

Primer

Motor learning

Daniel M. Wolpert¹
and J. Randall Flanagan²

Although learning a motor skill, such as a tennis stroke, feels like a unitary experience, researchers who study motor control and learning break the processes involved into a number of interacting components. These components can be organized into four main groups. First, skilled performance requires the effective and efficient gathering of sensory information, such as deciding where and when to direct one's gaze around the court, and thus an important component of skill acquisition involves learning how best to extract task-relevant information. Second, the performer must learn key features of the task such as the geometry and mechanics of the tennis racket and ball, the properties of the court surface, and how the wind affects the ball's flight. Third, the player needs to set up different classes of control that include predictive and reactive control mechanisms that generate appropriate motor commands to achieve the task goals, as well as compliance control that specifies, for example, the stiffness with which the arm holds the racket. Finally, the successful performer can learn higher-level skills such as anticipating and countering the opponent's strategy and making effective decisions about shot selection. In this Primer we shall consider these components of motor learning using as an example how we learn to play tennis.

Information extraction


Movement allows us to determine when and where to place our sensory receptors — for example, those in our retinas or on our fingertips — and this allocation of our sensory resources can be made in a task-specific manner. For example, when facing a bowler in cricket, experienced batters will fixate the point at which the bowler releases the ball, make a saccade to the anticipated bounce point, and then use pursuit eye movement to track the ball to the contact point with the bat. Studies show that experts generate

eye movements sooner and more accurately than novices, supporting the idea that learning where and when to direct gaze is a key component of learning a motor skill [1]. Moreover, in the laboratory it has been shown that, even in simple visual search tasks in which the subject must find a visual target among distractors, each gaze fixation location is selected to maximize information — and hence minimize uncertainty — about the target location given the sequence of previous fixations [2].

Similarly, when manipulating an object, the type of tactile exploration

depends on the information we are trying to extract. Even when our sensors receive the same sensory input, attentional and other constraints mean that the task determines which sensory information is actually processed. Task-irrelevant information is often not perceived, leading to the phenomena of inattention blindness [3].

Finally, the extraction of information does not simply rely on the incoming sensory stream but is strongly shaped by previous experience. This is captured by the Bayesian framework in which noisy sensory



Current Biology

Figure 1. Structural learning.

The two rackets in the upper panel share a similar structure in terms of their geometry and dynamics but with different parameter such as length and weight. However, the frisbee has a different structure from rackets in terms of both its geometrical and dynamic properties. Structural learning involves acquiring knowledge of the way in which different objects or tasks share similar properties. Parametric learning involves setting the particular parameters for a given object or task having identified the structure.




Figure 2. Control of stiffness.

By varying the activation of the set of muscles in the arm the stiffness properties of the racket can be controlled. The red ellipsoid shows the stiffness with the long axis representing the directions of high stiffness. For a smash the racket is held stiff so as to maintain the energy in the ball whereas for a drop-shot the stiffness perpendicular to the racket head is low to allow any viscosity to absorb the ball's energy.

evidence, such as the visual estimate of where a tennis ball is likely to bounce, is combined with prior knowledge, such as the distribution of where the ball has bounced from your opponent's previous shots, to produce an optimal (most accurate) estimate of the bounce location [4]. Therefore, extracting task-relevant information is a highly active and learned process in which we can decide what sensory information to sample, what information to process from the sample and how to extract the information in an efficient, and perhaps even optimal, manner.

Learning features of the task

A key component that underpins our actions is learning the features relevant for a task. For example, in tennis we must learn the transformation between muscle

commands and the motion of the racket head, learn how to credit errors to different aspects of our performance and determine how the context — such as court surface or prevailing wind conditions — affects the task.

Structural and parametric learning

The appropriate motor commands in tennis will depend on the geometry and mechanics of the tennis racket that specify the relation between the hand's posture and the location of the racket head (termed the kinematic or visuomotor transformation) as well as the way the racket responds to forces and torques applied by the hand and ball (dynamic transformation). In general, there are two levels we can consider when learning such transformations. The first, structural learning, involves identifying the

appropriate inputs (for example, motor commands or forces) and output (for example, racket motion) of the system and the form of the equations that link the two. For example, the set of all rackets have common structural properties (Figure 1) in terms of their geometrical and dynamic (for example, inertia and viscosity) features which distinguish them from the set of all frisbees or set of all scissors, each of which has a different structural form. When faced with a novel task, the structure of the task needs to be learned through experience.

