

Computations in Sensorimotor Learning

DANIEL M. WOLPERT

*Computational and Biological Learning, Department of Engineering,
University of Cambridge, Cambridge CB2 1PZ, United Kingdom*

Correspondence: wolpert@eng.cam.ac.uk

Our cognitive abilities can only be expressed on the world through our actions. Here we review the computations underlying the way that the sensorimotor system converts both low-level sensory signals and high-level decisions into action, focusing on the behavioral evidence for the theoretical frameworks. We review recent work that determines how motor memories underlying sensorimotor learning are activated and protected from interference, the role of Bayesian decision theory in sensorimotor control including sources of suboptimality, the role of risk sensitivity in guiding action, and how rapid motor responses may underlie the robustness of the motor system to the vagaries of the world.

Although sensorimotor learning feels to us like a unitary process, its study is often broken down into a number of interacting processes. Here we review several aspects of sensorimotor learning that have seen progress in recent years and focus on computational approaches that have been studied behaviorally in humans. We start by examining how motor memories for different skills are organized, focusing on the rules by which different motor memories are parcellated and accessed. We then focus on the issue of online control and trajectory generation reviewing the theory of optimal feedback control and behavioral tests of this theory. Finally, we focus on the interaction between decision-making and sensorimotor control, examining their bidirectional interaction.

MOTOR MEMORIES

Our ability to learn and store multiple motor skills underlies the extraordinary flexibility with which we can revisit old skills, tools, and environments. A great deal of effort has been expended in trying to understand the rules by which novel skills are parcellated, stored, and protected from interference but yet reactivated and amended as a particular task is reexperienced. The use of motor tasks that are in some sense opposite to each other and therefore interfere has been a powerful tool to study motor memories within the laboratory setting (Brashers-Krug et al. 1996; Krakauer et al. 1999; Karniel and Mussa-Ivaldi 2002; Caithness et al. 2004). For example, robotic interfaces can be used to apply novel dynamics on the arm by having a subject grasp a handle that generates forces on the hand that typically depend on the state of the hand. Using such devices one can, for example, generate velocity-dependent forces that act at right angles to the direction of current motion—the so-called velocity-dependent curl field (Gandolfo et al. 1996). Such a field can be applied in either a clockwise (CW) or a counterclockwise (CCW) direction. When experienc-

ing a single force field, subjects learn to compensate for the dynamics, reverting to the approximately straight-line movements with bell-shaped velocity profiles before the perturbation. However, if CW and CCW fields are intermingled (either randomly or in batches) minimal learning is seen with such opposing force fields leading to strong interference. Even with contextual cues such as overlaying color or a virtual tool whose features fully predict the direction of the force field, learning is at best modest (Howard et al. 2013). This suggests that such opposing skills tend to compete for the same motor memory.

Contextual Rules for Motor Memories

Recently, it has been shown that more dynamic features of a task can allow separate motor memories to form. For example, it is possible to learn opposing force fields in one arm if motion of the other arm (present or absent) indicates the identity of the force field (Nozaki et al. 2006) or if the other arm makes concurrent and different movements for the two force fields (Howard et al. 2010). These results suggest that the coding of motor memories can depend on the conjunction of the movements of the two arms. Other features of movement can also be used to learn opposing perturbations. For example, if similar movements are made that are discrete or rhythmic, opposing perturbations can be learned when each is associated with a different type of movement (Howard et al. 2011). Separating the location of learning either proprioceptively or visually also facilitates learning of opposing force fields (Hwang et al. 2006; Hirashima and Nozaki 2012; Howard et al. 2013). Modeling approaches have suggested that such contextual effects arise from the engagement of separate neural populations (Nozaki and Scott 2009).

Recently, we have examined the effect of the lead-in to a movement. Although there was already evidence for some sequential effects on the ability to learn opposing

force fields (Wainscott et al. 2005), we examined a task in which we systematically varied the dwell time after a lead-in to a movement that would then be made in a force field (Howard et al. 2012). For example, the lead-in could be made from two different starting locations, followed by a single movement through either a CW or CCW force field. Critically the starting location was fully predictive of the force-field direction that would subsequently be experienced. When subjects were required to dwell for 1000 msec after the lead-in movement before proceeding to the movement through the field, they were unable to learn to compensate for the opposing force fields when randomly interleaved. However, as the dwell time was reduced to < 500 msec, learning increased, and with minimal dwell time there was substantial learning. This suggests that the recent history (~ 500 msec) of the state of the arm before experiencing the field affects the currently active motor memory and the more different the history is the greater the ability to separate motor memories. Interestingly, this effect did not require active movement, as it was seen even if the hand was passively moved for the lead-in.

