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

CHAPTER 24

Towards a computational neuropsychology of action

John W. Krakauer[1] and Reza Shadmehr[2,*]

[1]*The Motor Performance Laboratory, Department of Neurology, Columbia University College of Physicians and Surgeons, New York, NY 10032, USA*
[2]*Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, MD 21205, USA*

Abstract: From a computational perspective, the act of using a tool and making a movement involves solving three kinds of problems: we need to learn the costs that are associated with our actions as well as the rewards that we may experience at various sensory states. We need to learn how our motor commands produce changes in things that we can sense. Finally, we must learn how to actually produce the motor commands that are needed so that we minimize the costs and maximize the rewards. The various computational problems appear to require different kinds of error signals that guide their learning, and might rely on different kinds of contextual cues that allow their recall. Indeed, there may be different neural structures that compute these functions. Here we use this computational framework to review the motor control capabilities of two important patients who have been studied extensively from the neuropsychological perspective: HM, who suffered from severe amnesia; and BG, who suffered from apraxia. When viewed from a computational perspective, the capabilities and deficits of these patients provide insights into the neural basis of our ability to willfully move our limbs and interact with the objects around us.

Keywords: motor control; computational models; forward models; optimal control

Introduction

In 1997, one of us performed a 2-day experiment on the severely amnesic patient HM to see how well he could learn a new motor task and retain it in memory (Shadmehr et al., 1998). The task was the standard reach adaptation task where subjects hold the handle of a robotic arm and learn to use it to guide a cursor to a sequence of targets (Shadmehr and Mussa-Ivaldi, 1994). When seated in front of the robot, HM, like all naïve volunteers, sat quietly and waited for instructions. We asked him to put his hand on the robot's handle and move it around a bit. Naturally, he kept his gaze on his hand as he moved the robot's handle. He was instructed to not look at his hand, but rather at the video monitor, where a cursor was present. After a minute or so of moving the cursor around, a center target was presented and he was asked to move the cursor to that location. Subsequently another target was shown and he was encouraged to move the cursor there. After a few minutes of practice in reaching to targets, the robot began to impose forces on HM's hand, perturbing the path of the cursor. With more practice, he altered his motor commands so to predicatively compensate for the forces, as evidenced by large errors in catch trials during which the field was turned off. We then thanked him for his time and he left to have lunch.

*Corresponding author. Tel.: +1 410-614-2458;
Fax: +1 410-502-2826; E-mail: reza@bme.jhu.edu

DOI: 10.1016/S0079-6123(06)65024-3

When he came back to the experiment room 4 h later, he claimed that he had never seen the robotic device or knew what it was for. We pushed the robotic arm aside and asked him to sit down. He sat down, but then voluntarily reached and grabbed the robot's handle, brought it toward him, and looked at the video monitor, where the cursor was present. It was clear that despite having no conscious recollection of having done the task before, some part of HM's brain recognized that the contraption was a tool that had a particular purpose: to manipulate cursors on a screen. Furthermore, his brain knew that in order to operate this tool, he had to hold its handle. When a target was presented, he showed strong after-effects of the previous training. That is, his brain expected the robot to perturb his movements, and so he generated motor commands that attempted to compensate for these forces. His brain knew all this, yet he was consciously unaware of this knowledge.

This experiment suggested to us that the brain could solve two general problems without conscious awareness and without the medial temporal lobes. First, the brain could learn how to use a tool in order to achieve a purpose. Second, the visual sight of the tool at a later time was sufficient to allow recall of both the purpose of the tool and the motor commands needed to achieve that purpose, although the latter may have required kinesthetic cues from the handle. This is despite the fact that the same visual information was not sufficient to recall conscious memory of having done the task.

Brenda Milner had of course made a similar observation in HM some 30 years earlier in a task where he drew on a piece of paper while looking in a mirror (Milner, 1968). In the novel visual feedback setting, HM adapted his motor output and learned to draw accurately. When he returned the next day, the visual and/or tactile cues associated with the experimental setup were sufficient to allow him to recall the motor skill that he had learned before. Over the years, a number of other investigators made similar observations in other amnesic patients (Gabrieli et al., 1993; Yamashita, 1993; Tranel et al., 1994), culminating in the theory that formation of motor memories are independent of the medial temporal lobe (Mishkin et al., 1984).

Subsequent research has made great strides in understanding the anatomical, physiological, and behavioral characteristics of motor memory. For example, many studies experimentally characterized and computationally modeled adaptation and re-adaptation of reaching movements in force-fields (Shadmehr and Brashers-Krug, 1997; Hwang et al., 2003; Krouchev and Kalaska, 2003; Caithness et al., 2004; Hwang and Shadmehr, 2005; Wainscott et al., 2005) and altered visuomotor environments (Krakauer et al., 1999, 2000, 2005; Wigmore et al., 2002; Caithness et al., 2004). In addition, single unit recording (Li et al., 2001; Padoa-Schioppa et al., 2002), lesion analysis (Maschke et al., 2004; Smith and Shadmehr, 2005), functional imaging (Krakauer et al., 2004; Diedrichsen et al., 2005), and some computational models (Donchin et al., 2003) successfully mapped adaptation processes onto the anatomy and physiology of the cerebral cortex, the basal ganglia, and the cerebellum. However, it appears to us that some of the fundamental implications of the HM experiments have been largely overlooked. For example, what cued him to recall the motor memory that he had acquired in the initial training session with the novel tool? In the original mirror drawing experiments, was it the visual cues associated with seeing the pencil and the mirror, or the act of holding the pencil? In the robot reaching task, was it the sight of the robot, or was it the act of holding the handle? What motivated him to learn these tasks and how did he sustain that motivation in subsequent sessions? Did he need to have a conscious desire to reduce errors for some perceived reward or did the motor system have its own implicit reward system?

