P. Cisek, T. Drew & J.F. Kalaska (Eds.)
Progress in Brain Research, Vol. 165
ISSN 0079-6123

CHAPTER 23

Trial-by-trial motor adaptation: a window into elemental neural computation

Kurt A. Thoroughman*, Michael S. Fine and Jordan A. Taylor

Department of Biomedical Engineering, Washington University, 1 Brookings Dr., Saint Louis, MO 63130, USA

Abstract: How does the brain compute? To address this question, mathematical modelers, neurophysiologists, and psychophysicists have sought behaviors that provide evidence of specific neural computations. Human motor behavior consists of several such computations [Shadmehr, R., Wise, S.P. (2005). MIT Press: Cambridge, MA], such as the transformation of a sensory input to a motor output. The motor system is also capable of learning new transformations to produce novel outputs; humans have the remarkable ability to alter their motor output to adapt to changes in their own bodies and the environment [Wolpert, D.M., Ghahramani, Z. (2000). Nat. Neurosci., 3: 1212–1217]. These changes can be long term, through growth and changing body proportions, or short term, through changes in the external environment. Here we focus on trial-by-trial adaptation, the transformation of individually sensed movements into incremental updates of adaptive control. These investigations have the promise of revealing important basic principles of motor control and ultimately guiding a new understanding of the neuronal correlates of motor behaviors.

Keywords: motor learning; motor control; psychophysics; reaching; neural network; generalization

Introduction

Despite the centrality of motor learning to basic and clinical neuroscience, we know very little about the quantitative role neural systems play in human motor behavior (Flash and Sejnowski, 2001). Motor behavior, of the hand and arm in particular, consists of many mathematical problems that the central nervous system solves effortlessly. Many reaching tasks seem only positionally dependent; bringing a glass to one's mouth, shaking a friend's hand, and writing with a pen all require the identification and acquisition of arm and hand placements. But this oversimplified positional description belies the complexity of these movements, because desired positions must be related to the velocities, torques, and forces necessary to execute the movement (Atkeson, 1989). For every desired movement of the arm, the central nervous system somewhere needs to calculate the motor neuronal activity appropriate to generate the muscle forces necessary to actuate the movement. The dynamic equations necessary to move an unencumbered arm are very complex (Chan and Moran, 2006); adding the interaction forces of manipulated objects makes the calculations even more daunting.

Motor learning of external dynamics

In the mid 1990s two research groups (Brandeis and MIT) discovered that important properties of

---

*Corresponding author. Tel.: +1 314 935 9094;
Fax: +1 314 935 7448;
E-mail: thoroughman@biomed.wustl.edu

DOI: 10.1016/S0079-6123(06)65023-1

human motor behavior could be discovered by bypassing the complexities of natural arm movements and adding externally generated, novel dynamic demands. The Brandeis group (Lackner and Dizio, 1994) generated these novel forces by seating participants in the center of a rotating room. When the room rotated at constant speed and was dark, participants had no sensory perception that they were moving. When making outward reaches, however, they experienced novel Coriolis forces due to the interaction of the rotation and the arm movement. Lackner and Dizio discovered that even without visual feedback, participants could adapt to the Coriolis forces to move straight to the target. When the room stopped its rotation, participants generated an after-effect, making pointing errors opposing the direction of the initial error. These after-effects demonstrated that people could build a lasting adaptation to predictive force control over mere minutes of training in a novel dynamic environment.

The MIT group (Shadmehr and Mussa-Ivaldi, 1994) achieved similar results with a very different paradigm. Shadmehr and Mussa-Ivaldi trained participants to make reaching movements while holding a robotic arm (termed a manipulandum) that generated unusual velocity-dependent (viscous) forces. As with the Brandeis findings, Shadmehr and Mussa-Ivaldi observed that participants could learn to counter the novel forces, and generated an after-effect when the forces were removed. The manipulandum, however, offered the flexibility of removing the perturbation in single trials within blocks of viscous force generation. Later, Shadmehr and Brashers-Krug (1997) would use the growth of after-effect magnitude as a novel metric of force field learning and as evidence of interference across training paradigms.

Shadmehr and Mussa-Ivaldi initially trained participants to reach with their hands in front of their torsos. They then asked if and how participants transferred this motor memory when controlling the robot with their arms outstretched, lateral to their shoulders. Shadmehr and Mussa-Ivaldi discovered that participants did transfer learned dynamics to this new posture, as evidenced by better-than-naïve performance with the viscous forces on, and lingering after-effects with the forces off. The experimenters also altered the dependence of the forces such that the viscosity was related to either hand velocity or joint velocity; upon changing postures, experienced forces therefore remained consistent in either hand coordinates or joint coordinates. Participants performed better when the forces were consistent in joint coordinates, providing evidence that the motor memory was represented in terms of joint velocity.

