

Feature Review

Learning to move machines with the mind

Andrea M. Green and John F. Kalaska

Groupe de Recherche sur le Système Nerveux Central, Département de Physiologie, Université de Montréal, C.P. 6128, Succursale Centre-ville, Montréal, Québec H3C 3J7, Canada

Brain–computer interfaces (BCIs) extract signals from neural activity to control remote devices ranging from computer cursors to limb-like robots. They show great potential to help patients with severe motor deficits perform everyday tasks without the constant assistance of caregivers. Understanding the neural mechanisms by which subjects use BCI systems could lead to improved designs and provide unique insights into normal motor control and skill acquisition. However, reports vary considerably about how much training is required to use a BCI system, the degree to which performance improves with practice and the underlying neural mechanisms. This review examines these diverse findings, their potential relationship with motor learning during overt arm movements, and other outstanding questions concerning the volitional control of BCI systems.

The BCI: transforming volition into action via assistive technology

The remarkable ease with which we interact with the world with our arms and hands is lost in patients with severe motor deficits. BCI technology (Glossary) shows great potential to help patients perform everyday tasks, such as feeding and grooming themselves or using a computer and entertainment devices, by recording their brain activity to extract signals about their motor intentions. These signals are converted into actions of remote devices (effectors) ranging from computer cursors to anthropomorphic robots [1–7] and even into contractions of the subject's own muscles by electrical stimulation [5,8,9].

The ideal BCI system should allow subjects to control effectors without extensive training by simply thinking about what they want the effector to do. In human subjects, such mental efforts activate neural circuits normally involved in overt motor control. For instance, cerebral cortical motor areas become active when paralyzed humans are asked to try to move body parts or to imagine that they are moving an effector [1–7,10–16]. The neural activity evoked during those covert volitional motor efforts has several properties in common with that recorded in able-bodied subjects during overt arm movements [14–16]. However, asking subjects to engage in introspective motor imagery is not as demanding as requiring them to impose efficient BCI control over the motions of an actual physical effector. This is especially true if the effector does not follow their

motor intentions as faithfully in the real world as it likely does in their mental imagery.

Most BCI development has been accomplished using able-bodied animals such as monkeys and rats [1–8,17–21] that cannot respond to verbal instructions to generate movement-related brain activity. Instead, they are typically first trained in behavioral tasks to control the motions of effectors by making free limb movements or operating a joystick or similar tool. Only then is the task-related brain

Glossary

Anterograde interference: the phenomenon whereby acquisition of a new skill is impeded by the prior acquisition of a different skill.

Brain–computer interface (BCI): a system which records neural activity, decodes that activity to extract information about a subject's intended actions and uses that decoded information to control a remote effector (see below). Other names for this technology include: brain-controlled, brain–machine or neural interface, and neural or neuromotor prosthesis.

Effector: the physical object controlled by the brain to achieve desired goals and actions. In overt arm motor-control, it is the arm itself. In the context of BCI control, it is usually a tool such as a computer cursor, television remote, wheelchair or robotic prosthetic device.

Feed-forward: in arm motor-control, feed-forward control refers to the predictive neural processes that generate the control signals expected to drive the arm to its desired goal state, such as a reaching movement to a target. Feed-forward control will produce errors unless it has perfect knowledge of how the arm will respond to a motor command across all conditions.

Feedback: in arm motor-control, visual or proprioceptive sensory inputs provide signals (feedback) about the response of the arm (current state) to the outgoing motor command. A feedback control loop is established when the feedback signal is used to modify the outgoing motor command to influence the continuing evolution of the movement or to permit adaptive changes to future motor commands in response to sensed errors.

Internal model: as used here, the term describes a neural representation of the properties of a physical system or a physical relationship. Such properties could often be described by one or more mathematical equations. Internal model refers to the implementation of the computations/transformations implied by such equations by neural populations. In limb motor-control, internal models refer to neural networks which implement various aspects of the computations that describe the properties of the arm and its interactions with the environment. Such internal models can be used to help transform desired goals into commands for action or to transform sensory information and efference copies of motor commands into current-state estimates of actual behavior and future-state predictions of how the arm will respond to a motor command.