Our aim here is to first pose the problem of motor control in a broad computational framework and then show that this framework sheds light on the neuropsychological basis of motor learning and motor memory.

Seeking rewards and observing the consequences of action

Our decisions to perform a task are guided by measures of costs and rewards. For example, in a

task when volunteers are asked to control a robotic handle in order to move a cursor to a target, they are often told to try to get to the target in a given amount of time. If they do so, they are "rewarded" by a target explosion, whereupon they hear the experimenter make an encouraging comment. They are also often paid for their time, and this payment may depend on performance.

Was HM's behavior driven by reward? For HM, the target explosion triggered a childhood memory of going bird hunting. As he was performing the task and was able to get a target explosion, he would spend the next few minutes describing the memory in detail: the type of gun that he used, the porch in the rear of his childhood home, the terrain of the woods in his backyard, and the kinds of birds that he hunted. He repeated these details to us many times during the 2-day experiment. Sue Corkin (our colleague who had examined HM for many decades) mentioned that she had not heard HM talk about this before. Although we did not record the conversation, the re-telling of the memory appeared to be related to the target explosions.

When HM came back on his second and third sessions with the robot, he did not need instructions on what the task was about; he sat down, reached for the robot, looked at the video screen and waited for the targets to appear. Perhaps the earlier experience with the task was such that it left a memory of expected reward: upon return, the sight and touch of the robot was sufficient to encourage a motor act that would be expected to be rewarding. If on the other hand, use of the robot in the first session had been paired with a shock or another noxious stimulus, it seems likely that he would have been reluctant to use the device again.

In a broader sense, experience with a tool provides us with information about how its use is associated with costs and rewards. In the robot reaching task, the reward is achieved only if the target explodes, and the cost is the energy spent doing the reaching. (For a typical volunteer, there are also costs associated with time away from their normal routine, but for simplicity let us ignore these realities.) Assuming that cursor position is $\mathbf{y}$ and target position is $\mathbf{r}$, then through experience we learn that the objective of the task is to minimize $(\mathbf{y}-\mathbf{r})^T(\mathbf{y}-\mathbf{r})$ at time N after the reach starts (this is the time that the target will explode if we are near it). Superscript T is the transpose operator. To denote the fact that this cost is zero except for time N, we write this cost as:

$$\sum_{n=1}^{N}(\mathbf{y}^{(n)}-\mathbf{r})^T T^{(n)}(\mathbf{y}^{(n)}-\mathbf{r})$$

where the superscript (n) refers to a discrete measure of time, and matrix T is a measure of our cost at each time step (which may be zero except at time N). We also have a cost associated with our motor commands $\mathbf{u}$, which here we assume grows as a quadratic function. Now the total cost becomes:

$$\sum_{n=1}^{N}(\mathbf{y}^{(n)}-\mathbf{r})^T T^{(n)}(\mathbf{y}^{(n)}-\mathbf{r})+\mathbf{u}^{(n)T}L^{(n)}\mathbf{u}^{n} \quad (1)$$

where matrix L is a time-dependent measure of the weighted costs associated with the motor commands. To get reward (explode the target), we need to find the motor commands that will minimize this cost. We learn this by observing that moving the robot handle will move the cursor. In particular, we learn that pushing on the robot will result in a specific proprioceptive and visual feedback regarding the state of our body and the state of the cursor. These are the sensory consequences of our actions. Grossly simplifying the problem, here we write these consequences as a linear function of our motor commands:

$$\begin{aligned}\mathbf{x}^{(n+1)} &= A\mathbf{x}^{(n)}+B\mathbf{u}^{(n)}+\varepsilon_u^{(n)}\\ \mathbf{y}^{(n)} &= C\mathbf{x}^{(n)}+\varepsilon_y^{(n)}\end{aligned} \quad (2)$$

where $\mathbf{x}^{(n)}$represents the state of the body and the world that we interact with, $\mathbf{u}^{(n)}$ is the motor command, $\varepsilon_u^{(n)}$ a stochastic variable representing motor noise, and $\varepsilon_y^{(n)}$ a stochastic variable representing sensory noise. If Eq. (2) is an accurate model of how motor commands to our arm muscles produce changes in the state of our body and the cursor, then we can use it as a set of constraints with which to minimize Eq. (1). This is the classic linear quadratic problem in optimal control. Solving this problem under the assumption that the noise variables are Gaussian yields a linear feedback control law that specifies the state-dependent motor commands that we should produce at each

time step:

$$\mathbf{u}^{(n)} = (L^{(n)} + B^T W^{(n+1)} B)^{-1} B^T (\boldsymbol{\theta}^{(n+1)} - W^{(n+1)} A \mathbf{x}^{(n)}) \quad (3)$$

The two new variables in this equation, W and $\mathbf{q}$, are time dependent quantities that depend on Lagrange multipliers that reflect the constraints in Eq. (2) in minimizing Eq. (1).