The second level, parametric learning, involves learning the particular parameter settings for a given structure. For example, when playing with a particular racket the player needs to learn its particular mass, weight distribution and string compliance. Recent studies have shown that, in the laboratory, structural learning can be induced by exposing participants to a randomly varying set of tasks that share a common structure but vary in their parameter settings [5]. Such structural learning dramatically speeds up learning of new tasks that share the same structure because once the learner identifies the structure they need only adjust the key parameters appropriate for that structure. Numerous studies of adaptation to relatively simple visual and force perturbations report immediate and monotonic improvements in performance, presumably reflecting parametric learning. In other more complex tasks, however, no improvement is seen during initial exposure. It is likely that here there is an initial exploratory period during which subjects must discover the structure of the task before any improvement in performance is seen [6].


Context and credit assignment

In motor learning, as in all learning, context is critical. When we learn new dynamics or kinematics, we must also be able to link this learning to appropriate objects, tasks or environments. Studies of motor learning have typically looked at the role of context by testing how learning in one context generalizes to other contexts, or whether interference in learning, which is often seen between motor tasks, can be diminished

when context is varied. It has been shown that arbitrary cues, such as the background light color, have little effect in helping learning of multiple tasks. However, more natural contextual cues, such as the visual orientation of a manipulated object (such as the location of the racket head) or whether the two arms act on one or two separate objects (such as a one *versus* two armed backhand) have strong contextual effects which facilitate motor learning [7].

Within a given context, an important issue in motor learning is the problem of credit assignment. If a tennis player starts hitting her shots into the net, the problem could be that the ball is heavy, that the racket strings are loose, that there is an oncoming wind, that she is fatigued, or that she has grown since last she played and has not yet adapted. Although the immediate response — hit the ball harder — may be the same in all of these scenarios, correctly crediting the problem is essential for shaping learning. For example, if the racket dynamics have changed, she would do well to learn and remember these dynamics if she plans to use the racket again. Conversely, if she is simply tired or it is windy she needs to make a temporary adjustment and perhaps learn how to calibrate herself for the wind or fatigue. Recent work has examined such credit assignment in terms of allotting the cause of errors to changes in the properties of the body *versus* the external world including objects such as a held tool [8]. Using a Bayesian formulation, the model allots the errors in proportion to the optimal estimate of where the errors arise from, and can account for a range of empirical data.

While credit assignment can allocate errors spatially across effectors and tools, recent work has shown that errors appear to be allocated across modules that learn with different time scales. Recent work on force-field and visuomotor adaptation has provided evidence that learning involves (at least) two parallel processes, a fast process that adapts and de-adapts quickly and a slower process that adapts and de-adapts more gradually [9]. Rapid learning mechanisms enable the performer to deal with potential short-lived perturbations. If these perturbations persist over time, then the slower mechanisms, which are longer lasting,



Current Biology

Figure 3. Minimum intervention principle.

Variation in the azimuth angle of the racket head (left panel) leads to variability in the ball landing location (red ellipse) that is distributed within the court and therefore is less deleterious to the task (or task-irrelevant for a novice player) compared to variations in elevation angle (right panel) which can lead to the ball landing outside the court. The minimum intervention principle suggests that azimuthal variation should be corrected for less strongly than elevation variability.

adapt. Just as credit assignment may enable learners to identify the causes of perturbations, credit assignment can also be used to identify the longevity of these perturbations. Of course, these two characteristics are often linked; perturbations due to fatigue or the wind are typically brief, whereas those due to a new racket or bodily change are long lasting. Moreover, laboratory studies have shown that rather than being fixed, the time constants themselves can be influenced by the rate of change of perturbations previously experienced in the task [10].

Classes of control

We can consider several processes that need to function together in order to generate the appropriate motor command: learning how to modulate the compliance of the hand to absorb or maintain the energy of the ball; how to make appropriate postural adjustments to support the stroke; and how to use sensory inputs to update outgoing commands.

Stiffness, reactive and predictive

control

One of the major challenges facing motor performance is the large time delays inherent in sensorimotor feedback loops that limit the rapidity with which the motor system can respond to sensory events. There are three mechanisms, each of which can undergo learning, that can alleviate the problems of time delays. The first mechanism is to vary the compliance of the body by co-contraction of specific muscles. For example, by varying the activations of a set of muscles in the arm it is possible to control the stiffness at the hand or racket head. Not only can people scale their overall stiffness (hence compliance), they can also shape the pattern of stiffness either by varying muscle activations or the posture of the arm [11]. For example, when using a knife to cut into an apple, lateral deviations of the blade from the top of the apple could lead to the knife slipping. In such a task, stability can be maintained by stiffening up in the axis perpendicular to the blade but not increasing stiffness parallel to the blade or in the vertical direction.