This effect of recent states on separating motor memories led us to explore whether this effect could potentially explain the value of the follow-through in sports. Despite the fact that motion after contact or release with a ball cannot affect the flight of the ball, coaches have learned that the follow-through is important for skilled performance. We hypothesized that future actions that you plan to make might select current motor memories, and therefore different follow-throughs would preactivate distinct motor memories. We showed that it is possible to learn opposing force fields if the follow-through after the movement in the force field is different for each of the fields (Howard et al. 2015). This suggests that the value of making a consistent follow-through is that it activates a single motor memory leading to faster learning. We confirmed this by showing that when learning a single force field, learning is faster if a consistent follow-through is made after each movement through the field compared with when the follow-through is variable across trials. This gives a new interpretation to the value of the follow-through, that rather being specified so that you can optimize planning it may in fact optimize the learning. The neural underpinning of such follow-through effects is not known. However, potential future motor plans will certainly affect neural activity (Cisek and Kalaska 2005) so that different follow-throughs may directly engage separate neural populations that lead to the generation of movement. An alternative, but complementary perspective is that by viewing motor generation as a dynamical system of interacting neurons (Churchland et al. 2012) different follow-throughs will require different initial states of the dynamical system, each of which could be associated with the different force fields.

Structural Learning

Rather than using cues to allow distinct memories to be learned, we have shown that experiencing a variable perturbation increases our ability to learn a new perturbation

that comes from the same structure but not from other structures (Braun et al. 2009; Turnham et al. 2012). For example, having experienced visuomotor rotations drawn uniformly from -90° to $+90^\circ$, subjects can then readily learn a fixed $+60^\circ$ rotation and show minimal interference when they transition to a -60° rotation. We proposed that this ability depends on learning the structure, for example, of rotations. Two recent studies have suggested a more mechanistic explanation for the ability to learn from random experience. One suggests that such learning simply leads to increasing error sensitivity globally (Herzfeld et al. 2014), whereas the other suggests that we learn the appropriate weights for primitives (for the set of visuomotor errors we might experience) and that these are activated by prospective (predicted) errors (Takiyama et al. 2015). However, currently there are two issues that these models fail to address, which leaves the question of structural learning open. First, the models already assume the structure such as rotations in the underlying model so half the problem is effectively pre-solved. More troublesome is that the ability to generalize critically depends not on simply experiencing errors that correlate across trials (both models require this) but that the structure be consistent. We found no increased learning of visuomotor rotations when subjects experienced a range of linear visuomotor transformations that includes rotations, shearing, and scalings (Braun et al. 2009). Similarly, for groups of subjects exposed to random shearing or to random scalings, we found facilitation for shearing and scalings, respectively. Several other groups have also provided evidence that the structure of tasks is learned and used to appropriately make feedback and feed-forward adjustments (Yousif and Diedrichsen 2012) and that the degree to which tasks share similar structures determines the ability of people to generalize between them (Ranganathan et al. 2014).

Tool Use

Although learning of force fields and arbitrary visuomotor perturbations has led to a great deal of understanding, it may not reflect most of the learning we perform as adults. That is because although we need to learn the dynamics of our arms or of tools we interact with, similar to force-field learning, this may happen early in life. As an adult we might expect to have rather good structural models of how arms, hammers, and bicycles behave, and what we have to learn is the particular parameters of these such as length, mass, and friction. We have proposed that in naturalistic settings (Ingram and Wolpert 2011), such as when you operate on a tool such as a hammer, using vision you can estimate the structure of the task (assuming you have experience with hammers) but what you lack is precise knowledge of the parameters of the hammer (mass, density distributions, etc.). Consistent with this is that when you operate on a virtual hammer (e.g., rotating it at the handle), the forces that you generate even before any experience of the dynamics is in the correct direction to compensate for the forces induced on the handle by visual motion of the hammer head but

for an estimated magnitude (Ingram et al. 2010, 2013). Learning requires adaptation of the magnitude across trials. Such magnitude learning appears to be local to the orientation of the handle in the hand, suggesting that learning of parameters of tools, unlike the structure, is grasp-specific rather than holistic. Consistent with a qualitative difference between these situations is the different time course of learning for force fields compared with more naturalistic learning with tools, with the former an interaction of two processes (a fast and slow) (Smith et al. 2006) and the latter mediated by a single process (Ingram et al. 2011).

