a simulation of a stroke survivor (VP) walking on a treadmill, and that if they pushed or pulled on the manipulandum, the force is transferred to the VP's ankle. Instructions were to make the VP's ankle pass through three gait targets (Figure 2). These targets were: early-swing, mid-swing, and late-swing, corresponding to 78, 88, and 100% of the VP's gait cycle, respectively. These targets were selected from the average ankle trajectory of the unaffected leg for the stroke survivors upon which the model was based (hitting the gait targets increased VP gait symmetry).

Visual Display

All subjects saw the VP's ankle (the interaction point), the three gait targets, a curved line representing the ground, and a magenta "tail" that followed the ankle to provide a visual indication of the ankle's two-dimensional path (Figure 2). Importantly, the VDC group also saw two linked segments representing the VP's thigh and shank. For both groups, a small rectangular tracker was presented when the VP's leg was in contact with the treadmill to help subjects maintain proper step timing (because other timing-related visual information normally available during real gait training, such as a patient's contralateral leg, was absent). The timing cue slid along the treadmill, becoming visible at heel-strike and disappearing at push-off. An audible beep was provided at heel-strike to assist with pacing. All programming associated with the virtual patient model and visual display was performed in C++ (Microsoft Visual Studio 2010 Professional, Version 10.0.30319.1 RTMRel; Microsoft Corporation, Redmond, WA).

Initial Practice

Both groups began practice in a null-dynamics condition (two trials of 80 steps; Figure 3). The VDC group saw the

Figure 2. Illustration of visual feedback seen by study participants. The visual dynamics cues (VDC) group saw the affected leg of a virtual patient locomoting on a treadmill (left image) while a control group only saw the end-point (i.e., the ankle; right image). The white dots show early-, mid-, and late-swing gait targets. The red curve at the bottom represents the treadmill. The images show the affected leg during the swing phase just after passing the mid-swing target, with the leg moving towards the right. The unaffected leg was not shown.

VP's leg and the control group did not. In this condition, subjects only experienced the treadmill model and felt no impedance forces from the VP (i.e., $F_D = F_M = 0$). The purpose of the null-dynamics condition was to allow subjects to become familiar with the task and interface. Next, subjects practiced the task with impaired-dynamics (Impaired Dynamics I). In this condition, subjects felt as if they were holding onto the leg of a miniature stroke patient (the VP) and had to apply appropriate forces to move the VP's ankle through the three gait targets. In this condition, subjects performed four 80-step trials; the leg remained visible for the VDC group and absent for the control group. A one-minute rest was provided between each trial.

Generalization Tests

After initial practice, subjects performed two generalization tests. The early-swing gait target was shifted 1.45 mm upwards (a 42% increase) and 0.24 mm to the right (see Figure 3). The mid-swing gait target height was increased by 1.21 mm. The late-swing target was not changed. These target shifts required subjects to make the VP perform an exaggerated heel-lift during push-off. To facilitate group comparisons the generalization test was performed without the VP's leg displayed so that both groups would have the same visual display (facilitating group comparisons). The first 30-step generalization test (Generalization I) was performed after Impaired Dynamics I. After more practice with the original targets (Impaired Dynamics II), a second 30-step generalization test was performed (Generalization II), but in this case, the VDC group had the opportunity to practice

Figure 3. Experimental protocol for virtual locomotor manipulation task. One group of subjects began practice able to see the virtual patient's thigh and shank segments (visual dynamics cues group); the other group only saw the ankle.

without seeing the leg before the target locations were changed (Impaired Dynamics III). This enabled the separation of two effects in the VDC group, losing the leg visual and switching target locations (Generalization I did both at the same time).

