

Prediction Precedes Control in Motor Learning

J. Randall Flanagan,^{1,*} Philipp Vetter,²
Roland S. Johansson,³ and Daniel M. Wolpert²

¹Department of Psychology
and Centre for Neuroscience Studies
Queen's University
Kingston, Ontario K7L 3N6
Canada

²Sobell Department of Motor Neuroscience
Institute of Neurology
University College London
Queen Square
London WC1N 3BG
United Kingdom

³Section for Physiology
Department of Integrative Medical Biology
Umeå University
SE-901 87 Umeå
Sweden

Summary

Skilled motor behavior relies on the brain learning both to control the body and predict the consequences of this control. Prediction turns motor commands into expected sensory consequences [1], whereas control turns desired consequences into motor commands. To capture this symmetry, the neural processes underlying prediction and control are termed the forward and inverse internal models, respectively [2–5]. Here, we investigate how these two fundamental processes are related during motor learning. We used an object manipulation task in which subjects learned to move a hand-held object with novel dynamic properties along a prescribed path. We independently and simultaneously measured subjects' ability to control their actions and to predict their consequences. We found different time courses for predictor and controller learning, with prediction being learned far more rapidly than control. In early stages of manipulating the object, subjects could predict the consequences of their actions, as measured by the grip force they used to grasp the object, but could not generate appropriate actions for control, as measured by their hand trajectory. As predicted by several recent theoretical models of sensorimotor control [6–8], our results indicate that people can learn to predict the consequences of their actions before they can learn to control their actions.

Results and Discussion

Subjects were required to grasp an object with a precision grip and move it along a straight line. During the movement, a novel viscous load, which perturbed the load experienced by the fingers and therefore the hand path, was applied to the object (see the Experimental

Procedures for details). Figure 1 shows, for a single subject, the hand path (top trace) and the grip (middle) and load (bottom) force records from the first 10 trials under the novel viscous load as well as trials 15 and 20 and every 10th trial thereafter. The first three trials are the warm-up trials in which the viscosity coefficient was incrementally increased. The hand path is shown in the coronal plane, and the start point of each path is on the left (as though viewing the path from in front of the subject). In early trials, the load caused an upward perturbation of the hand path. The curvature and length of the hand path decreased gradually, and roughly straight paths were often observed by trial 70. Thus, this subject only gradually exerted control over the load so as to produce the straight-line hand paths observed in point-to-point arm movements without unusual loads [9–11].

In contrast, good grip force prediction was quickly established. In the first few trials, clear reflex-mediated increases in grip force were observed. For example, in the first full load trial (trial 4; see inset in Figure 1), a sharp increase in grip force (see arrow) was observed during the movement, and grip force continued to increase after the load force peak and reached its peak about 90 ms later. Such reflexive grip force increases are indicative of poor load force prediction [12]. However, after about four full load trials, reflexive increases in grip force were seldom observed, and the grip and load force peaks coincided closely in time (see, for example, trial 10 in the inset).

Overall, there was a strong relationship between the grip force and load force magnitudes. A reliable relationship between the peak forces was observed ($r = 0.69$; $p < 0.001$), and the slope and intercept of the linear regression line were 2.10 and 0.62 N, respectively. Reliable correlations were also observed when fitting the data from each subject separately ($p < 0.001$ in all eight cases). To assess the temporal coordination of grip and load forces, we examined the timing of the peak rates of increase in grip force and load force that occurred during the initial phase of the movement. A reliable relationship between the peak force rate times was observed ($r = 0.82$; $p < 0.001$), and the slope and intercept of the linear regression line were 1.12 and -0.04 s, respectively. Thus, the peaks in grip force rate and load force rate coincided closely in time. When fitting the data from each subject separately, significant correlations were also observed ($p < 0.001$ in all eight cases).