The Brandeis and MIT experiments ushered in a new subdiscipline of exploring motor learning via the psychophysics of adapting to external dynamics. Other studies have considered the transfer of motor learning across postures (Shadmehr and Moussavi, 2000), speeds (Goodbody and Wolpert, 1998), movement directions (Gandolfo et al., 1996), tasks (Conditt et al., 1997), or from one hand to the other (Criscimagna-Hemminger et al., 2003). Studies also considered possible interference across learning multiple dynamic (force-generating) environments (Brashers-Krug et al., 1996; Caithness et al., 2004) or between learning visual and dynamic perturbations (Tong et al., 2002). All these studies provided behavioral evidence of motor memories built over one or several training sessions.

Trial-by-trial adaptation in dynamic environments

This body of work characterized the temporal and spatial properties of motor memories, but did not consider the influence of individual movements on motor adaptation. Motor adaptation was usually investigated with blocks of movements in which the dynamic perturbation was held fixed. This design enabled a measure of the rise time and asymptote of learning over many trials, as well as the transfer of that memory to different times, different movements, and different environments. Since most paradigms kept the movement environment fixed, they could not identify how errors in a single movement led to adaptation in the very next movement. Some experiments interspersed catch trials, single movements in which forces are unexpectedly removed, during training sets (Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Brashers-Krug, 1997), but did not examine how those catch trials altered adaptation. The consensus was that since the catch trials were rare, they only subtly affected adaptation over hundreds of movements and could therefore be ignored in their contribution to adaptation.

Two groups, at Northwestern and Johns Hopkins, provided novel analyses to begin the quantitative investigation of motor adaptation across individual trials. These new experiments focused on protocols and analyses that shifted the focus away from blocks of movements to the contribution of individual experiences. The block design of previous protocols provided stable environments from which to abstract relevant metrics. The trial-by-trial approaches of the Northwestern and Johns Hopkins groups altered the environmental dynamics within training, therefore preventing the settling of behavior into easily identifiable time constants and asymptotes. At first glance this design would seem to obscure rather than illuminate adaptive processes. The experimenter, however, controlled the trial-by-trial sequence of experienced forces. This sequence was later used as a template to identify subtle changes in control following each force presentation. The totality of the complete sequence of presentations and responses gave the trial-by-trial approach the same interrogative power as block designs, but with the additional vantage point of identifying the immediate transfer of sensory information into action.

The Northwestern group (Scheidt et al., 2001) trained subjects to reach in a single direction, directly away from their bodies, while experiencing forces whose strength was randomly drawn from a distribution in each trial. Scheidt et al. found that participants' performance in a trial, as quantified by maximal hand deviation, linearly depended on forces experienced in that movement, and also on forces and error experienced in the previous movement. The overall learning over a training session was found to consist simply of the summation of these increments of adaptation.

The Johns Hopkins group (Thoroughman and Shadmehr, 2000) did not draw forces from a distribution but rather reanalyzed previous data in which participants learned a viscous environment with occasional catch trials. New analyses shifted the focus to adaptation across single trials, enabled by the catch trial occurrences. Participants trained in eight directions of movement, so these analyses could newly investigate if and how incremental adaptation generalized across movement directions. The analysis of Thoroughman and Shadmehr revealed that participants generalized sensed error to improve subsequent control across many movement directions.

A neural network model mimicked this generalization if the force estimate relied upon neural units broadly tuned to movement direction and speed. This model exemplified a broader theory of neural computation, that the weighted linear combination of broadly tuned neurons allows for generalization and that learning may occur solely in changing the weights (Poggio, 1990). Further studies have indicated that this broadly tuned neural network model mimicked motor learning in several viscous environments (Donchin et al., 2003) as well as position-dependent environments (Hwang et al., 2003). The constancy of this broad trial-by-trial generalization supports the theory that fixed neuronal tuning simplifies learning and, in motor adaptation, provides a simple, consistent expectation of environmental complexity (Pouget and Snyder, 2000).

In the aforementioned experiments, people demonstrated trial-by-trial adaptation that depended co-linearly on two factors: a signal proportional to error, and a generalization of that error across the movement space. We have recently reported that both of these dependencies are not fixed, but rather change as a function of short-term experience; adaptation can be markedly disproportional and generalization can be narrow or broad depending on short-term environmental experience.