Probe Trials

Errors made during unexpected changes in dynamics can reveal information about human sensorimotor control strategies[15]. Thus, each subject experienced four unexpected null-steps during the last practice trial before the second generalization test (during Impaired Dynamics III; see Figure 3). In these steps, the VP rigid-body and neuromuscular impedance forces were turned off during the middle of the stance phase (the treadmill model remained in effect). Mid-stance was chosen as the initiation point because kinematic errors (and hence F_D and F_M impedance forces) are negligible at this time. The VP impedance forces remained off throughout the stance phase and were turned back on after the VP's leg completed its swing (when the forward velocity of the VP's ankle changed sign). The null-step numbers were selected from uniformly distributed pseudo-random integers in the span of 5-80. Each null-steps had to be at least 4 steps from neighboring probe steps, if not a new random number draw was performed until the requirement was satisfied.

Data Reduction and Analysis

The root-mean-squared error (RMSE) quantified task performance, determined as the smallest Euclidean distance between the manipulandum trajectory (the VP's ankle) and the early-, mid-, and late-swing gait targets (in two-dimensions). The RMSE was averaged across the first four steps of each practice condition to capture early practice performance, and across the last 20 steps to capture late practice performance (a shorter time scale was used for early practice because the changes are more rapid). Outliers were detected using Tukey fences[16]. If for a given condition and time point (early or late practice), an RMSE was outside ± 2.7 standard deviations (a 99.3% confidence interval), the value was excluded from further analysis. Analysis performed in MATLAB (ver. 9.4, R2018a; Mathworks, Natick MA).

Statistics

The effects of visual dynamics cues on learning locomotor dynamics were assessed with a mixed models analysis. Group (VDC or control), Condition (Null Dynamics, Impaired Dynamics I, Generalization I, Impaired Dynamics II, Impaired Dynamics III, and Generalization II), and time (early vs. late practice within a condition) were included as fixed effects, with repeated measurements on condition and time (first-order autoregressive covariance structure). Subject ID was included as a random effect. The dependent variables were the RMSE for the early-, mid-, and late-swing gait targets. A second mixed-models analysis was performed for the random null step trials during Impaired Dynamics III, with Group and Null Step Number (1-4) as fixed effects (repeated measurements on Null Step Number). The alpha level was .05 for all tests, which were performed in SPSS (ver. 25; IBM, Armonk, New York) and all tests were two-tailed. The Benjamini-Hochberg step-up procedure[17] was used to control for multiple comparisons with a tolerable false discovery rate (FDR) of 0.1[18].

RESULTS

Subject skill was assessed by the RMSE between the manipulandum (the VP's ankle) trajectory and the early-, mid-, and late-swing gait targets. There were no group differences in RMSE for the early swing target (Figure 4; no group main effect; $F_{1,79} = 0.50$, p = .482), and both groups reduced their error from early to late practice within each condition (time main effect; $F_{1,151} = 45.71$, p < .001). For the middle- and late-swing targets, there was a three-way interaction, indicating that the group differences depended on condition and time (middle-swing: $F_{5,204} = 8.48$, p < .001; late-swing: $F_{5,280} = 2.43$, p = .047).

Pairwise comparisons showed that for the mid-swing target during early practice in Null Dynamics, the VDC group had larger errors than the control group (early practice p < .001), such that the VDC group significantly overshot the target (signed error = 0.88 ± 0.82 mm; mean ± s.d.). The mid-swing RMSE became similar between the groups by the end of Null Dynamics (p = .128), and the VDC group had a larger RMSE than controls during Impaired Dynamics II, but only for early practice (early p = .024; late p = .562). For the late-swing target, the VDC group had a larger RMSE during Generalization I (early p < .001; late p = .085), and this larger error persisted through Impaired Dynamics II (early p = .043; late p = .015) and Impaired Dynamics III (early p = .016). By the end of Impaired Dynamics III, the group RMSE's were similar (late p = .094). Larger late-swing errors for the VDC group were associated with early termination of the VP's leg swing (e.g., the ankle fell short of the target by 0.60 ± 0.57 mm in early Generalization I; mean ± s.d.). Pairwise comparisons remain significant after Benjamini-Hochberg corrections for multiple comparisons.