To quantify trajectory learning, we computed, for each trial, the distance traveled by the hand path. Figure 2A shows the mean path distance, averaged across subjects, as a function of trial. The open circles represent the three warm-up trials, and the points to the right represent the ten replication trials. The figure shows that trajectory learning was gradual. An exponential of the form $y = ae^{bx} + c$ fit to the mean data points yielded a half-life of 19.8 trials. Nonlinear regression revealed that all three parameters of the exponential were significant ($p < 0.05$). (Note that, for each of the exponential fits

*Correspondence: flanagan@psych.queensu.ca




Figure 1. Hand Paths and Force Profiles for Selected Trials Moving the Object with Novel Dynamics

For each trial, the hand path is shown and the grip force (thick trace) and load force (thin trace) records are shown below. The dashed line represents zero force. The first three trials are warm-up trials in which the load was incrementally increased. The inset shows grip force and load force records from trials 4 and 10. For both trials, the left margin of the gray bar is aligned with peak load force, and the width of the bar is 100 ms. The open circles indicate the grip and load force at the start and end of this 100-ms epoch.

described in this paper, very similar parameter estimates and confidence intervals were obtained when fitting the mean data points averaged across subjects and the individual data points from all subjects.)

Figure 2B shows the percentage change in grip force (from the time of peak load to 100 ms later) as a function of trial. In sharp contrast to the gradual learning curve observed for trajectory learning, this measure decreased rapidly over the initial 5–10 full load trials and then leveled off. The exponential fit revealed a half-life of 2.6 trials, and all parameters were significant ($p < 0.05$). Thus, although grip force failed to accurately predict the novel load in the first few trials, good prediction developed rapidly such that reactive increases in grip force were no longer observed.

At the end of learning, replication trials were performed in which subjects were required to produce trajectories that matched those from the first ten full load trials in terms of the path (see the Experimental Procedures for details). Importantly, the percentage change in grip force in the replication trials was far lower than in the initial few full load trials. Thus, despite moving the object with similar kinematics as in the initial trials, at the end of learning, subjects adjusted grip force predictively. This result indicates that the large percent increases in grip force observed in the initial learning trials were not merely a byproduct of curved hand paths. The percentage change in load force (from the time of peak load to 100 ms later) is shown in Figure 2C. On average, there was about a 10% decrease in load force 100 ms after the peak. In contrast, the percentage change in grip force – in the later trials, when reactive increases in grip force were no longer observed – was close to zero. This reflects the fact that reductions in grip force tend to be more sluggish than reductions in load force (e.g., [12–14]).

Warm-up trials were included so as not to startle our subjects, as this might lead to excessive grip forces when first experiencing the novel dynamics. The results shown in Figure 2D indicate that this approach was successful. The figure shows the mean grip force to load force ratio averaged across each movement. The mean ratio was slightly elevated in the first few full force trials and then leveled off at a value of about 3.5. The heightened ratio in the initial trials was largely due to increases in grip force; the average load force remained quite constant across all full force trials. The half-life

yielded by an exponential fit to the mean force ratios was 3.71 trials. This corresponds closely with our estimate of the rate of learning of grip force prediction (above). The mean force ratios in the replication trials were clearly lower than in the initial trials being replicated and were, if anything, smaller than in the later full force trials. It is important to stress that the large reactive grip force increases observed in the first few trials are unlikely to be due simply to the larger overall grip force used in these trials. First, our measure of grip force increase, from the time of peak load to 100 ms later, was normalized to grip force at the time of peak load. Second, with greater overall grip force, smaller reactive grip force increases are observed [15].

These above results demonstrate that subjects learned to predict the behavior of the object with novel dynamics, so as to generate appropriate grip forces, before they learned to control the behavior of the object, so to as achieve the desired movement trajectory. It is important to emphasize several differences between learning to control the trajectory of the object and learning to modulate grip force appropriately. In our task, it is the desired trajectory of the object that is specified by the goal; namely, to move the object in a straight line from the start location to the target within a specified time. Although maintaining a stable grasp may also be viewed as a goal, the aim is simply to preserve an adequate ratio of grip force to load force. The desired grip force profile depends solely on the trajectory of the object and the dynamics of the object (which together determine the load force profile), and it is not directly specified by the task. Although people will slow down their movements if excessive grip force would otherwise be required [16], there is no evidence that the form of the hand trajectory is determined by constraints on grip force production. To the contrary, very similar hand trajectories are observed with and without objects in hand (e.g., [13, 17–19]). Thus, grip force responses may be viewed as postural adjustments that provide support and stabilization for the task at hand [20, 21].