Experience-dependent flexibility of error generalization

Previous research has suggested that the extent of trial-by-trial generalization remains fixed during learning, which provides a stable adaptive platform through which people can process sensed signals into motor output. These experiments have generated forces that mimicked natural environments in their low spatial complexity, where complexity describes the change in force direction across movement space. However, if forces change rapidly across movement space and the human motor system generalizes broadly and remains fixed, then generalization itself can be maladaptive. To directly test this theory, we (Thoroughman and Taylor, 2005) trained human participants in dynamics with low, medium, and high complexity, as determined by the equations:

$$F = -15\sqrt{\dot{x}^2 + \dot{y}^2}\begin{bmatrix} -\sin(m\phi) \\ \cos(m\phi) \end{bmatrix} \qquad (1)$$
$$\phi = \arctan\left(\frac{\dot{y}}{\dot{x}}\right)$$

where $\dot{x}$, $\dot{y}$, and ϕ are the Cartesian components and direction of hand velocity. By increasing m, the spatial complexity of the function mapping velocity direction into force direction could be varied (Fig. 1A–C). On different days, the force direction was changed as fast ($m = 1$; Field One), twice as fast ($m = 2$; Field Two), or four times as fast ($m = 4$; Field Four) as the velocity direction. We found that people could learn all three environments over training. Participants reached the same degree of learning in Fields 1 and 2, while slightly less learning occurred in Field 4.

Surprisingly, participants quickly changed the way they adapted in each environment (Thoroughman and Taylor, 2005). We quantified the trial-by-trial generalization of error into adaptation (Fig. 1D) by using a state-space model. A vector parameter B determined the strength of adaptation dependent on the angular separation (θ) between the direction of the movement in which error is sensed and the direction of the next controlled movement. Our parameterization did not capture a significant trial-by-trial adaptation in Field Four; this may be because learning plateaued early in this environment leaving little trial-by-trial signal throughout the training sets. Trial-by-trial generalization in Field One exhibited the strongest and broadest generalization. The generalization function was always positive, such that an error sensed in one direction generated the same sign (positive or negative) of adaptation in all subsequent movement directions. For example, a rightward force experienced while reaching away from the body would generalize to an expectation of a rightward force, even while reaching toward the right or toward the body. In contrast, Field Two had a remarkably different shape; it was narrower above the X-axis and featured negative components in directions far away from sensed error. A rightward force experienced while reaching away from the body would not generalize to movements toward the right, but would generalize to an expectation of a leftward force for movements toward the body (as $B \approx 0$ for $\theta = 90^\circ$ and $B < 0$ for $\theta = 180^\circ$).

When the forces changed rapidly across movement space, people changed the way they learned the environment. This reshaping of motor adaptation suggests that people can rapidly change their internal mapping of the movement space that informs transformations of sense into action. Coupled with our previous and current computational models and with prominent theories linking neuronal tuning to generalization, we suggest that the functional tuning of motor space may be plastic, to induce either appropriately broad or appropriately narrow trial-by-trial generalization.

Trial-by-trial adaptation in response to pulsatile perturbations

While these studies quantified motor adaptation across single trials, the continuous nature of the perturbing force was ill suited to investigate how different time points within individual movements informed subsequent adaptation. By perturbing participants with brief pulses of force, we recently (Fine and Thoroughman, 2006) explored if and how the feedback experienced at the beginning, middle, or end of a single movement differentially affected the control of the next movement.

Pulses were applied in 20% of movements, either to the left or right, 2, 3, 4, 5, 6, or 7 cm into a 10 cm movement. Pulses were never experienced in two consecutive movements. Prepulse movements (movements immediately before a pulsed movement) were relatively straight, curved only a couple of millimeters. The pulse significantly perturbed these prepulse movements. Participants adapted in movements immediately after pulsed movements (termed postpulse movements). Participants adapted to the left after rightward pulses