Near the end of practice, during Impaired Dynamics III, subjects were exposed to four unexpected null-dynamics steps to probe the strength of participants' internal representations. During these steps, subjects did not feel any impedance forces as they tried to lift the VP's foot from the ground. For the early-swing target, the control group made larger errors (higher RMSE) than the VDC group, overshooting the target on the null-steps (see Figure 5; group main effect; $F_{1,116} = .02$, p = .018). For the early-swing target, there were no main effects of time ($F_{3,116} = 1.83$, p = .146) and no interactions between group and time ($F_{3,116} = 0.41$, p = .748). There were no main effects or interactions for the mid- and late-swing targets (Figure 5).

DISCUSSION

Despite expectations to the contrary, visual dynamics cues provided no learning benefit for humans interacting with a complex physical system in the form of a virtual stroke

Figure 4. Task performance. Average root-mean-squared error (RMSE) for the early, middle, and late-swing gait targets for the visual dynamics cues (VDC) group (red) and the control group (blue) while practicing the virtual locomotor training task. Significant pairwise group differences indicated by (*). Standard errors shown.

survivor walking on a treadmill. Moreover, subjects who began practice with visual dynamics cues (the virtual patient's leg) were unable to fully adapt when the visual cues were removed in a generalization task. Full adaptation, defined as low trajectory errors on all gait targets, only occurred after subjects gained experience practicing *without* the visual cues. This suggests that learning with visual dynamics cues may reinforce a dynamics representation associated with those cues, which may not generalize well if the visual context is changed.

Influence of Visual Dynamics Cues on Initial Dynamics Expectations

Showing the leg during the initial null-dynamics trials appeared to instantiate a specific set of expectations about what impedance forces would be felt by participants. Subjects who could see the VP's leg (VDC group) lifted it too high during the swing phase, while those who only saw the ankle interaction point (control group) did not show this bias. It could be that upon seeing the leg, subjects expected to feel its weight and adjusted by lifting upwards to provide (unnecessary) support. This occurred even though subjects were explicitly told that, during the initial null-dynamics trials, they would not feel forces from the VP's leg and they only had to trace through the air. This is consistent with evidence that explicit knowledge about an object's properties (e.g., mass) supplied by non-proprioceptive modalities is not sufficient to override prior experiences, i.e., sensorimotor errors are needed to update internal models[19].

An alternative explanation is that the large mid-swing target errors arose from the two-thirds power law, i.e., the inverse relation between movement curvature and speed[20]. After passing the early swing target, high curvature of the ankle trajectory is needed: the VP's ankle must arc downward to hit the mid-swing target. Although the VP leg's pendular dynamics would normally facilitate this motion, they were absent in the initial null-dynamics trials. Therefore, subjects may have rounded-off the VP's ankle trajectory instead of slowing down. However, if this were true, it would be expected to affect both groups similarly, yet the control group was able to achieve the necessary curvature. Therefore, it is more likely that the high mid-swing error in subjects who initially saw the VP's leg arose from their prior experiences with rigid bodies (i.e., large body parts are usually heavy), brought about from the visual cues of the leg.

Effects on Task Performance and Generalization

With experience in object manipulation, humans develop internal models that represent an object's dynamic structure (relation between forces and motion as a function of time) and parametric properties (e.g., length, mass, inertia)[21]. One way of evaluating these representations is with generalization tests that assess the degree of carry-over between learning contexts[22]. Generalization was tested by shifting the gait target positions, which required the manipulandum to be moved along a different path. This changed the forces experienced by subjects. The results showed that subjects who had previously practiced seeing the VP's leg (VDC group) fared no better on the generalization tests than those who never saw the leg at all (control group). Interestingly, for the late-swing target, subjects who began practice seeing the VP's leg had greater generalization errors than those that did not, and

Figure 5. Null-step performance. Results for early-, mid-, and late-swing gait target root-mean-squared error (RMSE) for the unexpected null-step trials experienced during the Impaired Dynamics III practice condition. The visual dynamics cues (VCD) group (red) had prior task practice with visual feedback of the virtual patient's leg, the control group (blue) did not. Significant pairwise group differences indicated by (*). Standard errors shown.

only improved late-swing target generalization after they were able to practice *without* seeing the VP's leg for a period.