To summarize, the computational problem of learning motor control may be described as having three components:

1. To perform any action, we need to know the costs that are associated with our actions as well as the sensory states that are rewarding [Eq. (1)]. In the reaching task, through instruction or observation we might learn that target explosions are a rewarding act and they occur only when the cursor reaches the target at a specific time. The relative benefit of this reward with respect to the cost of the motor commands will dictate an internal value of this act.
2. We need to know how our motor commands produce changes in things that we can observe [Eq. (2)]. That is, through experience we must learn that when we are holding the robot in hand, our motor commands will result in a specific change in the state of our arm and the state of the cursor. This is a system identification problem associated with the particular tool. Learning of this map is called forming a forward model.
3. Finally, we must learn how to actually produce the motor commands that are needed so that we minimize the costs and maximize the reward [Eq. (3)]. That is, we need to figure out the "best" motor commands that bring the cursor to the target and get it to explode. This is the constrained minimization problem [minimize Eq. (1) under the constraints of Eq. (2)]. The result of the minimization is a feedback control law that specifies the motor response to the sensory states that we observe in our body and the environment. Learning of this feedback control law is called forming an inverse model.

This computational framework for representing the problem of biological motor control is largely due to the pioneering work of Todorov and Jordan (2002). At the heart of the approach is a departure from a framework in adaptive control where agents are provided with desired trajectories. Indeed, here we have no one to tell us what actions to follow. Rather, the goal is to acquire rewards, and the means is through observing the consequences of our actions.

What motivated HM to learn the reach adaptation task?

We expected that a severely amnesic individual who was performing a novel task would have to be regularly reminded of the task's instructions: "try to move the cursor to the target fast enough so it explodes." However, after HM had exploded a few targets, he no longer needed verbal reminders. Strikingly, when he returned the next day he voluntarily reached for the robot handle and began preparing for onset of targets by moving the cursor to the center location. This remarkable behavior suggests that during the first session, he learned the reward basis of the task [Eq. (1)] implicitly, while during the later sessions, the visual appearance of the machine, and the act of holding its handle, was sufficient to trigger a recall of this reward structure.

Why should target explosions be an implicitly rewarding action for HM? Perhaps because he associated them with his earlier experience with bird hunting, a memory that he very much enjoyed retelling. It is entirely possible that if we had chosen another mode of task performance feedback, say a numeric score, the intrinsic value of the task for HM might have been much lower. This might result in a subject who is ambivalent or even reluctant to perform a task.

We would conjecture that the reward mechanism that is a part of Eq. (1) is implemented in the brain through the release of dopamine. We have recently shown (Mazzoni et al., 2007) that although patients with Parkinson's disease (PD) reach more slowly than age-matched controls, they are nevertheless able to make fast movements without a loss in accuracy if forced to move fast by the experimental task. However, the patients take longer (require more trials) to accumulate a set number of movements at the required speed. Thus,

bradykinesia (slowness of movements), one of the hallmarks of PD, may represent decreased intrinsic value of a given motor task because of an imbalance between an estimate of an implicitly-determined cost for making a fast movement and the expected rewards.

Learning sensory consequences of motor commands vs. learning optimum control policies

In order for HM to explode targets, he had to learn how motor commands to his arm produced changes in the proprioceptive state of his limb and the visual state of the cursor [Eq. (2)]. This equation represents a forward model, an association between actions and sensory consequences. We now know that patients with cerebellar dysfunction are often severely impaired in learning this association, whereas basal ganglia damage appears to spare this ability (Maschke et al., 2004; Smith and Shadmehr, 2005). If the cerebellum is the crucial site for this learning, then the error signals that form forward models are likely delivered through climbing fiber activity.

With an accurate forward model [Eq. (2)], one can produce actions that achieve fair performance levels even without learning an optimal inverse model [Eq. (3)]. This is because imperfect motor commands can be rapidly compensated through internal feedback of the forward model. Indeed, for control of saccadic eye movements, the cerebellum appears to fit the role of a forward model that ''steers'' ongoing oculomotor commands rather well. However, optimal performance requires forming both an accurate forward model and an optimal inverse model. For example, if the robot is producing a force field, then learning an accurate forward model without an optimal inverse model will result in nearly straight trajectories. In theory, however, the optimum cursor trajectory that will result in the maximum probability of exploding the target is not a straight trajectory. Rather, the optimum cursor path is one that is slightly curved, resulting in an overcompensation of the forces early in the reach and an under compensation near the end. The reason for this is that because of noise associated with motor commands, producing larger commands early in the movement is a better policy than late in the movement because feedback allows one to correct the early errors but not the late errors. Interestingly, these curved cursor paths are similar to those that we have recorded in healthy volunteers (Thoroughman and Shadmehr, 2000).