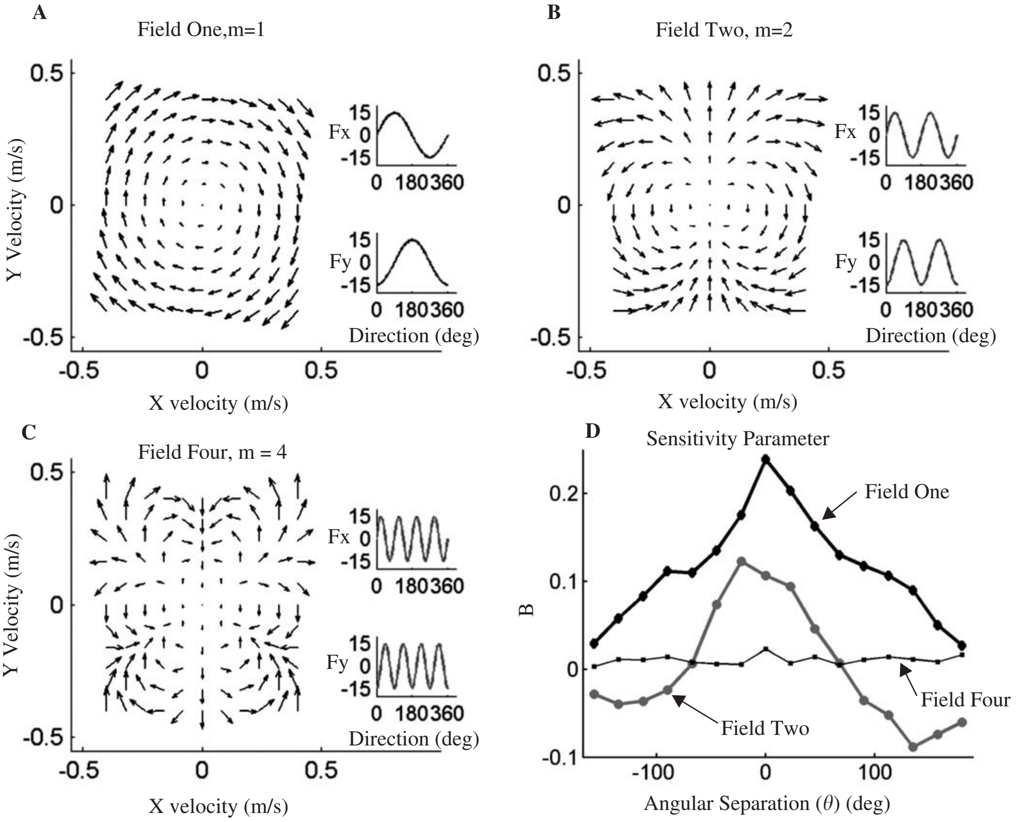

Fig. 1. (A–C) Trial-by-trial sensitivity was estimated in Fields One (A), Two (B), and Four (C). The main subfigures show the magnitude and direction of the force (arrows) as a function of the *X*- and *Y*-components of hand velocity. Insets show the dependence of the *X*- and *Y*-components of force on hand velocity direction. (D) The sensitivity function (*B*) plotted against the angular separation (*θ*) between sensed and adapted movement directions for each environment. Adapted with permission from Thoroughman and Taylor (2005), © 2005 Society for Neuroscience.

and to the right after leftward pulses. Surprisingly, the timing of postpulse adaptation did not depend significantly on the position of the pulse; pulses early, middle, and late in movement all altered subsequent control from the very beginning of the next movement. Postpulse trajectories also followed similar paths; adaptation to leftward and rightward pulses was statistically indistinguishable regardless of the position of the pulse, whether displacements were measured early, in the middle of, or late in the postpulse movement.

We also investigated the dependence of postpulse adaptation on the magnitude of the pulse. Arm position error (Jordan and Rumelhart, 1992), stiff and viscous feedback (Wolpert and Ghahramani, 2004), and predictive torque error (Jordan and Wolpert, 1999) have all been hypothesized to drive adaptation. All three of these hypotheses have, at their core, an assumption that adaptation strives to regress across experience to minimize the overall magnitude of error. Even reinforcement learning (Sutton and Barto, 1998), which lacks an explicit training signal, presumes regression to drive performance to an optimum. Each of these hypotheses, then, requires the magnitude of adaptation to vary proportionally with the magnitude of error. Current state-space models have included this presumption and have successfully reproduced human behavior when people experienced novel force perturbations throughout the extent of a movement.

We found that the human response to pulsatile force perturbations, in contrast, contained no proportionality to any error signal. We tested participants with force pulses of 70 ms duration and 6, 12, or 18 N peak force (Fig. 2A–C). Pulses were again experienced 3, 5, or 7 cm into a 10 cm movement, either to the left or the right. The error induced by the forces sensibly scaled with the magnitude of the perturbation. If the pulse pushed participants to the left, they adapted to the right in the very next movement; if the forces pushed participants to the right, they adapted by moving to the left. The size of the response, however, was constant, regardless of the magnitude of the force pulse (Fig. 2D–F).