Learning to control the object trajectory involves learning a new mapping between the desired trajectory (which does not change in our task) and the motor commands required to achieve this trajectory. Initially, the actual trajectory will be disturbed by the novel object dynamics such that there will be a discrepancy between the actual and desired trajectories. In contrast, the motor

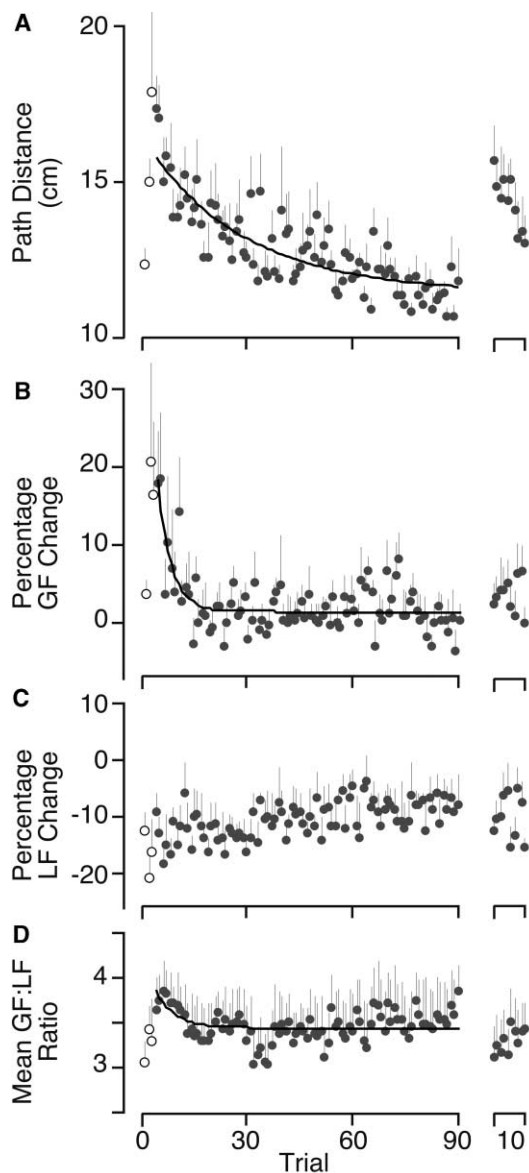


Figure 2. Measures of Grip Force Prediction and Trajectory Control (A) Hand path distance plotted as a function of trial. Each point represents the average of eight subjects, and the vertical lines represent 1 SE. The open circles represent the three warm-up trials in which the novel load acting on the object was incrementally increased. The ten points shown to the right represent the replication trials in which the subject had to reproduce the trajectories of the first ten trials with the full novel load. The solid curve represents an exponential fit to the mean values averaged across subjects. (B–D) Corresponding plots showing the percentage changes in (B) grip force and (C) load force from the time of peak load force to 100 ms later; (D) the mean ratio of grip force to load force within a trial.

system does not have to learn a new mapping between desired grip force and grip motor commands. This mapping depends only on the internal dynamics of the object (e.g., object compliance) that do not change in our task. That is, grip force is generated against a familiar, rigid object throughout the experiment. Thus, there will be no discrepancy between desired and actual grip force

profiles. However, knowledge of the (external) dynamics of the object is required to determine the desired grip force profile since the load force acting at the fingertips depends on these dynamics in combination with the trajectory of the object.

Given that the novel object dynamics will affect both the trajectory of the object and the desired grip force, it follows that learning the dynamics of the object is essential for both grip force and trajectory control. How, then, can we explain why the former is established far more rapidly than the latter? One possibility is that grip and trajectory learning involve the adaptation of separate inverse models, with one adapting more rapidly. Thus, one inverse model would map the desired object trajectory onto arm motor commands, whereas the other would map the desired object trajectory onto grip force motor commands. Note that the motor commands needed to control the arm depend on the dynamics of the object as well as the dynamics of the arm itself, whereas the motor commands required for grip force depend on the dynamics of the object as well as the internal dynamics of the object. Given that both the dynamics of the arm and the internal dynamics of the object are familiar, the only new learning required by each inverse model would be the dynamics of the object (i.e., the external or motion-dependent dynamics). Thus, it seems unlikely, a priori, that the two inverse models would be learned at different rates. Moreover, adaptation of the internal model for grip force should depend on adaptation of the inverse model for trajectory control because the mapping between the desired object trajectory and required grip force, which is determined by the actual trajectory, will change as the inverse model for trajectory control adapts and the discrepancy between the actual and desired object trajectories changes.