In case of the reaching/robot task, theory predicts that the curved movements that appear to be a signature of optimal inverse models will not occur if training includes catch trials (randomly interleaved trials in which the force field is turned off). The over-compensation, which results from the large motor commands early in the trajectory, and the resulting curved movements, are optimal only in the case where one is certain that the force field will be present. If there are catch trials, then the optimal feedback control law is a path that is fairly close to a straight line.

Catch trials were part of the protocol with HM, and therefore we do not know if learning of optimal inverse models was intact in him. However, the error signals that guide formation of optimal inverse models need to link changes in expectations of cost and reward with the control policy that links sensory states to motor commands. As dopamine appears to be closely linked to reward prediction, it may play a central role in this learning. An interesting prediction of this hypothesis is that patients with basal ganglia damage should learn forward models normally, but be unable to learn optimal inverse models. To test this idea, current theoretical models of action need to advance so that they can predict tasks and conditions where subtle modifications of task characteristics can alter probability of success, resulting in clear changes in behavior in the optimal (presumably normal) learner.

The dissociation that we have made between learning motor commands that can produce rewarding states (optimal inverse models) and learning to predict sensory consequences of those commands (accurate forward models) can help explain a recent experimental result (Mazzoni and Krakauer, 2006). In our task, subjects reached to a target but could not see their hand. Instead, a cursor was presented on a vertical screen. The screen displayed eight targets, arranged around a

circle. On each trial one of these targets would be highlighted and the subject was instructed to take the cursor to that target. The novel part of this experiment was that the subjects were told that there is a 45° counter-clockwise (CCW) rotation of visual feedback during their reaching movements. Importantly, they were given a cognitive strategy to counter this perturbation: aim for the target 45° clockwise (CW) from the desired target in order to ensure that the cursor enters the desired target. The prediction was that successful implementation of the strategy would result in an abrupt stepwise cancellation of errors and the absence of after effects. The term aftereffect, seen in catch trials, refers to a trajectory deviation in the direction opposite to the imposed perturbation and indicates learning of an internal model.

Subjects were indeed initially effective in canceling the rotation with errors returning immediately to near zero. Surprisingly however, as subjects continued to make movements they made increasingly large directional errors, leading the cursor away from the desired target. This indicates that the verbal instructions provided them with the means to produce motor commands that resulted in rewarding states. However, because the rotation altered the relationship between the motor commands and the implicitly expected cursor path, the forward model began to adapt. This adaptation resulted in internal feedback during the reach, altering the trajectory and producing consistent errors. When subjects were asked about what they thought was happening, they often expressed frustration at the fact that they became progressively worse at hitting the target (the desired reward) and were unaware that they were adapting to the rotation. We interpret this result as evidence for the idea that despite the fact that they initially produced the optimum motor commands, never the less the forward model began to adapt, resulting in gradually incorrect movement. We think that the process of learning forward models may be distinctly separate from learning optimum control policies.

One way to view this is to assume that the motor system has two independent reward systems, one that evaluates the reward basis of a task and determines the success or failures of a control policy, and another that compares predicted sensory consequences of motor commands with their measured values, resulting in forward models. Our evidence from patients with PD who show a reluctance to make fast accurate reaching movements even though they are capable of making them, and from healthy subjects who adapt to a visuomotor rotation despite this being contrary to the explicit goal of the task, appears consistent with this framework.

Apraxia and the case of patient BG

With the exception of the idea that the motor system might have its own implicit reward system, the computational framework described thus far has dealt with relatively low-level aspects of motor control, namely execution and adaptation of reaching trajectories. However, patients with focal lesions of the CNS, especially of premotor and parietal regions, display a variety of higher-order motor disorders that have not been addressed to any great degree by this framework. Here we will focus on apraxia as an example of a syndrome where neuropsychological observations might be informed by a motor control perspective and vice versa.

Apraxia is a foreboding topic for a non-clinically oriented motor physiologist. Dictionary definitions of apraxia are usually a variation on the following: the inability of a person to perform voluntary and skillful movements of one or more body parts to command even though there is no muscle weakness, incoordination, sensory loss, aphasia or dementia. The unsatisfactory nature of this definition is apparent in the seemingly innumerable overlapping definitions and subtypes of apraxia that abound in the literature. This is undoubtedly because attempts to ascertain a core behavioral manifestation for apraxia have proven frustrating. The extant literature is largely descriptive and non-quantitative with few attempts to incorporate clinical phenomena into the emerging framework provided by basic research on sensorimotor integration and the parietal lobe (Leiguarda and Marsden, 2000). In contrast, recent research on motor control has used restrained, and arguably oversimplified and unnatural laboratory

based tasks to study motor learning, motor generalization and the role of context (Shadmehr and Wise, 2005). Thus, these tasks might be inadequate to elicit apraxic deficits. A caveat, however, is that patients with ideomotor apraxia have been shown to have kinematic abnormalities even for simple natural movements (Clark et al., 1994; Poizner et al., 1995). These kinematic abnormalities tend to be ignored by the neuropsychological and neurological literature (although see Ietswaart et al., 2006). Thus, there is an apparent divide between the approach to motor control taken by physiologists, in which the emphasis is on sensorimotor integration for simple movements, and the work on apraxia undertaken by clinical neuropsychologists, where the emphasis is on "higher order" motor deficits (see Table 1).