Sensitivity to pulse direction, but insensitivity to pulse magnitude, is most apparent by averaging across pulse position, and plotting adaptation against the direction and magnitude of the force pulse (Fig. 3). This result was surprising, given the proportional responses measured in novel dynamic environments. This component of adaptation therefore cannot scale with any previously hypothesized error signal, nor can it depend on a real-valued critical metric. The magnitude of the force perturbations and positional errors were well within the ranges experienced in previous studies (e.g., Scheidt et al., 2001), so our categorical result is very unlikely to arise due to saturation effects. These results, when combined with previous studies, suggest that people can adopt different modes of adaptation that either can or cannot scale with error magnitude (Fine and Thoroughman, 2007). Both the transformation of error to adaptation and the generalization of error across movement space can therefore change with the environmental demands of a task.

Discussion

Monkey neurophysiologists have correlated the activity of neurons in specific brain areas to many different features of arm movement. Foundational work identified a strong correlation between the activity of primary motor neurons and the direction of arm movement (Georgopoulos et al., 1986). More recent work has investigated how cortical neurons predict both hand speed and movement direction (Moran and Schwartz, 1999) and may also encode several other movement parameters, such as posture of the arm (Scott and Kalaska, 1997) or muscle force (Li et al., 2001; Sergio et al., 2005). Activity of Purkinje cells in the cerebellum, meanwhile, can not only correlate to hand movement speed and direction (Coltz et al., 1999), but also to joint position and muscle force (Thach, 1978). The specific role of cerebral or cerebellar activity in normal motor behavior remains elusive largely due to the correlative nature of these experiments, the high number of neurons that participate in each movement, and the continuous time series of a large number of movement states (hand and joint position and velocity; joint torque; muscle force) that can all contribute to positive post-hoc correlations.

The goal of trial-by-trial analyses is to identify the particular computations used by participants to adapt across individual trials. These and further studies will determine specific signal processing computations that will provide neurophysiologists with protocols with which they can identify the specific contribution a particular brain area makes to the trial-by-trial transformation of sense into incremental adaptation.

Our two new results provide two examples of possible experimental correlates. The observed experience-dependent narrowing or broadening of trial-by-trial adaptation, when compared to theories of generalization, suggests that the neuronal activities underlying adaptation change their tuning with environmental demands. Neurophysiologists could train monkeys in similar environments first to confirm that non-human primates replicate the flexibility in trial-by-trial behavior. Neuronal recordings could then seek to identify areas of activity that change their tuning to movement direction when the environmental complexity changes.

Our observed categorical adaptation to force pulses, in turn, could be used to differentiate areas that encode error from areas that encode adaptation. The bevy of theories that posited a linear relation between sensed error and incremental adaptation could never, because of the linearity, dissociate error and adaptation encoding. If monkeys replicate our human psychophysical response to force pulses, mid-movement feedback