Another possible explanation for our results is that grip force learning involves the adaptation of a forward model that is distinct from the inverse model adapted for trajectory control [22, 23]. By combining a forward model of the object and arm with a copy of the arm motor command (efferent copy [24]), the load force acting at the fingertips (or, more precisely, the expected sensory consequences of the load force) could be predicted. The output of the forward model would then be sent to a grip force controller to generate grip motor commands appropriate for the expected load force. An important feature of this control scheme is that adaptation of the forward model does not require information about the desired object trajectory. Adaptation of the forward model is based on the error between actual sensory feedback and sensory feedback predicted from the arm motor commands, regardless of whether the latter achieve the desired trajectory or not. Thus, in theory, the forward model can be adapted independently of the inverse model.

Several motor control architectures have recently been proposed that have explicit representations of separate forward and inverse models for prediction and control [6–8]. Our results suggesting faster learning of the predictor over the controller are consistent with models that incorporate a predictor that is used to train the controller [6, 7, 25]. In fitting exponentials to the grip force and trajectory learning curves, we obtained a ratio

of half-lives of 2.64–19.8. Within the framework of combined forward and inverse models, this suggests that the forward model is adapted 7.5 times more quickly than the inverse model. A recent simulation of trajectory adaptation to novel dynamics [8] found that forward and inverse models are learned at similar rates when adapting to a novel force field, but that the forward model is learned five times more quickly when subsequently adapted to an opposing force field. However, the authors caution that these estimates “can only be taken as preliminary evidence, because the ability to estimate the rate of adaptation of the inverse model was hampered by the relative insensitivity of movement parameters to changes in this part of the adaptive controller.”

In summary, we have shown that, in a motor learning task involving the manipulation of an object with novel dynamics, subjects can learn to predict the behavior of the object before they master control over the behavior. The ability to quickly learn prediction enables us to stabilize our limbs and the load and may also play an important role in training the controller.

Experimental Procedures

Two six-axis cylindrical force transducers (Nano F/T, ATI Industrial Automation) were embedded in a custom-built cylindrical test object with two parallel vertical grip surfaces 3 cm in diameter and 6.4 cm apart. The two grip surfaces were covered with fine grade 300 sandpaper. Each force transducer measured translational forces in three dimensions (0.05 N resolution) at 500 Hz. The test object was attached to a light-weight robotic manipulator (Phantom Haptic Interface 1.5, Sensable Devices) that could generate forces up to 8 N in any of three dimensions. Three optical encoders, placed on the three motors of the robot, were used to measure the object's position (0.1 mm resolution) at 500 Hz. The three-dimensional force exerted by the manipulandum on the hand was servo controlled at 1000 Hz in order to simulate novel object dynamics. Specifically, we used the manipulandum to create an unusual force field; as subjects moved the object from right to left in the horizontal plane, as required in our task, an upward vertical force was generated proportional to the horizontal velocity of the object.

Eight right-handed subjects who were naive as to the purpose of this experiment gave their informed consent. A local ethics committee approved the experimental protocol. Using their right hand, subjects grasped the test object by using a precision grip with the tips of the index finger and thumb on the grip surfaces. The grip axis was orthogonal to the subjects' coronal plane, and the digits were oriented horizontally. The object was positioned approximately 30 cm in front of the subject at shoulder level. In each trial, the subject was required to move the object between the start position and the target positioned 10 cm to the left and in the same coronal plane. The locations of the start position, the target, and the current position of the object were continuously displayed on a computer monitor located in front of the subject at eye level.

To initiate a trial, the subject had to position the object within 0.5 cm of the start position, and the object had to be moving at less than 1 cm/s for 500 ms. The color of the target then changed from green to red, providing the signal to move. When the object had moved 1 cm from the start position, a timer was started, and a tone was delivered when the timer reached 500 ms. Subjects were instructed to reach the target coincident with this tone.

To familiarize subjects with the task, they initially performed 30 trials in which the force-servo was turned off and in which the object behaved as a simple mass load (65 g). Subjects then performed 90 trials in which the object had novel dynamics. For these trials, a horizontal line between the start and target positions was also displayed on the monitor, and subjects were explicitly instructed to move along the line. In the first four trials, the coefficient relating

horizontal velocity to upward force was increased from 6.25 N/m/s to 25 N/m/s in increments of 6.25 N/m/s. The coefficient then remained at 25 N/m/s for the remaining 86 trials. We used warm-up trials to prevent subjects from overgripping the test objects when first encountering the novel dynamics. After completing these 90 trials, subjects were asked to replicate, with the same object dynamics, the movement paths they produced in their first 10 full load trials with the novel object (trials 34 to 43). The path to be replicated was displayed on the monitor together with the current position of the object, and the same auditory timing cue was provided. Each of the ten traces was presented six times in a block fashion, and the block order was randomized. The replication trial that most closely matched the original trial in terms of velocity was selected as the replication trial. Specifically, we computed, for both the horizontal and vertical velocity, the root mean square errors between the target trial and each replication trial and selected the replication trial with the smallest summed error. By comparing grip forces in the first ten full load trials with the replication trials, we could examine the effect of learning grip force prediction independently of the form of the trajectory.