However, if there is to be progress it is critical for the two fields to learn from one another. In particular, it will allow us to determine whether there are categorical differences between sensorimotor integration seen, for example, during prism adaptation and sensorimotor integration when observed behavior has to be imitated. The case of HM attests to this need. Detailed neuropsychological description of this single patient over 50 years has yielded innumerable insights and avenues of investigation for neurophysiologists interested in the medial temporal lobe system (Corkin, 2002). Analogously, we would argue that motor control scientists stand to benefit greatly from taking a more careful look at apraxia.

Patients with "ideomotor" apraxia show all or a subset of the following motor execution abnormalities (see Leiguarda and Marsden, 2000; Koski et al., 2002, for review):

1. Spatiotemporal errors in the production of both over-learned and novel motor acts. For example, Poizner and colleagues have shown that patients asked to make a slicing movement as though cutting bread fail to show the correct phase relationships between pairs of joint angles (Poizner et al., 1995).
2. Impaired ability to pantomime symbolic gestures or tool-use to command or when given the tool itself. For example, a patient asked to imitate how they would use a pair of scissors may oppose their forefinger and index finger together as though they are themselves the blades of the scissors.
3. Impaired ability to imitate meaningful and meaningless motor acts. For example, a patient may not be able to copy a teeth-brushing motion or copy a random dance move.
4. An inability to adopt complex hand postures. For example, a dog-shaped shadow puppet made with the hands cannot be copied by a patient.

Patients with "ideational' and "conceptual" apraxia make higher-level errors compared to patients with ideomotor apraxia. They seem to have lost the semantic meaning (concept) of an action and can have trouble performing an action in the correct sequence. For example, while making a cup of tea they might pour the water into the cup before it has boiled and forget the tea bag. They can show choice of the wrong movement for a given

Table 1. An apparent divide between the approach to motor behavior taken by motor control physiologists versus that taken by clinical neuropsychologists

	Motor control approach	Neuropsychological approach
Questions	Trajectory control, adaptation, precision grip	Higher order motor deficits, e.g., apraxia, optic ataxia, and neglect
Experimental population	Healthy college students	Patients
Tasks	Restrained movements, e.g., planar reaching, sequential finger movements	3D everyday natural tasks, e.g., using knife and fork, mailing a letter, writing
Planning	Spatial accuracy, reaction time, scaling of early trajectory variables	Perform a task in the right sequence, e.g., making cup of tea
Execution	Continuous kinematic measures	Nominal scales, goal accomplishment
Conceptual framework	Quantitative: control theory, biomechanics, computational	Qualitative: neuroanatomical, descriptive, diagrammatic

transitive act. For example, they might make a chopping action when asked to pantomime brushing their teeth. They may also be impaired in recognizing a gesture or the purpose of a tool.

The majority of patients with apraxia have parietal lobe damage in the hemisphere contralateral to their dominant hand, but cases have been described with lesions in the premotor cortex bilaterally or lesions in the non-dominant parietal cortex (Halsband et al., 2001). There is currently no unifying explanatory framework for apraxic phenomena, which is perhaps not surprising given the multiplicity of parietal functions in motor control (Fogassi and Luppino, 2005; Culham and Valyear, 2006). We would argue that it would be more fruitful to map identified and putative physiological and computational processes in motor control onto specific apraxic phenomena. The goal is to have the clinical phenomena inform future experiments and models of motor control and parietal function.

Let us focus on two central features of ideomotor apraxia, both of which were present and rigorously described in patient BG, an important case study described by Buxbaum et al. (2000). We use BG as an illustrative case whilst recognizing that unlike HM, she did not have an identifiable circumscribed lesion. She had a primary progressive ideomotor apraxia, resulting in an inability to pantomime tool use either by verbal command, by viewing the object, or through imitation. However, when she held the tool in her hand, she was able to gesture and demonstrate its use near normally. That is, holding the tool allowed BG to recall the purpose of that tool and demonstrate its use through correct motor commands. This showed that the motor memory was present but not retrievable by verbal command or by looking at the object. When we compare this ability with HM, we see that it is rather important that with visual cues alone, HM demonstrated an ability to recall that the tool was associated with a rewarding behavior, resulting in his voluntary act of reaching for it, holding the handle in hand, and waiting for the task to begin. Furthermore, recall that once he held the tool, he was able to recall the motor commands necessary to control it.

For BG, why was semantic knowledge or visual observation of the tool insufficient to recall the memory of its purpose or how to hold it? As we described in the first section, knowledge of the behavior of a rotated cursor tool did not allow subjects to learn how to use it and in fact the motor system overrode healthy subjects' explicit knowledge (Mazzoni and Krakauer, 2006). Another clue to the failure of verbal commands or vision to call up the correct gesture comes from a large number of experiments that have tried to use visual cues to recall a particular internal model of a tool, either a robotic force field or rotated cursor (Cunningham and Welch, 1994; Gandolfo et al., 1996; Miall et al., 2004; Shadmehr et al., 2005). In these experiments, subjects learn to associate a force field or visuomotor rotation with a color or some other arbitrary symbolic cue. For example, blue for a CW rotation and red for a CCW rotation. Unexpectedly, subjects are either unable or require extensive training [monkeys had to be trained on a color-force field association for 12 months (Krouchev and Kalaska, 2003)] to use color as a cue to call up the associated internal model. This is strikingly reminiscent of the inability of BG to use visual cues to recall the appropriate gesture.