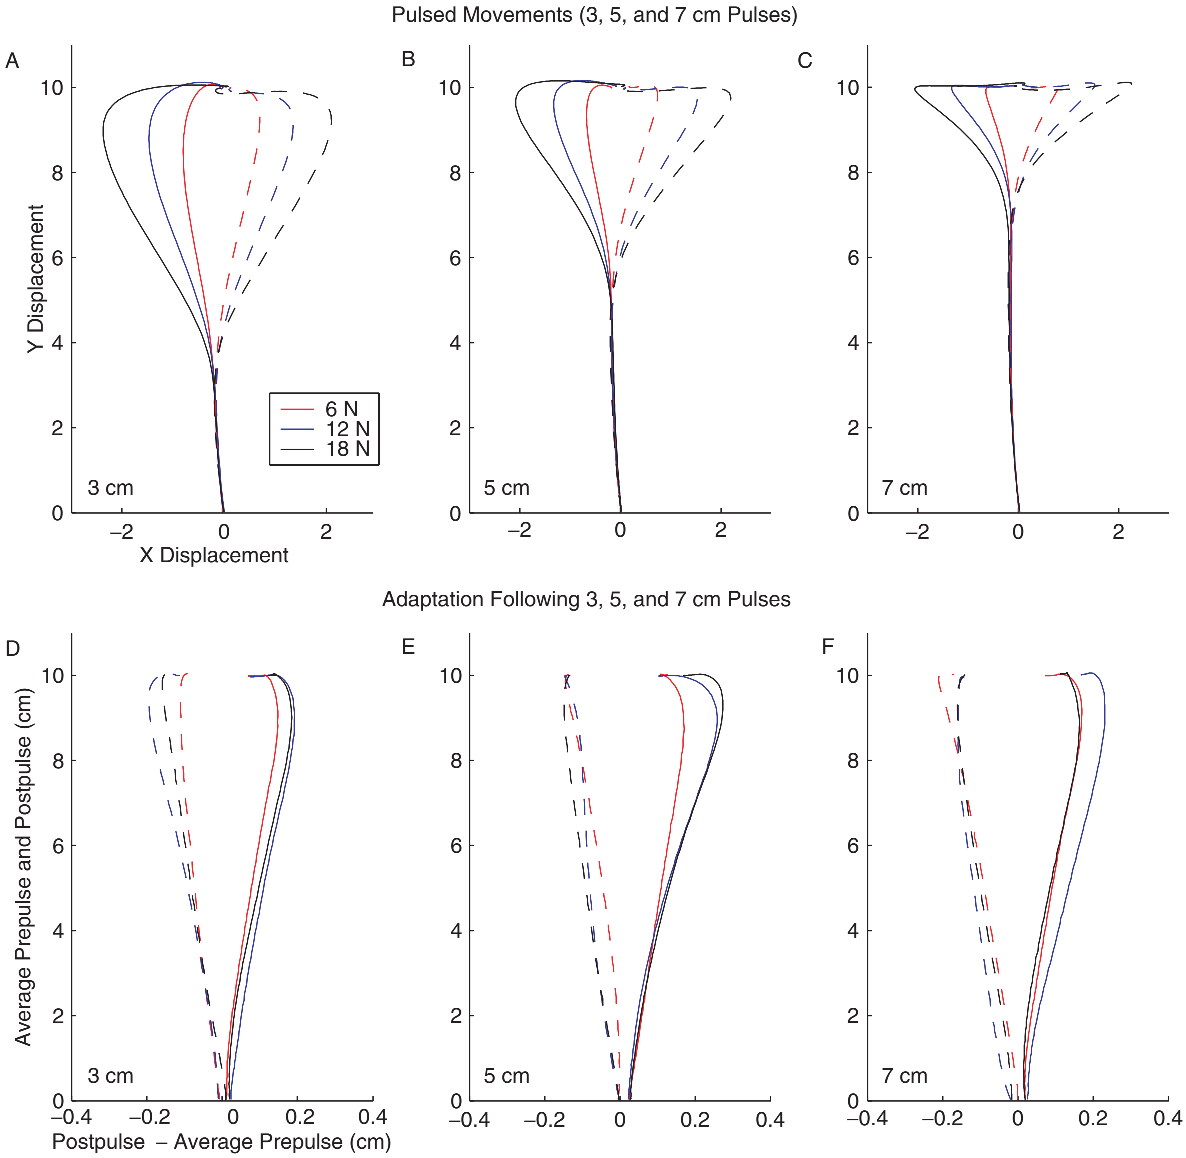

Fig. 2. (A–C) Movement trajectories, averaged across movements and participants, when pulses (C, inset) were applied 3, 5, or 7 cm into a 10 cm movement. The positional error and mid-movement correction were proportional to the magnitude of the pulse (indicated by color). (D–F) To quantify adaptation, we subtracted movements before a pulse (prepulse) from movements after a pulse (postpulse), and averaged across all participants within a particular pulsetype. Notice the *X*-axis is magnified. Adaptation counters the pulse direction but does not scale with the pulse magnitude. Adapted with permission from Fine and Thoroughman (2006), © 2006 American Physiological Society.

responses to the pulses will scale with pulse amplitude, but adaptation in the next movement will depend on the direction and not the magnitude of the pulse. Neuronal recordings could then determine whether a particular brain region's activity responded in proportion to pulse magnitude (as in Fig. 2A–C) or in response to direction but not magnitude (as in Figs. 2D–F and 3). Activity

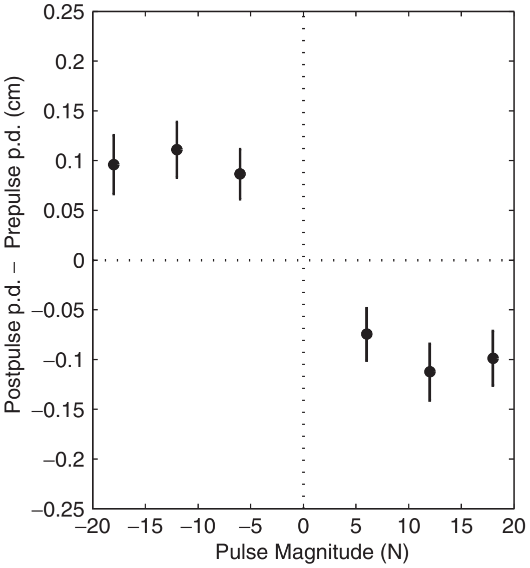

Fig. 3. The magnitude, quantified as the change in mid-movement perpendicular displacement, of aftereffects following force pulses of varying magnitude and direction. Data are averaged across pulse position and participants. Dots in the upper-left quadrant reflect the average aftereffect magnitude after leftward pulses; in the lower-right, after rightward pulses. If the average participant response was proportional, aftereffects following 18 N pulses would be three times the size of those following 6 N pulses. Adapted with permission from Fine and Thoroughman (2006), © 2006 American Physiological Society.

that scaled with pulse magnitude would indicate encoding of error; activity that responded to direction but not magnitude would indicate encoding of the incremental adaptation of subsequent control.