Seeing the leg may have promoted a leg-geometry specific dynamics representation. That is, subjects may have represented the dynamic structure of the VP using the angles of the VP's thigh and shank. Thus, when the leg visual was removed in the generalization test, it required subjects to modify their internal representations to be compatible with an endpoint-based (Cartesian) dynamics representation. This was not a problem for the control group, who never saw the leg and therefore had more experience with a singular dynamics representation (and singular visual context). In contrast, a recent study by Sadeghi and colleagues[6] showed that subjects who saw the configuration of a virtual tool had better generalization than those who saw an ambiguous tool configuration. Differences in results are likely because the Sadeghi et al. study tested participants under the same visual conditions, i.e., those who learned while seeing the tool geometry also saw it on generalization tests, and vice-versa. The present study used a mixed study design, such that the group that saw the VP' leg during initial practice performed the generalization test *without* seeing the leg, thereby assessing the within-group effect of visual dynamics cue removal. Together, these results suggest that the generalization benefits of visual dynamics cues may only apply in the same sensory context.

Finally, the observed group differences in the unexpected null step responses support the notion that participants' internal representations were specific to the visual context.

The group who practiced seeing the VP's leg had smaller null-step errors for the early-swing target. This would be expected if they were disadvantaged by the need to learn multiple visual context-specific dynamics representations. In general, smaller errors on null-dynamics catch trials (or in this study, catch steps) may reflect high antagonistic muscular co-activation due to greater uncertainty in dynamics representations[23]. Group differences for null-step responses were only observed for the early-swing target, which is likely because subjects could initiate corrective actions by the time the mid- and late-swing targets were reached. In this regard, force channels are a useful type of probe because they clamp kinematic errors and allow longer time scales to be investigated[24]. A previous study using the same VP gait training task implemented force channels by locking the VP's ankle motion to a predefined target trajectory[11]. However, in the present study force channels could not be used because there was no target trajectory, only discrete gait targets.

Limitations

The results discussed above should be interpreted in the context of the study's limitations: 1) Without further experiments, it is unclear whether the results would extend to different motor tasks, including the physical manipulation of another real person. 2) It is impossible to know precisely what subjects were looking at without eye-tracking measurements, i.e., whether they foveated on the leg segments or something else. 3) The extent to which the control group's performance on the task depended on them knowing that they were interacting with a locomoting VP is unknown. The task instructions were the same for all subjects so that any observed differences would be due to the manipulation of the visual dynamics cues and not asymmetrical withholding of information.

Implications

These results suggest that it is best to practice under the same visual conditions one would expect to encounter normally. If at some point visual dynamics cues may become unavailable or unreliable (e.g., driving in low-visibility), it may be better to learn the task without them, as there did not appear to be any learning deficit in the group that never saw the VP's leg. This may be of particular relevance in virtual reality (VR) and augmented reality (AR), which can readily add or overlay visual information that is normally unavailable, for example through the use of interactive mirrors[25,26]. To date, the generalization and transfer of skills learned in VR or AR have been equivocal, with some showing a lack of transfer[27,28] but others showing some carryover for simple actions[29]. The present results may offer an alternative interpretation. It may be that with more complex tasks there is more opportunity for augmented visual feedback

to impose a coordinate system on the learner that does not generalize when augmented feedback is absent.

ACKNOWLEDGMENTS

Funded by Northeastern University and the National Center for Advancing Translational Sciences (NIH# UL1TR002544).

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

Wolpert, D. M., Ghahramani, Z. & Flanagan, J. R. Perspectives and problems in motor learning. *Trends in Cognitive Sciences* **5**, 487-494 (2001).

Krakauer, J. W., Ghilardi, M.-F. & Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neuroscience* **2**, 1026 (1999).

Sainburg, R. L., Ghilardi, M. F., Poizner, H. & Ghez, C. Control of limb dynamics in normal subjects and patients without proprioception. *Journal of Neurophysiology* **73**, 820-835 (1995).

Tong, C., Wolpert, D. M. & Flanagan, J. R. Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. *Journal of Neuroscience* **22**, 1108-1113 (2002).