ONLINE CONTROL

Our actions are undoubtedly variable in that even if we wish to make the same movement repeatedly, we clearly have some variability that cannot be fully controlled. There is currently a debate as to the source of such variability, with different proponents suggesting that substantial variability can arise from sensory noise (Osborne et al. 2005), from suboptimal computation (Beck et al. 2012), or from motor noise (van Beers 2007; Faisal et al. 2008; Chaisanguanthum et al. 2014). However, even if only a proportion of variability arises from motor noise, it still places a premium on planning movement so as to reduce this component of variability, and current models of trajectory control place an emphasis on minimizing costs, which are a mixture of energy and accuracy (for review, see Scott 2012). Such models fall within the framework of optimal feedback control (OFC; Todorov and Jordan 2002; Diedrichsen et al. 2010), which suggests that people do not plan desired trajectories but instead generate time-dependent feedback controllers with the aim of minimizing the movement cost in the presence of potential disturbances. A number of studies have confirmed predictions of OFC in terms of the task-dependent corrections that it predicts (e.g., Diedrichsen 2007; White and Diedrichsen 2010). In addition, the framework of OFC has placed reflexes, or rapid motor responses, at the center of online control, as these reflect the feedback gains and OFC proposes that these should be highly tuned to the task at hand.

Rapid Motor Responses

Several studies have shown elegantly that many rapid motor responses are highly tuned to task demands (Kurtzer et al. 2008; Pruszynski et al. 2008; Yang et al. 2011; Pruszynski and Scott 2012). We have further explored task-dependent reflex tuning and found that both visuomotor and stretch reflexes are highly modifiable by task demands. For example, the visuomotor reflex can be both up-regulated and down-regulated when visuomotor perturbations are task-relevant or task-irrelevant, respectively (Franklin and Wolpert 2008). Moreover, even at the same point within a movement, the visuomotor reflex gain can be regulated to respond differently for leftward or rightward displacements relative to the movement direction (Franklin et al. 2014). Indeed such up-regulation

of visuomotor gains is also seen under the initial period of learning dynamic perturbations (Franklin et al. 2012). Such feedback gains can also be examined over the course of a movement and show a time-dependent modulation through the movement that is qualitatively in agreement with predictions of OFC (Dimitriou et al. 2013). Moreover, these gains can adapt within a movement to changes in task demands such as target jumps within 100 msec. Taken together, studies of reflexes within a limb show a level of sophistication that suggests a good match to task demands.

Recently reflexes between limbs have been studied and these show that stretch reflexes between the limbs are also modifiable by task demands (Dimitriou et al. 2012; Omrani et al. 2013). Consider holding a tray with both hands that needs to be maintained level. Vertical perturbations of one hand can be compensated by both reflex responses in the opposite direction to the perturbation in that arm or a response in the same direction in the other arm, both of these will restore the level of the tray. However, if both hands are perturbed vertically in the same direction then there is less need for reflex responses and conversely when they are perturbed in opposite direction reflex responses are needed in both arms. We examined just such a situation in which we could perturb one arm or both arms in the same or opposite direction and examined the rapid motor response (through electromyography [EMG] activity) and found clear task-dependent tuning of reflexes modulated online within a single trial based on the interaction across the arms and the nature of the perturbation (Dimitriou et al. 2012).

Risk Sensitivity

In general, optimal feedback control suggests that what we care about is the expected, or average, cost of the movement. Recently, we have examined the extent to which the variance of the cost is also important to subjects (for review, see Braun et al. 2011). Theories of risk sensitivity posit that the cost is a trade-off between the mean cost and variance, with those who are willing to accept a higher cost to either reduce variance or increase variance being risk-averse or risk-seeking. In a series of sensorimotor tasks we have shown that indeed subjects are often risk-averse or risk-seeking and do not simply optimize the expected reward (Nagengast et al. 2010, 2011a,b).