What was it about holding the tool itself, rather than the vision of the tool, which allowed BG to recall the appropriate gesture? An experiment we have recently performed in healthy subjects sheds light on this (Krakauer et al., 2006). In this experiment we hypothesized that the contextual cues subjects use to recall internal models are kinesthetic and implicit rather than visual and explicit. Specifically, we hypothesized that the relevant contextual cue is an implicit memory of action with a particular body part. To test this hypothesis we had subjects to learn a visuomotor rotation (the screen cursor's trajectory was rotated 30° CCW to the hand's trajectory) in what we conjectured would be two different contexts: by moving their hand through motion of their shoulder and elbow, or through motion of their wrist.

Our hypothesis was confirmed by the demonstration that subjects could recall opposite visuomotor rotations when each rotation was associated with a different body part. We would argue that the kinesthetic memory of the body part used to learn each rotation operates in a mechanistically related manner to how kinesthetic information

from handling of a tool allowed BG to recall the appropriate gesture. For example, we would predict that BG would be able to pantomine tool use if she first held the tool in her other hand. In the case of HM, we think that viewing the robot allowed him to remember the costs and rewards associated with the task, while the act of holding the robot allowed him to remember the forward model associated with control of the machine.

What remains to be explained in this framework is how healthy subjects use visual cues or semantic knowledge to recall internal models of everyday tools. We do not know the answer to this but suggest that it is a serial process in which multiple iterations with kinesthetic cues are required first before transitioning to visual cues. This makes intuitive sense when one considers the overshoot that occurs when we lift a cup that we think is full but is actually empty. We overshoot because we have lifted full cups in the past and retained the kinesthetic memory of the act.

A second prominent abnormality was BG's inability to imitate meaningless gestures. This finding, seen in some patients with ideomotor apraxia, has been a puzzle to neuropsychologists because a popular view of apraxia is that patients are impaired in their ability to recall stored motor programs. However, imitation of a meaningless gesture clearly cannot be related to impaired retrieval. Therefore, the explanation is likely separate from the contextual cuing effects discussed above.

In the study of BG, the authors conjectured that the difficulty with imitating gestures, nonsense gestures in particular, is due to a deficit in representing the relative positions of body parts (Buxbaum et al., 2000). In support of this notion is the observation that patients with posterior parietal lesions, unlike normal subjects or patients with motor cortex damage, show poor correlation between imagined and executed sequential finger movements, which suggests that they are not good at simulating movements. The hand laterality task is often used to assess such simulation. In this task one measures the time it takes for a subject to determine whether the picture is of a right hand or a left hand, has been shown to be a function of the position of the participant's own hand. Therefore, when one sees a picture of a hand, our ability to determine if it is a left or right hand is based on our ability to imagine translating and rotating our hand from its current posture to the viewed posture. Interestingly, in a recent study of 55 patients with left hemispheric stroke, there was a correlation between performance on the hand laterality task and the ability to imitate meaningless gestures (Schwoebel et al., 2004). This type of finding is consistent with the conjecture that the basis for BG's problem with meaningless gesture imitation may be due to a problem with her notion of body schema, which is defined as a representation of the relative positions of body parts derived from multiple sensory inputs and, perhaps, efference copy (Schwoebel and Coslett, 2005).

There is a problem with this viewpoint, however. The problem can be appreciated by considering patient PJ, an individual with an extra-axial cyst encroaching upon the left superior parietal lobule (Wolpert et al., 1998). Without vision of her right arm, PJ was unsure of where it was in space and in fact would feel as though it had disappeared. The unique aspect to her symptoms, as compared to patients who are deafferented from either central or peripheral lesions, was that she had normal tactile sensation and proprioception but that these proprioceptive sensations faded without vision of her right arm. The interpretation was that without vision, PJ had an inability to store a proprioceptively derived internal estimate of the state of her right arm and thus accumulated error over time. Notably, however, PJ was not apraxic. Thus, a deficit in limb state estimation that is correctible by vision is not likely to be the explanation for ideomotor apraxia in which viewing of tool use does not help imitation. Thus, notions of body schema (or estimation of limb state) are not sufficient explanation for the inability of patients with ideomotor apraxia to imitate transitive and nonsense gestures. Instead, it seems more plausible to relate these apraxic abnormalities in imitation of hand-held tool use and hand postures to the specific role of the IPL (and ventral premotor cortex) in visuomotor transformations for grasping (Sakata, 2003).