In particular, the trial-by-trial approach correlates well with established models of cerebellar motor adaptation, models that to date have not been tested with trial-by-trial analyses of arm movements. These models have their origins in the Marr notion of adaptation of parallel fiber synapses onto Purkinje cells (Marr, 1969). The foundational Marr hypothesis posits that synaptic efficacy changes as a function of two terms: an error signal, carried by the climbing fibers and encoded as a difference in complex firing rates from baseline, and the tuning curves of the parallel fiber input. The climbing fiber input forms a global error signal that might serve to change synaptic efficacy in many parallel fiber synapses; synaptic specificity comes via the dependency on the parallel input itself. Similar to a backpropagation framework (Jordan and Rumelhart, 1992), the synapse specificity enables the fine tuning of synapses in proportion to their contribution to the neuronal causation of the error. Recent neurophysiology has corroborated these theories; the complex spikes of the Purkinje cells seem to encode sensory error signals in motor command coordinates (Kawato et al., 1987; Kawato and Gomi, 1992). Neural correlates of this model have been identified in ocular following response (Shidara and Kawano, 1993; Gomi et al., 1998; Kobayashi et al., 1998) and in arm movements (Gilbert and Thach, 1977; Kitazawa et al., 1998).

The trial-by-trial approach enables a highly specific computational identification as it relies upon individual errors as inputs and resultant movement modifications as outputs. The careful analysis of current and emerging trial-by-trial adaptation protocols will make possible the construction of very precise hypotheses of neuronal processing on four levels: the tuning of the sensory representation, the representation of mid-movement corrective control, trial-by-trial adaptation, and the across-trial reshaping of that adaptation. The specificity of the model predictions, trial-by-trial monkey behavior, and the corresponding neural activity will foster a careful identification of the computations underlying sensorimotor adaptation and will enable a richly supported, biologically detailed integration of human psychophysics, computation, and primate neurophysiology.

References

Atkeson, C.G. (1989) Learning arm kinematics and dynamics. Annu. Rev. Neurosci., 12: 157–183.

Brashers-Krug, T., Shadmehr, R. and Bizzi, E. (1996) Consolidation in human motor memory. Nature, 382: 252–255.

Caithness, G., Osu, R., Bays, P., Chase, H., Klassen, J., Kawato, M., Wolpert, D.M. and Flanagan, J.R. (2004) Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. J. Neurosci., 24: 8662–8671.

Chan, S.S. and Moran, D.W. (2006) Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces. J. Neural Eng., 3: 327–337.

Coltz, J.D., Johnson, M.T. and Ebner, T.J. (1999) Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. J. Neurosci., 19: 1782–1803.

Conditt, M.A., Gandolfo, F. and Mussa-Ivaldi, F.A. (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. J. Neurophysiol., 78: 554–560.

Criscimagna-Hemminger, S.E., Donchin, O., Gazzaniga, M.S. and Shadmehr, R. (2003) Learned dynamics of reaching movements generalize from dominant to nondominant arm. J. Neurophysiol., 89: 168–176.

Donchin, O., Francis, J.T. and Shadmehr, R. (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. J. Neurosci., 23: 9032–9045.

Fine, M.S. and Thoroughman, K.A. (2006) Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. J. Neurophysiol., 96: 710–720.

Fine, M.S. and Thoroughman, K.A. (2007) The trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. J. Neurophysiol., (published online July 5, 2007). doi:10.1152/jn.00196.2007.

Flash, T. and Sejnowski, T.J. (2001) Computational approaches to motor control. Curr. Opin. Neurobiol., 11: 655–662.

Gandolfo, F., Mussa-Ivaldi, F.A. and Bizzi, E. (1996) Motor learning by field approximation. Proc. Natl. Acad. Sci. U.S.A., 93: 3843–3846.

Georgopoulos, A.P., Schwartz, A.B. and Kettner, R.E. (1986) Neuronal population coding of movement direction. Science, 233: 1416–1419.

Gilbert, P.F. and Thach, W.T. (1977) Purkinje cell activity during motor learning. Brain Res., 128: 309–328.