Extrinsic motor parameter: parameters that describe the spatial properties of a movement, such as the spatial location of targets and the direction or velocity of hand movement through space.

Intrinsic motor parameter: parameters that describe a movement in terms centered on different body parts, such as joint angles and joint rotations or muscle lengths and muscle length changes.

Retrograde interference: the phenomenon whereby the acquisition of a new skill interferes with the retention or recall of a previously learned skill.

Viscous–curl field: an external force field whose magnitude is proportional to movement velocity (viscous force) and is applied in the direction perpendicular to the direction of movement (curl).

Corresponding author: Kalaska, J.F. (kalaskaj@physio.umontreal.ca).

activity appropriated by the BCI to directly control effector motions, bypassing the intermediary of the arm and tool.

Reports of the time required for subjects to acquire BCI control range from minutes to weeks [1–21]. Furthermore, much like an infant’s arm movements, the BCI-mediated effector motions initially achieved by subjects are usually slow, often clumsy and sometimes unsuccessful. Many studies report improvements in BCI control with practice [1–8,10–13,17–21], suggesting that it is a skill that can be acquired in the same way as overt motor skills. However, the improvement is often poorly documented, and no subject to date has achieved the same precision and speed of effector control as during normal overt arm movements. Finally, motor skill acquisition is widely assumed to require changes in neural activity. However, the activity changes that have been reported after subjects begin to use a BCI system are diverse and often difficult to interpret. The diversity of reports on these issues in part reflects whether the primary objectives of each study concern technical development of BCI systems or the investigation of the brain mechanisms underlying BCI use. We will examine these and other issues related to the neural

mechanisms of BCI control in this review. To provide some context, however, we will first provide an overview of some basic concepts in relation to the brain mechanisms underlying voluntary arm movements. For practical technical reasons, most BCI systems extract brain signals from the cerebral cortex. Therefore we will focus on cerebral cortical mechanisms of overt motor control, while acknowledging that subcortical circuits also make essential contributions.

Voluntary control and adaptation of overt arm movements

To make a voluntary movement, the motor system must convert a desired goal (e.g. drink some coffee) into a plan of action (reach to your coffee cup) and ultimately into the spinal motoneuron activity that produces the required muscle contractions. This involves coupled open-loop feed-forward and closed-loop feedback control mechanisms. Feed-forward processes performed by neurons distributed throughout the supraspinal motor system convert the goal into a motor command that is further transformed by spinal cord circuits into muscle activity. Proprioceptive and visual feedback closes the control loop by informing the