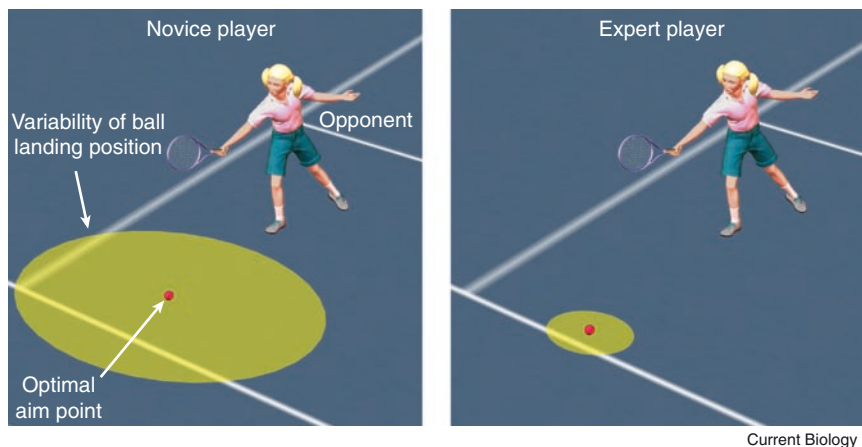


Figure 4. Optimal aim location depends on variability.

The closer a shot is aimed to the line, the further it will be from the opponent, making it less likely she will be able to return the ball. However, due to variability (yellow area) the closer to the line the greater the chance that the ball will land outside the court. There is, therefore, an optimal location to aim (red ball) to maximize the chance of winning the point, which trades off the probability of the ball landing inside the court with the probability of the ball being successfully returned by the opponent. For a novice player (left) who has a large amount of variability the optimal location is further inside the court than for an expert (right) who has a small amount of variability.

It has been shown in such tasks that subjects are able, albeit to a limited extent, to shape their stiffness to match the task requirements.

To return to our tennis example, the player can increase stiffness perpendicular to the racket head when hitting a smash forehand or reduce it to take the energy out of the ball for a drop-shot (Figure 2). By modulating stiffness, the motor system can exercise rapid control over the response to external perturbations. Although stiffness can be used to deal with some perturbations, it is limited in its flexibility and, because it requires co-contraction, can be an effortful solution to maintaining stability. Therefore, in many instances skilled performance requires the neural processing of sensory information during the task.

A second mechanism which can alleviate some problems with time delays is to use fast reactive feedback loops to drive motor responses. While the fastest of these, such as the mono-synaptic stretch reflex, is not modifiable by experience, longer loop reflexes that can involve supraspinal mechanisms can be modified in a task-dependent manner [12]. Thus, there is a trade-off across the set of sensorimotor responses with the fastest being less adaptable and the slower being more flexible. Ideally, these responses work in concert

with the most rapid reflexes holding the fort waiting for the cavalry in the shape of the slower and more task-dependent responses. Considerable work has shown that the gain and even sign of these longer reflexes can be tuned in a time-dependent manner to match properties of the task being performed.

Finally, the third mechanism, that of prediction can be used to generate the appropriate command to compensate for upcoming and predictable perturbations. That is, motor commands can be generated in anticipation of the task requirements [13].

How these three mechanisms interact during learning has become a key question in the field. Most tasks involve a combination of all three mechanisms, with the contribution of each depending on the nature of the task. In terms of the interaction between stiffness and predictive control it has been proposed that when errors are large, stiffness increases but as the predictive component begins to learn, and errors reduce, stiffness decreases [14].

Optimal feedback control

There are many examples of tasks, ranging from precision lifting to locomotion, where the processing of sensory information in feedback loops is highly task-dependent and goal-directed. Indeed, the selection

of task-appropriate sensorimotor feedback mechanisms designed to handle errors is part and parcel of planning the motor task. Recently, the idea of optimally shaping the control mechanisms to task goals has been formalized in the theoretical framework of optimal feedback control [15] and this has led to a resurgence of interest in feedback control. This framework suggests that the central nervous system sets up feedback controllers that continuously convert sensory inputs into motor outputs that are optimally tuned to the task being performed by trading off energy consumption with constraints on performance, such as accuracy.

An important feature of the model is the concept of minimum intervention; that is, setting up feedback controllers that only correct for variation that is deleterious to the task. In the tennis example shown in Figure 3, variations in the angle of the racket head about the azimuth (left panel) have little effect on whether the ball will land in the court whereas variations about the elevation angle (right panel) can threaten the goal of landing the ball in the court. Therefore, variations in the angle about the azimuth do not need to be corrected for as strongly as variations in elevation. Not only are corrections of task-irrelevant errors wasteful, they can also generate task-relevant errors. A key component of motor learning involves exploring the task so as to learn which errors threaten the goal and need to be dealt with via sensorimotor feedback mechanisms.