Raw position and force data were filtered offline with a 4th order, zero phase lag, 14 Hz low-pass Butterworth filter. The grip force was computed as the average of the normal forces at the two grip surfaces. To compute load force, we first determined, for each grip surface, the resultant of the two tangential forces, and we then summed these resultant forces. To assess the ability of subjects to control the object with novel dynamics, we simply measured the hand path distance. The start and end of the movement were defined at the times at which the resultant velocity of the hand first exceeded and last dropped below 2 cm/s, respectively. To assess subjects' ability to predict these dynamics, we measured the magnitude of the reactive grip force response. When manipulating objects with familiar dynamics [13, 26, 27], or novel dynamics that have been learned [19, 28], peak grip force coincides closely in time with peak load force. Such predictive modulation of grip force may be contrasted with reactive grip adjustments that typically occur 60–90 ms after unexpected or poorly predicted changes in load [12, 29–32]. As a consequence of these reactive grip responses, peak grip force lags behind the load force peak by some 100 ms [33–36].

To capture such reactive changes, we subtracted the grip force at the time of peak load force from the grip force measured 100 ms later and normalized this difference by dividing it by the grip force at the time of peak load force. For comparison, we computed the relative load force change by subtracting the peak load force from the load force 100 ms later and dividing by peak load force. We expressed these relative changes in force as percentages. We also computed, for each trial, the correlation between the rate of change of grip force and the rate of change of load force and the average grip force to load force ratio. Both of these measures were computed over the period of time from the start of movement to the end of movement.

Acknowledgments

This project was supported by grants from the Canadian Natural Sciences and Engineering Research Council, the Human Frontier Science Program, the Wellcome Trust, and the Swedish Medical Research Council. P.V. was funded by the Wellcome 4 year PhD Programme in Neuroscience at University College London. We thank Paul Bays and James Ingram for technical and programming assistance.

Received: September 2, 2002

Revised: November 7, 2002

Accepted: November 7, 2002

Published: January 21, 2003

References

1. Miall, R.C., and Wolpert, D.M. (1996). Forward models for physiological motor control. *Neural Networks* 9, 1265–1279.
2. Wolpert, D.M., and Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nat. Neurosci.* 3, 1212–1217.