DECISION-MAKING AND SENSORIMOTOR CONTROL

Both decision-making and motor control require acting in real time on streams of noisy evidence. Thus both rely on inference, termination rules, time constraints, and value/effort costs. Up to now, these two fields have evolved in parallel. Decision-making has focused on complex streams of evidence from the environment but simple, ballistic movements. Conversely, motor control has focused on complex control problems but relatively simple stimuli. Both fields have computational frameworks

that underpin them. Decision-making has benefited from the framework of bounded evidence accumulation (e.g., drift-diffusion models) whereby a stream of noisy information is integrated to a threshold, which terminates the decision and establishes the choice outcome (for review, see Shadlen and Kiani 2013). This framework can explain both the timing and the accuracy of behavioral decisions and is supported by a wealth of physiology. Sensorimotor control has benefited from a mathematical framework of optimal feedback control, whereby costs associated with accuracy and effort are combined to produce a time-varying feedback controller. Bayesian decision theory has rapidly become a unifying framework in which to understand how sensorimotor processes deal with uncertainty to make inferences, make decisions, and generate actions.

Bayesian Decision Theory

Key to this framework is that humans learn priors (i.e., the statistical distribution) for tasks they need to perform. Many studies have shown that when a prior is experimentally placed upon a task subjects tend to learn the prior and use it in Bayesian way (e.g., Adams et al. 2004; Körding and Wolpert 2004). However, this is true only for simple priors such as Gaussians, and people have much more trouble learning complex priors, at least in a laboratory setting. To examine the sources of such suboptimality we examined whether it was the learning of the prior or the use of the prior that was limited. Subjects performed a Bayesian integration task in which the prior was explicitly presented and therefore did not need to be learned (Acerbi et al. 2014). Our results suggest suboptimality may be primarily due to a problem of acquiring the priors rather than computing with them.

Rather than applying priors it is possible to try to infer the subjects' natural priors, although such examination has until recently been limited to rather simple representations of the prior (Weiss et al. 2002; Turnham et al. 2011). We have recently developed a new method, termed cognitive tomography, which takes simple low-dimensional responses from subjects to calculate high-dimensional priors (Houlsby et al. 2013). As an example, we applied this method to the prior over the physical structure of faces by asking subjects to perform two different tasks. In one task they were presented with computer-generated pairs of faces and had to respond as to which face was more familiar, and in a second task they saw three faces and had to select which face they regarded as the odd one out. Given a generative model of these tasks we could infer each participant's subjective prior over the structure of faces (represented in the first two principal components of facial structure). We found that the priors were complex (not simply Gaussian) and varied widely across participants yet were very similar within each participant across the tasks. Moreover, we could use the prior extracted from one task to predict the performance on the other task. These results provide support for the ideas that there is a single high-dimensional prior for faces that guides behavior in multiple tasks.

In general, we tend to perceive our own actions differently to how we perceive actions of others, and most people overrate their ability to drive (most people rate themselves as above average). Recently we have examined whether the prior we have for our own actions are the same as priors for other agent's actions (Wolpe et al. 2014). We found that when performing an action ourselves we tend to have a prior for errors that is narrower than our true error distribution, leading to us perceiving our own performance as better than it really is. In contrast, when we watch another agent perform the same task (in fact our own movements played back to us) the prior now matches our error distribution leading to veridical perception. This suggests that in a simple sensorimotor task overoptimism may arise from our subjective prior being narrower than it should be.

Bidirectional Interplay of Decisions and Control

To examine the relation between high-level decision-making and sensorimotor control we used a task in which we could estimate both the evolving decision variable (from drift-diffusion models) and stretch reflex (from OFC) to relate these before a decision was made. The question we addressed is whether decision-making and motor control are a sequential process, accumulation to a bound followed by motor preparation and action, or whether there could be a more continuous flow from the decision process to the motor system. To do this, subjects were required to make a decision of the direction of motion of a random dot kinematogram and indicate the direction by an elbow flexion or extension movement when the display was extinguished (Selen et al. 2012). At the moment the display was extinguished a robotic device rapidly extended the elbow, setting up a stretch reflex whose magnitude could be measured using EMG. By varying the stimulus duration and difficulty (motion coherence) and fitting a model to the behavioral performance we could estimate the decision variable on each trial and relate this to the magnitude of the stretch reflex. We found that the stretch reflex tracked the evolving decision variable from even the smallest amounts of information, suggesting a continuous flow from the decision process (even before a decision has been made) to the motor system, allowing the motor system to pretune its reflexes so as to get to the appropriate target faster.