In their seminal work, Sakata and colleagues found that there were object-specific visual and visuomotor cells in the anterior intraparietal area

(AIP) of the IPL that coded for object shape and hand preshaping, respectively (Sakata, 2003). However, patients with ideomotor apraxia show a clear dissociation between impaired gesture imitation and intact reaching and grasping (Ietswaart et al., 2006). Thus, the ability to perform simple visuomotor transformations for grasping is relatively intact in patients with ideomotor apraxia. This suggests that visuomotor transformations related to actions *with* an object are distinct from visuomotor transformations needed to merely grasp an object. In retrospect this distinction is perhaps not surprising given that many of these patients perform much better with the tool itself and therefore will need to have grasped it in the first place. A recent fMRI experiment found that AIP was activated during both object-manipulation and observation of object manipulation by others (Shmuelof and Zohary, 2006). Their findings map well onto a recent study that showed a strong correlation between impairments in the imitation of transitive gestures and the recognition of object-related gestures and hand postures (Buxbaum et al., 2005). Finally, an fMRI study showed that left IPL activation was associated specifically with the somatic perception of hand-object interactions (Naito and Ehrsson, 2006), which suggests a particular form of sensorimotor integration between limb and object states. Damage to this area might explain why some patients with ideomotor apraxia have trouble making transitive gestures even when holding a tool.

Thus, the ability to improve ideomotor apraxia by holding a tool suggests that motor memories of tool use can be contextually triggered by implicit kinesthetic cues. However, the inability to imitate transitive and nonsense gestures, and the correlation of this inability with impairment in hand posture recognition, may suggest a specific role of the IPL (and ventral premotor cortex) in visuomotor transformations for object-related actions.

It should be apparent that the concepts derived from the current computational framework for motor control could only partially explain components of the ideomotor apraxia syndrome. These patients can be considered to have the implicit procedural analogs of the declarative abnormalities in HM. The study of higher order aspects of motor behavior in healthy subjects is difficult in the laboratory setting because complex perturbations are likely to be required. A productive synergy can therefore be envisaged whereby the computational models and experimental approaches derived from studies in healthy subject are applied and adapted to patients with higher-order motor abnormalities. Such an approach is a necessary complement to functional imaging studies in healthy subjects, which rely on activation differences to gain insight into higher order motor processes. Transcranial magnetic stimulation may prove effective in inducing higher-order errors in healthy subjects but as of yet has not yielded insights comparable to those obtained from patients. A rigorous computational approach in patients with motor disorders has been and will continue to be a fruitful avenue of investigation.

Acknowledgments

The work was supported by National Institutes of Health (NIH) grant K02 NS-048099 (JWK), NIH R01-037422 (RS), and a grant from the Human Frontiers Science Foundation (RS).

References

Buxbaum, L.J., Giovannetti, T. and Libon, D. (2000) The role of the dynamic body schema in praxis: evidence from primary progressive apraxia. Brain Cogn., 44: 166–191.

Buxbaum, L.J., Kyle, K.M. and Menon, R. (2005) On beyond mirror neurons: internal representations subserving imitation and recognition of skilled object-related actions in humans. Brain Res. Cogn. Brain Res., 25: 226–239.

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.

Clark, M.A., Merians, A.S., Kothari, A., Poizner, H., Macauley, B., Gonzalez Rothi, L.J. and Heilman, K.M. (1994) Spatial planning deficits in limb apraxia. Brain, 117(Pt 5): 1093–1106.

Corkin, S. (2002) What's new with the amnesic patient H.M.? Nat. Rev. Neurosci., 3: 153–160.

Culham, J.C. and Valyear, K.F. (2006) Human parietal cortex in action. Curr. Opin. Neurobiol., 16: 205–212.

Cunningham, H.A. and Welch, R.B. (1994) Multiple concurrent visual-motor mappings: implications for models of adaptation. J. Exp. Psychol. Hum. Percept. Perform., 20: 987–999.

Diedrichsen, J., Hashambhoy, Y., Rane, T. and Shadmehr, R. (2005) Neural correlates of reach errors. J. Neurosci., 25: 9919–9931.

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.

Fogassi, L. and Luppino, G. (2005) Motor functions of the parietal lobe. Curr. Opin. Neurobiol., 15: 626–631.

Gabrieli, J.D., Corkin, S., Mickel, S.F. and Growdon, J.H. (1993) Intact acquisition and long-term retention of mirror-tracing skill in Alzheimer's disease and in global amnesia. Behav. Neurosci., 107: 899–910.

Gandolfo, F., Mussa-Ivaldi, F.A. and Bizzi, E. (1996) Motor learning by field approximation. PNAS, 93: 3843–3846.

Halsband, U., Schmitt, J., Weyers, M., Binkofski, F., Grutzner, G. and Freund, H.J. (2001) Recognition and imitation of pantomimed motor acts after unilateral parietal and premotor lesions: a perspective on apraxia. Neuropsychologia, 39: 200–216.

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.

Hwang, E.J. and Shadmehr, R. (2005) Internal models of limb dynamics and the encoding of limb state. J. Neural Eng., 2: S266–S278.

Ietswaart, M., Carey, D.P. and Della Sala, S. (2006) Tapping, grasping and aiming in ideomotor apraxia. Neuropsychologia, 44: 1175–1184.

Koski, L., Iacoboni, M. and Mazziotta, J.C. (2002) Deconstructing apraxia: understanding disorders of intentional movement after stroke. Curr. Opin. Neurol., 15: 71–77.