3. Kawato, M., Furawaka, K., and Suzuki, R. (1987). A hierarchical neural network model for the control and learning of voluntary movements. *Biol. Cybern.* *56*, 1–17.
4. Jordan, M.I. (1996). Computational aspects of motor control and motor learning. In *Handbook of Perception and Action: Motor Skills*, Second Edition, H. Heuer and S. Keele, eds. (New York: Academic Press), pp. 71–118.
5. Desmurget, M., and Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends Cogn. Sci.* *4*, 423–431.
6. Jordan, M.I., and Rumelhart, D.E. (1992). Forward models: supervised learning with a distal teacher. *Cogn. Science* *16*, 307–354.
7. Wolpert, D.M., and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks* *11*, 1317–1329.
8. Bhushan, N., and Shadmehr, R. (1999). Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern.* *81*, 39–60.
9. Morasso, P. (1981). Spatial control of arm movements. *Exp. Brain Res.* *42*, 223–227.
10. Soechting, J.F., and Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *J. Neurosci.* *1*, 710–720.
11. Flash, T., and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* *5*, 1688–1703.
12. Johansson, R.S., and Westling, G. (1988). Coordinated isometric muscle commands adequately and erroneously programmed for the weight during lifting task with precision grip. *Exp. Brain Res.* *71*, 59–71.
13. Flanagan, J.R., and Wing, A.M. (1993). Modulation of grip force with load force during point-to-point arm movements. *Exp. Brain Res.* *95*, 131–143.
14. Flanagan, J.R., and Tresilian, J.R. (1994). Grip-load force coupling: a general control strategy for transporting objects. *J. Exp. Psychol. Hum. Percept. Perform.* *20*, 944–957.
15. Cole, K.J., and Johansson, R.S. (1993). Friction at the digit-object interface scales the sensorimotor transformation for grip responses to pulling loads. *Exp. Brain Res.* *95*, 523–532.
16. Saels, P., Thonnard, J.L., Detrembleur, C., and Smith, A.M. (1999). Impact of the surface slipperiness of grasped objects on their subsequent acceleration. *Neuropsychologia* *37*, 751–756.
17. Atkeson, C.G., and Hollerbach, J.M. (1989). Kinematic features of unrestrained vertical arm movements. *J. Neurosci.* *9*, 2318–2330.
18. Flanagan, J.R., and Lolley, S. (2001). The inertial anisotropy of the arm is accurately predicted during movement planning. *J. Neurosci.* *21*, 1361–1369.
19. Flanagan, J.R., and Wing, A.M. (1997). The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J. Neurosci.* *17*, 1519–1528.
20. Wing, A.M., Flanagan, J.R., and Richardson, J. (1997). Anticipatory postural adjustments in stance and grip. *Exp. Brain Res.* *116*, 122–130.
21. Winstein, C.J., Horak, F.B., and Fisher, B.E. (2000). Influence of central set on anticipatory and triggered grip-force adjustments. *Exp. Brain Res.* *130*, 298–308.
22. Flanagan, J.R., and Wing, A.M. (1996). Internal models of dynamics in motor learning and control. *Soc. Neurosci. Abstr.* *22*, 897.
23. Wing, A.M., and Flanagan, J.R. (1998). Anticipating dynamic loads in handling objects. *Proc. Am. Soc. Mechan. Engin. Dynamic Systems and Control Division* *64*, 139–143.
24. Von Holst, E. (1954). Relations between the central nervous system and the peripheral organs. *Brit. J. Anim. Behav.* *2*, 89–94.
25. Haruno, M., Wolpert, D.M., and Kawato, M. (2001). Mosaic model for sensorimotor learning and control. *Neural Comput.* *13*, 2201–2220.
26. Johansson, R.S., and Westling, G. (1984). Roles of glabrous skin receptors and sensorimotor memory in automatic-control of precision grip when lifting rougher or more slippery objects. *Exp. Brain Res.* *56*, 550–564.
27. Flanagan, J.R., and Wing, A.M. (1995). The stability of precision grip forces during cyclic arm movements with a hand-held load. *Exp. Brain Res.* *105*, 455–464.
28. Hermsdorfer, J., Marquardt, C., Philipp, J., Zierdt, A., Nowak, D., Glasauer, S., and Mai, N. (2000). Moving weightless objects. Grip force control during microgravity. *Exp. Brain Res.* *132*, 52–64.
29. Johansson, R.S., and Westling, G. (1987). Signals in tactile afferents from the fingers eliciting adaptive motor-responses during precision grip. *Exp. Brain Res.* *66*, 141–154.
30. Cole, K.J., and Abbs, J.H. (1988). Grip force adjustments evoked by load force perturbations of a grasped object. *J. Neurophysiol.* *60*, 1513–1522.
31. Johansson, R.S., Hager, C., and Riso, R. (1992). Somatosensory control of precision grip during unpredictable pulling loads. II. Changes in load force rate. *Exp. Brain Res.* *89*, 192–203.
32. Johansson, R.S., Riso, R., Hager, C., and Backstrom, L. (1992). Somatosensory control of precision grip during unpredictable pulling loads. I. Changes in load force amplitude. *Exp. Brain Res.* *89*, 181–191.
33. Johansson, R.S., and Westling, G. (1988). Programmed and triggered actions to rapid changes during precision grip. *Exp. Brain Res.* *71*, 72–86.
34. Blakemore, S.J., Goodbody, S.J., and Wolpert, D.M. (1998). Predicting the consequences of our own actions: the role of sensorimotor context estimation. *J. Neurosci.* *18*, 7511–7518.
35. Witney, A.G., Goodbody, S.J., and Wolpert, D.M. (1999). Predictive motor learning of temporal delays. *J. Neurophysiol.* *82*, 2039–2048.
36. Witney, A.G., Goodbody, S.J., and Wolpert, D.M. (2000). Learning and decay of prediction in object manipulation. *J. Neurophysiol.* *84*, 334–343.