Such online interactions between the decision and motor systems can also be seen in a change-of-mind phenomenon. We examined subjects performing a simple motion direction discrimination task in which they had to indicate their decision (leftward versus rightward motion) by reaching to one of two choice targets (Resulaj et al. 2009). On initiation of the movement the stimulus was extinguished. Nevertheless, on a proportion of trials subjects appeared to change their mind, reaching initially for one target before deviating to the other. Such changes of mind improved performance even though the stimulus was extinguished. A consideration of the time delays in the sensory, motor, and decision process suggests that even though the stimulus is absent there is further infor-

mation on the processing pipeline (350 msec) that could be accumulated during the movement. We accounted for the pattern of the changes of mind with a simple model that assumed that after reaching a bound to make the initial decision, if a change-of-mind bound is then reached (on the opposite side of zero from the original bound but not requiring as much evidence as for the initial decision), then subjects will change their initial decision. Moreover, we and others have shown that the placement of this change-of-mind bound is dependent on the physical effort required to change one's mind (Burk et al. 2014; Moher and Song 2014). For example, more evidence is required to change one's mind if the choice targets are widely separated compared with when they are close together.

CONCLUSIONS

We have reviewed several aspects of sensorimotor control from motor memories underlying learning, online control, and the interaction of decision-making and control. Several themes have started to emerge within the field. First, the field of sensorimotor control is making links to others fields such as decision-making, working memory (e.g., Ma et al. 2014), and explicit and implicit processes (e.g., Taylor et al. 2014; Trewartha et al. 2014), and such links are likely to become increasingly important. Second, new theories are starting to emerge that can be tested on increasingly large sets of experimental data (e.g., Takiyama et al. 2015), and this is likely to provide an important framework for future theory development. Finally, as the behavioral theories develop, they are able to make increasingly important links to neurophysiological data (e.g., Azim et al. 2014; Zhou et al. 2014).

ACKNOWLEDGMENTS

This work was supported by Wellcome Trust, the Human Frontier Science Program, and a Royal Society Noreen Murray Professorship in Neurobiology to D.M.W.

REFERENCES

Acerbi L, Vijayakumar S, Wolpert DM. 2014. On the origins of suboptimality in human probabilistic inference. *PLoS Comput Biol* **10**: e1003661.

Adams WJ, Graf EW, Ernst MO. 2004. Experience can change the “light-from-above” prior. *Nat Neurosci* **7**: 1057–1058.

Azim E, Jiang J, Alstermark B, Jessell TM. 2014. Skilled reaching relies on a V2a propriospinal internal copy circuit. *Nature* **508**: 357–363.

Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A. 2012. Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron* **74**: 30–39.

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

Braun DA, Aertsen A, Wolpert DM, Mehring C. 2009. Motor task variation induces structural learning. *Curr Biol* **19**: 352–357.

Braun DA, Nagengast AJ, Wolpert DM. 2011. Risk-sensitivity in sensorimotor control. *Front Hum Neurosci* **5**: 1.

Burk D, Ingram JN, Franklin DW, Shadlen MN, Wolpert DM. 2014. Motor effort alters changes of mind in sensorimotor decision making. *PLoS ONE* **9**: e92681.

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

Chaisanguanthum KS, Shen HH, Sabes PN. 2014. Motor variability arises from a slow random walk in neural state. *J Neurosci* **34**: 12071–12080.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012. Neural population dynamics during reaching. *Nature* **487**: 51–56.

Cisek P, Kalaska JF. 2005. Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron* **45**: 801–814.

Diedrichsen J. 2007. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* **17**: 1675–1679.

Diedrichsen J, Shadmehr R, Ivry RB. 2010. The coordination of movement: Optimal feedback control and beyond. *Trends Cogn Sci* **14**: 31–39.

Dimitriou M, Franklin DW, Wolpert DM. 2012. Task-dependent coordination of rapid bimanual motor responses. *J Neurophysiol* **107**: 890–901.

Dimitriou M, Wolpert DM, Franklin DW. 2013. The temporal evolution of feedback gains rapidly update to task demands. *J Neurosci* **33**: 10898–10909.

Faisal AA, Selen LP, Wolpert DM. 2008. Noise in the nervous system. *Nat Rev Neurosci* **9**: 292–303.