Decisions and strategies

Historically there has been relatively little interaction between work in motor control and work in cognitive domains such as decision making and attention. However, as sensorimotor researchers have broadened the scope of the tasks under study, the distinctions among sensorimotor, perceptual, and cognitive components of the task, including action selection and decision making, have become blurred. For example, one interesting area in which this interaction has recently manifest itself is research using motor tasks to look at decision making. In many explicit cognitive tasks, people often make suboptimal judgments when faced with a set

of decisions each of which has an uncertain outcome. In contrast, when people encounter motor variants of these tasks, they are often close to optimal in their behavior

For example, when pointing to target configurations that have different reward and penalty regions, it has been shown that subjects are able to choose their average pointing location so as to minimize the loss that accrues through the variability of pointing [16]. Such optimization applies in tennis where there is a trade-off between placing the ball far away from your opponent and keeping the ball within the court. For a novice player, who has a lot of variability, it is optimal to aim quite far inside the lines to maximize the chances of winning the rally – getting the ball in and preventing the opponent from returning it (Figure 4 left panel). On the other hand, the optimal location for an expert player, with less variability, will be closer to the lines and further away from the opponent (Figure 4 right panel).

At a higher level, tennis requires strategic decisions in multi-player interactions. Such decision making is typically examined within the framework of game theory and work in the cognitive domain has shown that when people have to make decisions based on a set of rules, they are typically sub-optimal. Recently, classic game theoretic problems such as prisoners' dilemma have been examined in the motor domain by translating years-in-prison into movement effort. In such two-player motor games, subjects rapidly develop near-optimal game-theoretic solutions, that is the players adopting Nash equilibrium solutions in which they choose actions so that neither has anything to gain by changing only his or her strategy [17]. Such motor game-theoretic interactions arise naturally, for example, in doubles where each player must cooperate with their partner while competing with their opponents.

Tennis enthusiasts invest considerable time, energy, and money into trying to improve their game. The secret is [to sharpen your] game [theory], set [your controllers optimally] and match [your sensory processing to the task].

Acknowledgments

This work was supported by the Wellcome

Trust, SENSOPAC (IST-2005- 028056), the Canadian Institutes for Health Research and Natural Science and Engineering Research Council of Canada.

References

1. Land, M.F., and Tatler, B.W. (2009). *Looking and acting*, (Oxford, UK: OUP).
2. Najemnik, J., and Geisler, W.S. (2005). Optimal eye movement strategies in visual search. *Nature* 434, 387–391.
3. Most, S.B., Scholl, B.J., Clifford, E.R., and Simons, D.J. (2005). What you see is what you set: sustained inattention blindness and the capture of awareness. *Psychol. Rev.* 112, 217–242.
4. Körding, K.P., and Wolpert, D.M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* 10, 319–326.
5. Braun, D.A., Aertsen, A., Wolpert, D.M., and Mehring, C. (2009). Motor task variation induces structural learning. *Curr. Biol.* 19, 352–357.
6. Sailer, U., Flanagan, J.R., and Johansson, R.S. (2005). Eye-hand coordination during learning of a novel visuomotor task. *J. Neurosci.* 25, 8833–8842.
7. Diedrichsen, J. (2007). Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr. Biol.* 17, 1675–1679.
8. Berniker, M., and Körding, K.P. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci.* 11, 1454–1461.
9. Smith, M.A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 4, e179.
10. Huang, V.S., and Shadmehr, R. (2009). Persistence of motor memories reflects statistics of the learning event. *J. Neurophysiol.* 102, 931–940.
11. Franklin, D.W., Liaw, G., Milner, T.E., Osu, R., Burdet, E., and Kawato, M. (2007). Endpoint stiffness of the arm is directionally tuned to instability in the environment. *J. Neurosci.* 27, 7705–7716.
12. Pruszynski, J.A., Kurtzer, I., and Scott, S.H. (2008). Rapid motor responses are appropriately tuned to the metrics of a visuospatial task. *J. Neurophysiol.* 100, 224–238.
13. Johansson, R.S., and Flanagan, J.R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nat. Rev. Neurosci.* 10, 345–359.
14. Franklin, D.W., Burdet, E., Tee, K.P., Osu, R., Chew, C.-M., Milner, T.E., and Kawato, M. (2008). CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.* 28, 11165–11173.
15. Todorov, E. (2004). Optimality principles in sensorimotor control. *Nat. Neurosci.* 7, 907–915.
16. Trommershäuser, J., Maloney, L.T., and Landy, M.S. (2008). Decision making, movement planning and statistical decision theory. *Trends Cogn. Sci.* 12, 291–297.
17. Braun, D.A., Ortega, P.A., and Wolpert, D.M. (2009). Nash equilibria in multi-agent motor interactions. *PLoS Comput. Biol.* 5, e1000468.

¹Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK. ²Department of Psychology and Centre for Neuroscience Studies, Queen's University, Kingston, ON K7L 3N6, Canada.
E-mail: wolpert@eng.cam.ac.uk