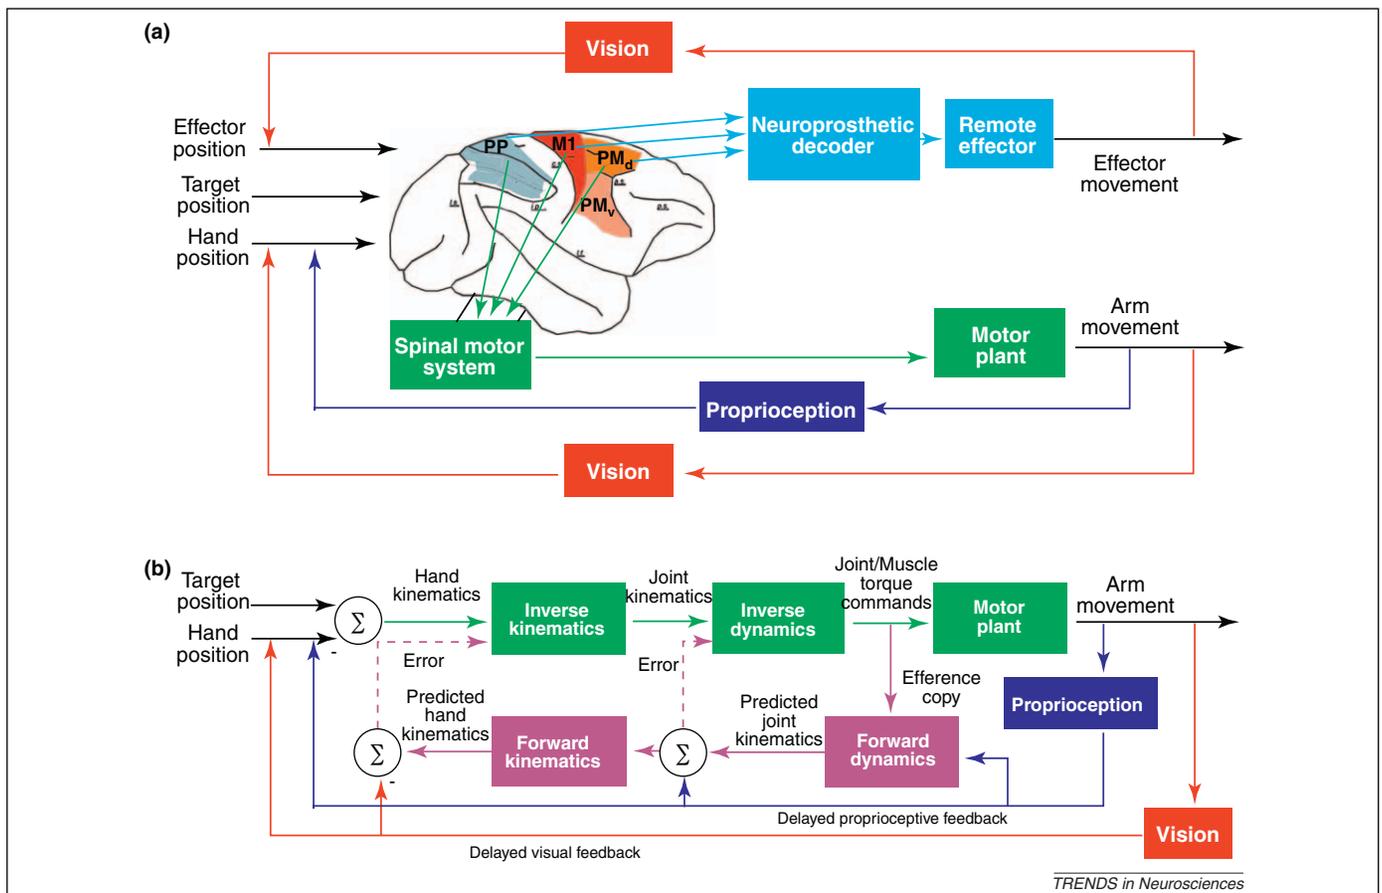


Figure 1. Comparison of the control systems involved in overt arm motor-control versus those for BCI control. (a) During normal overt arm control (bottom), feed-forward motor command signals (green pathways) generated in part in the primary motor cortex and premotor cortical areas are transmitted via multiple descending motor pathways to the spinal cord, where the information is transformed by spinal motor circuits into muscle activation patterns that are then converted by the musculoskeletal motor plant into arm movements. Visual (red) and proprioceptive (dark blue) sensory feedback signals provide information about hand position. By contrast, during BCI control (top) of a remote effector with no overt arm movements, neural activity recorded in cortical motor areas is transformed by a decoder into control signals to an effector (light blue pathways). Visual (red) feedback provides information about the resulting effector motions. M1, primary motor cortex; PM_d, dorsal premotor cortex; PM_v, ventral premotor cortex; PP, posterior parietal cortex. (b) Schema illustrating different internal models that describe various aspects of the properties of the limb (e.g. inverse kinematics: the relationship between hand position in space and limb joint angles; inverse dynamics: the relationship between joint angles and muscle force). These models are used to perform feed-forward sensorimotor transformations (green) to convert signals about target location and current hand position into descending motor commands or to transform sensory information and motor-command (efference copy) signals within feedback pathways (purple) into estimates of current behavioral state. Such state estimates can be compared with information about desired behavior to generate on-line error-correction signals.

motor system about the physical consequences of the motor command (Figure 1a).