Franklin DW, Wolpert DM. 2008. Specificity of reflex adaptation for task-relevant variability. *J Neurosci* **28**: 14165–14175.

Franklin S, Wolpert DM, Franklin DW. 2012. Visuomotor feedback gains upregulate during the learning of novel dynamics. *J Neurophysiol* **108**: 467–478.

Franklin DW, Franklin S, Wolpert DM. 2014. Fractionation of the visuomotor feedback response to directions of movement and perturbation. *J Neurophysiol* **112**: 2218–2233.

Gandolfo F, Mussa-Ivaldi FA, Bizzi E. 1996. Motor learning by field approximation. *Proc Natl Acad Sci* **93**: 3843–3846.

Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R. 2014. A memory of errors in sensorimotor learning. *Science* **345**: 1349–1353.

Hirashima M, Nozaki D. 2012. Distinct motor plans form and retrieve distinct motor memories for physically identical movements. *Curr Biol* **22**: 432–436.

Houlsby N, Huszár F, Ghassemi M, Orbán G, Wolpert D, Lengyel M. 2013. Cognitive tomography reveals complex, task-independent mental representations. *Curr Biol* **23**: 2169–2175.

Howard IS, Ingram JN, Wolpert DM. 2010. Context-dependent partitioning of motor learning in bimanual movements. *J Neurophysiol* **104**: 2082–2091.

Howard IS, Ingram JN, Wolpert DM. 2011. Separate representations of dynamics in rhythmic and discrete movements: Evidence from motor learning. *J Neurophysiol* **105**: 1722–1731.

Howard IS, Ingram JN, Franklin DW, Wolpert DM. 2012. Gone in 0.6 seconds: The encoding of motor memories depends on recent sensorimotor states. *J Neurosci* **32**: 12756–12768.

Howard IS, Wolpert DM, Franklin DW. 2013. The effect of contextual cues on the encoding of motor memories. *J Neurophysiol* **109**: 2632–2644.

Howard IS, Wolpert DM, Franklin DW. 2015. The value of the follow-through derives from motor learning depending on future actions. *Curr Biol* **25**: 397–401.

Hwang EJ, Smith MA, Shadmehr R. 2006. Dissociable effects of the implicit and explicit memory systems on learning control of reaching. *Exp Brain Res* **173**: 425–437.

Ingram JN, Wolpert DM. 2011. Naturalistic approaches to sensorimotor control. *Prog Brain Res* **191**: 3–29.