Krakauer, J.W., Ghez, C. and Ghilardi, M.F. (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. J. Neurosci., 25: 473–478.

Krakauer, J.W., Ghilardi, M.F. and Ghez, C. (1999) Independent learning of internal models for kinematic and dynamic control of reaching. Nat. Neurosci., 2: 1026–1031.

Krakauer, J.W., Ghilardi, M.F., Mentis, M., Barnes, A., Veytsman, M., Eidelberg, D. and Ghez, C. (2004) Differential cortical and subcortical activations in learning rotations and gains for reaching: a PET study. J. Neurophysiol., 91: 924–933.

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

Krakauer, J.W., Pine, Z.M., Ghilardi, M.F. and Ghez, C. (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. J. Neurosci., 20: 8916–8924.

Krouchev, N.I. and Kalaska, J.F. (2003) Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues. J. Neurophysiol., 89: 1165–1175.

Leiguarda, R.C. and Marsden, C.D. (2000) Limb apraxias: higher-order disorders of sensorimotor integration. Brain, 123(Pt 5): 860–879.

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.

Maschke, M., Gomez, C.M., Ebner, T.J. and Konczak, J. (2004) Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. J. Neurophysiol., 91: 230–238.

Mazzoni, P., Hristova, A. and Krakauer, J.W. (2007) Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. J. Neurosci., 27: 7105–7116.

Mazzoni, P. and Krakauer, J.W. (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. J. Neurosci., 26: 3642–3645.

Miall, R.C., Jenkinson, N. and Kulkarni, K. (2004) Adaptation to rotated visual feedback: a re-examination of motor interference. Exp. Brain Res., 154: 201–210.

Milner, B. (1968) Pathologie de la memoire. La Memoire, Geneve, pp. 185–212.

Mishkin, M., Malamut, B. and Bachevalier, J. (1984) Memories and habits: two neural systems. In: Lynch G. and MacGaugh J. (Eds.), Neurobiology of Learning and Memory. Guilford Press, pp. 65–77.

Naito, E. and Ehrsson, H.H. (2006) Somatic sensation of hand-object interactive movement is associated with activity in the left inferior parietal cortex. J. Neurosci., 26: 3783–3790.

Padoa-Schioppa, C., Li, C.S. and Bizzi, E. (2002) Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. Neuron, 36: 751–765.

Poizner, H., Clark, M.A., Merians, A.S., Macauley, B., Gonzalez Rothi, L.J. and Heilman, K.M. (1995) Joint coordination deficits in limb apraxia. Brain, 118(Pt 1): 227–242.

Sakata, H. (2003) The role of the parietal cortex in grasping. In: Siegel A.M., Andersen R.A., Freund H.-J. and Spencer D.D. (Eds.), Advances in Neurology, Vol. 93. Lippincott Williams & Wilkins, Philadelphia, PA, pp. 121–139.

Schwoebel, J., Buxbaum, L.J. and Coslett, H.B. (2004) Representations of the human body in the production and imitation of complex movements. Cogn. Neuropsychol., 21: 285–298.

Schwoebel, J. and Coslett, H.B. (2005) Evidence for multiple, distinct representations of the human body. J. Cogn. Neurosci., 17: 543–553.

Shadmehr, R., Brandt, J. and Corkin, S. (1998) Time-dependent motor memory processes in amnesic subjects. J. Neurophysiol., 80: 1590–1597.

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., Donchin, O., Hwang, E., Hemminger, S. and Rao, A. (2005) Learning dynamics of reaching. In: Riehle A. and Vaadia E. (Eds.), Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions. CRC Press, Boca Raton, FL, pp. 297–328.

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

Shadmehr, R. and Wise, S.P. (2005) The Computational Neurobiology of Reaching and Pointing: A Foundation for Motor Learning. The MIT press, Cambridge, MA.

Shmuelof, L. and Zohary, E. (2006) A mirror representation of others' actions in the human anterior parietal cortex. J. Neurosci., 26: 9736–9742.

Smith, M.A. and Shadmehr, R. (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J. Neurophysiol., 93: 2809–2821.

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

Todorov, E. and Jordan, M.I. (2002) Optimal feedback control as a theory of motor coordination. Nature Neurosci., 5: 1226–1235.

Tranel, D., Damasio, A.R., Damasio, H. and Brandt, J.P. (1994) Sensorimotor skill learning in amnesia: additional evidence for the neural basis of nondeclarative memory. Learn. Mem., 1: 165–179.

Wainscott, S.K., Donchin, O. and Shadmehr, R. (2005) Internal models and contextual cues: encoding serial order and direction of movement. J. Neurophysiol., 93: 786–800.

Wigmore, V., Tong, C. and Flanagan, J.R. (2002) Visuomotor rotations of varying size and direction compete for a single internal model in motor working memory. J. Exp. Psychol. Hum. Percept. Perform., 28: 447–457.

Wolpert, D.M., Goodbody, S.J. and Husain, M. (1998) Maintaining internal representations: the role of the human superior parietal lobe. Nat. Neurosci., 1: 529–533.

Yamashita, H. (1993) Perceptual-motor learning in amnesic patients with medial temporal lobe lesions. Percept. Mot. Skills, 77: 1311–1314.