Ingram, J. N., Howard, I. S., Flanagan, J. R. & Wolpert, D. M. Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Current Biology* **20**, 618-623 (2010).

Sadeghi, M., Sheahan, H. R., Ingram, J. N. & Wolpert, D. M. The visual geometry of a tool modulates generalization during adaptation. *Scientific Reports* **9**, 2731 (2019).

Cholewiak, S. A., Fleming, R. W. & Singh, M. Perception of physical stability and center of mass of 3-D objects. *Journal of Vision* **15**, 13-13 (2015).

Samuel, F. & Kerzel, D. Is this object balanced or unbalanced? Judgments are on the safe side. *Journal of Experimental Psychology: Human Perception and Performance* **37**, 529 (2011).

Buckingham, G., Cant, J.S. & Goodale, M. A. Living in a material world: how visual cues to material properties affect the way that we lift objects and perceive their weight. *Journal of Neurophysiology* (2009).

Ahmed, A. A., Wolpert, D. M. & Flanagan, J. R. Flexible representations of dynamics are used in

object manipulation. *Current Biology* **18**, 763-768 (2008).

Hasson, C. J. & Goodman, S. G. Learning to shape virtual patient locomotor patterns: Internal representations adapt to exploit interactive dynamics. *Journal of Neurophysiology* **121**, 321-335 (2019).

Olney, S. J. & Richards, C. Hemiparetic gait following stroke. Part I: Characteristics. *Gait & Posture* **4**, 136-148 (1996).

English, C. *Stiffness behaviour in two degree of freedom mechanisms* Doctor of Philosophy thesis, Carleton University, (1999).

Winters, J. M. in *Multiple Muscle Systems, pp.* 69-93 (Springer, 1990).

Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* **14**, 3208-3224 (1994).

Tukey, J. W. *Exploratory data analysis*. Vol. 2 (Reading, Mass., 1977).

Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)* **57**, 289-300 (1995).

McDonald, J. Multiple comparisons. *Handbook of Biological Statistics* **3** (2014).

Nowak, D. A. & Hermsdörfer, J. Sensorimotor memory and grip force control: does grip force anticipate a self-produced weight change when drinking with a straw from a cup? *European Journal of Neuroscience* **18**, 2883-2892 (2003).

Lacquaniti, F., Terzuolo, C. & Viviani, P. The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica* **54**, 115-130 (1983).

Braun, D. A., Mehring, C. & Wolpert, D. M. Structure learning in action. *Behavioural Brain Research* **206**, 157-165 (2010).

Krakauer, J. W., Mazzoni, P., Ghazizadeh, A., Ravindran, R. & Shadmehr, R. Generalization of motor learning depends on the history of prior action. *PLoS Biology* **4**, e316 (2006).

Mitrovic, D., Klanke, S., Osu, R., Kawato, M. & Vijayakumar, S. A computational model of limb impedance control based on principles of internal model uncertainty. *PloS One* **5**, e13601 (2010).

Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z. & Mussa-Ivaldi, F. A. Persistence of motor adaptation during constrained, multi-joint, arm movements. *Journal of Neurophysiology* **84**, 853-862 (2000).

Anderson, F., Grossman, T., Matejka, J. & Fitzmaurice, G. in *Proceedings of the 26th annual ACM symposium on user interface software and technology.* 311-320 (ACM).

Tang, R., Yang, X.-D., Bateman, S., Jorge, J. & Tang, A. in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems.* 4123-4132 (ACM).

Wierinck, E., Puttemans, V., Swinnen, S. & van Steenberghe, D. Effect of augmented visual feedback from a virtual reality simulation system on manual dexterity training. *European Journal of Dental Education* **9**, 10-16 (2005).

de Mello Monteiro, C. B. *et al.* Transfer of motor learning from virtual to natural environments in individuals with cerebral palsy. *Research in Developmental Disabilities* **35**, 2430-2437 (2014).

Vander Linden, D. W., Cauraugh, J. H. & Greene, T. A. The effect of frequency of kinetic feedback on learning an isometric force production task in nondisabled subjects. *Physical Therapy* **73**, 79-87 (1993).