Gomi, H., Shidara, M., Takemura, A., Inoue, Y., Kawano, K. and Kawato, M. (1998) Temporal firing patterns of Purkinje cells in the cerebellar ventral paraflocculus during ocular following responses in monkeys I: simple spikes. J. Neurophysiol., 80: 818–831.

Goodbody, S.J. and Wolpert, D.M. (1998) Temporal and amplitude generalization in motor learning. J. Neurophysiol., 79: 1825–1838.

Hwang, E.J., Donchin, O., Smith, M.A. and Shadmehr, R. (2003) A gain-field encoding of limb position and velocity in the internal model of arm dynamics. PLoS Biol., 1: E25.

Jordan, M.I. and Rumelhart, D.E. (1992) Forward models: supervised learning with a distal teacher. Cogn. Sci. Multidisciplinary J., 17: 463–496.

Jordan, M.I. and Wolpert, D.M. (1999) Computational motor control. In: Gazzaniga M.S. (Ed.), The Cognitive Neurosciences. MIT Press, Cambridge, MA, pp. 601–620.

Kawato, M., Furukawa, K. and Suzuki, R. (1987) A hierarchical neural-network model for control and learning of voluntary movement. Biol. Cybern., 57: 169–185.

Kawato, M. and Gomi, H. (1992) The cerebellum and VOR/OKR learning models. Trends Neurosci., 15: 445–453.

Kitazawa, S., Kimura, T. and Yin, P.B. (1998) Cerebellar complex spikes encode both destinations and errors in arm movements. Nature, 392: 494–497.

Kobayashi, Y., Kawano, K., Takemura, A., Inoue, Y., Kitama, T., Gomi, H. and Kawato, M. (1998) Temporal firing patterns of Purkinje cells in the cerebellar ventral paraflocculus during ocular following responses in monkeys II: complex spikes. J. Neurophysiol., 80: 832–848.

Lackner, J.R. and Dizio, P. (1994) Rapid adaptation to Coriolis force perturbations of arm trajectory. J. Neurophysiol., 72: 299–313.

Li, C.S., Padoa-Schioppa, C. and Bizzi, E. (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. Neuron, 30: 593–607.

Marr, D.A. (1969) Theory of cerebellar cortex. J. Physiol., 202: 437–470.

Moran, D.W. and Schwartz, A.B. (1999) Motor cortical representation of speed and direction during reaching. J. Neurophysiol., 82: 2676–2692.

Poggio, T. (1990) A theory of how the brain might work. Cold Spring Harb. Symp. Quant. Biol., 55: 899–910.

Pouget, A. and Snyder, L.H. (2000) Computational approaches to sensorimotor transformations. Nat. Neurosci., 3(Suppl): 1192–1198.

Scheidt, R.A., Dingwell, J.B. and Mussa-Ivaldi, F.A. (2001) Learning to move amid uncertainty. J. Neurophysiol., 86: 971–985.

Scott, S.H. and Kalaska, J.F. (1997) Reaching movements with similar hand paths but different arm orientations I: activity of individual cells in motor cortex. J. Neurophysiol., 77: 826–852.

Sergio, L.E., Hamel-Paquet, C. and Kalaska, J.F. (2005) Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. J. Neurophysiol., 94: 2353–2378.

Shadmehr, R. and Brashers-Krug, T. (1997) Functional stages in the formation of human long-term motor memory. J. Neurosci., 17: 409–419.

Shadmehr, R. and Moussavi, Z.M. (2000) Spatial generalization from learning dynamics of reaching movements. J. Neurosci., 20: 7807–7815.

Shadmehr, R. and Mussa-Ivaldi, F.A. (1994) Adaptive representation of dynamics during learning of a motor task. J. Neurosci., 14: 3208–3224.

Shidara, M. and Kawano, K. (1993) Role of Purkinje cells in the ventral paraflocculus in short-latency ocular following responses. Exp. Brain Res., 93: 185–195.

Sutton, R.S. and Barto, A.G. (1998) Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA.

Thach, W.T. (1978) Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. J. Neurophysiol., 41: 654–676.

Thoroughman, K.A. and Shadmehr, R. (2000) Learning of action through adaptive combination of motor primitives. Nature, 407: 742–747.

Thoroughman, K.A. and Taylor, J.A. (2005) Rapid reshaping of human motor generalization. J. Neurosci., 25: 8948–8953.

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

Wolpert, D.M. and Ghahramani, Z. (2004) Computational motor control. In: Gazzaniga M.S. (Ed.), The Cognitive Neurosciences (3rd Ed.). MIT Press, Cambridge, MA, pp. 485–494.