- Ingram JN, Howard IS, Flanagan JR, Wolpert DM. 2010. Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Curr Biol* **20**: 618–623.
- Ingram JN, Howard IS, Flanagan JR, Wolpert DM. 2011. A single-rate context-dependent learning process underlies rapid adaptation to familiar object dynamics. *PLoS Comput Biol* **7**: e1002196.
- Ingram JN, Flanagan JR, Wolpert DM. 2013. Context-dependent decay of motor memories during skill acquisition. *Curr Biol* **23**: 1107–1112.
- Karniel A, Mussa-Ivaldi FA. 2002. Does the motor control system use multiple models and context switching to cope with a variable environment? *Exp Brain Res* **143**: 520–524.
- Körding KP, Wolpert DM. 2004. Bayesian integration in sensorimotor learning. *Nature* **427**: 244–247.
- Krakauer JW, Ghilardi MF, Ghez C. 1999. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* **2**: 1026–1031.
- Kurtzer IL, Pruszynski JA, Scott SH. 2008. Long-latency reflexes of the human arm reflect an internal model of limb dynamics. *Curr Biol* **18**: 449–453.
- Ma WJ, Husain M, Bays PM. 2014. Changing concepts of working memory. *Nat Neurosci* **17**: 347–356.
- Moher J, Song JH. 2014. Perceptual decision processes flexibly adapt to avoid change-of-mind motor costs. *J Vis* **14**: 1.
- Nagengast AJ, Braun DA, Wolpert DM. 2010. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. *PLoS Comput Biol* **6**: e1000857.
- Nagengast AJ, Braun DA, Wolpert DM. 2011a. Risk-sensitivity and the mean-variance trade-off: Decision making in sensorimotor control. *Proc Biol Sci* **278**: 2325–2332.
- Nagengast AJ, Braun DA, Wolpert DM. 2011b. Risk sensitivity in a motor task with speed-accuracy trade-off. *J Neurophysiol* **105**: 2668–2674.
- Nozaki D, Scott SH. 2009. Multi-compartment model can explain partial transfer of learning within the same limb between unimanual and bimanual reaching. *Exp Brain Res* **194**: 451–463.
- Nozaki D, Kurtzer I, Scott SH. 2006. Limited transfer of learning between unimanual and bimanual skills within the same limb. *Nat Neurosci* **9**: 1364–1366.
- Omrani M, Diedrichsen J, Scott SH. 2013. Rapid feedback corrections during a bimanual postural task. *J Neurophysiol* **109**: 147–161.
- Osborne LC, Lisberger SG, Bialek W. 2005. A sensory source for motor variation. *Nature* **437**: 412–416.
- Pruszynski JA, Scott SH. 2012. Optimal feedback control and the long-latency stretch response. *Exp Brain Res* **218**: 341–359.
- Pruszynski JA, Kurtzer I, Scott SH. 2008. Rapid motor responses are appropriately tuned to the metrics of a visuospatial task. *J Neurophysiol* **100**: 224–238.
- Ranganathan R, Wieser J, Mosier KM, Mussa-Ivaldi FA, Scheidt RA. 2014. Learning redundant motor tasks with and without overlapping dimensions: Facilitation and interference effects. *J Neurosci* **34**: 8289–8299.
- Resulaj A, Kiani R, Wolpert DM, Shadlen MN. 2009. Changes of mind in decision-making. *Nature* **461**: 263–266.
- Scott SH. 2012. The computational and neural basis of voluntary motor control and planning. *Trends Cogn Sci* **16**: 541–549.
- Selen LP, Shadlen MN, Wolpert DM. 2012. Deliberation in the motor system: Reflex gains track evolving evidence leading to a decision. *J Neurosci* **32**: 2276–2286.
- Shadlen MN, Kiani R. 2013. Decision making as a window on cognition. *Neuron* **80**: 791–806.
- Smith MA, Ghazizadeh A, Shadmehr R. 2006. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* **4**: e179.
- Takiyama K, Hirashima M, Nozaki D. 2015. Prospective errors determine motor learning. *Nat Commun* **6**: 5925.
- Taylor JA, Krakauer JW, Ivry RB. 2014. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *J Neurosci* **34**: 3023–3032.
- Todorov E, Jordan MI. 2002. Optimal feedback control as a theory of motor coordination. *Nat Neurosci* **5**: 1226–1235.
- Trewartha KM, Garcia A, Wolpert DM, Flanagan JR. 2014. Fast but fleeting: Adaptive motor learning processes associated with aging and cognitive decline. *J Neurosci* **34**: 13411–13421.
- Turnham EJ, Braun DA, Wolpert DM. 2011. Inferring visuomotor priors for sensorimotor learning. *PLoS Comput Biol* **7**: e1001112.
- Turnham EJ, Braun DA, Wolpert DM. 2012. Facilitation of learning induced by both random and gradual visuomotor task variation. *J Neurophysiol* **107**: 1111–1122.
- van Beers RJ. 2007. The sources of variability in saccadic eye movements. *J Neurosci* **27**: 8757–8770.
- Wainscott SK, Donchin O, Shadmehr R. 2005. Internal models and contextual cues: Encoding serial order and direction of movement. *J Neurophysiol* **93**: 786–800.
- Weiss Y, Simoncelli EP, Adelson EH. 2002. Motion illusions as optimal percepts. *Nat Neurosci* **5**: 598–604.
- White O, Diedrichsen J. 2010. Responsibility assignment in redundant systems. *Curr Biol* **20**: 1290–1295.
- Wolpe N, Wolpert DM, Rowe JB. 2014. Seeing what you want to see: Priors for one's own actions represent exaggerated expectations of success. *Front Behav Neurosci* **8**: 232.
- Yang L, Michaels JA, Pruszynski JA, Scott SH. 2011. Rapid motor responses quickly integrate visuospatial task constraints. *Exp Brain Res* **211**: 231–242.
- Yousif N, Diedrichsen J. 2012. Structural learning in feed-forward and feedback control. *J Neurophysiol* **108**: 2373–2382.
- Zhou K, Wolpert DM, De Zeeuw CI. 2014. Motor systems: Reaching out and grasping the molecular tools. *Curr Biol* **24**: R269–R271.