The neural mechanisms that convert a goal into a movement are often described as a sequence of sensorimotor transformations from sensory inputs to motor outputs (Figure 1b) [22–26]. Two putative transformations are of interest here – the conversion of visual input about a target location into the desired reach trajectory, and the subsequent transformation of the desired trajectory into motor commands that evoke the required muscle activity. According to an influential hypothesis in the field, the motor system implements those transformations via adaptive internal models (IMs) – neural representations of the physical properties of the arm and its interactions with the world. Those properties determine the computations required to transform desired goals into motor commands and predict the physical consequences of those commands [27–30] (Figure 1b).

The transformation implemented by an IM in turn determines the response properties of its constituent neurons, including the combinations of sensory and motor-related signals to which they respond and how they discharge during movement [30–34]. Each neuron usually displays a particular activation pattern (e.g. Gaussian, cosine, planar) over a broad but limited range of different extrinsic or intrinsic motor parameters such as movement direction, arm posture, joint angles, output forces or muscle activity [22,25,26]. Motor adaptation would involve synaptic input weight changes onto neurons involved in one or more transformations, resulting in changes in their task-related activation patterns and thus in the transformations that they implement (Box 1). This suggests that the response properties of the neurons in the IMs undergoing adaptive changes determine how the motor system learns new skills. The mechanisms underlying motor skill learning in turn determine how neural activation patterns change [30,35–39].

Behavioral studies of motor adaptation

One widely studied type of motor learning is adaptation of reaching movements to a change in the relationship be-

tween desired goals and required motor commands. For instance, when an external force field repeatedly deviates the arm from its intended trajectory, the motor system learns to generate new patterns of forces and muscle activity to restore the originally-intended movement [29,35–38] (Figure 2a–f). By contrast, when a visuomotor dissociation (e.g. prisms, visual rotations) alters the relationship between visual input and arm movements, the motor system learns to make a different movement in response to a given visual input without altering the forces and muscle activity that cause that movement [39–43]. Force-field and visuomotor adaptation presumably require changes in different sensorimotor mappings in different IMs, but share several key properties.

First, in both situations the motor system uses sensory feedback about errors arising in a given trial to make incremental changes to the appropriate sensorimotor mapping to permit predictive feed-forward compensation for the conditions expected in the next trial [35,36,44]. This results in a learning curve of gradual improvement across many trials whose time-course can reflect multiple processes with different rates of learning and forgetting [45,46]. Adaptation can also be influenced by such factors as the degree of predictability of task conditions and whether the motor system assigns credit for the errors to a change in the world, a change in the properties of the limb, or incorrect calculation of the motor command [45–53].

Second, skills learned in one reach direction or arm posture can generalize (i.e. transfer) to untrained directions or postures, but the extent of generalization decreases as the directions and postures become increasingly different from those in which the skill was learned [29,35,36,41,42,54–57]. The observed generalization patterns resemble the approximately sinusoidal directional and planar postural activation patterns of many motor cortical neurons, consistent with theoretical predictions that generalization patterns reflect the properties of the neurons that are adapting [22–26,30,35,36,41,42,54–57].

Box 1. The origin and functional implications of neural activation patterns and their changes during adaptation

A fundamental issue for understanding the functional role of a given neuron is the causal origin of its observed activation pattern in a given task and the implications of changes in that pattern during adaptation [22–26,102–104,107,125].

From one perspective, each neuron receives synaptic inputs related to certain desired properties of a movement, such as its direction, speed or forces. Its activation pattern in a given task reflects the range of quantitative values of those parameters in those conditions. Changes in activation patterns across conditions reflect only task-dependent changes in the values of the parameters that modulate the neuron's discharge, rather than a fundamental change in the relationship of that neuron's activity to motor output.

Muscles behave in this way. Muscles generate forces and torques across joints. Each muscle's anatomy and biomechanical properties determine the relationship between its contractile activity and output forces at different muscle lengths or joint angles. As a result, an arm muscle's contractile activity level varies with the direction of reaching movements made in a given arm posture. This directionally-tuned activation pattern changes when the arm is in different postures or confronted with different external forces. Those changes in activation patterns only reflect changes in how the motor system activates the muscle in different task conditions to produce the required output

forces, not changes in the underlying relationship between the muscle's contractile activity and force output.

Alternatively, the relationship between neural discharge and intended movements could be labile. Changes in task-related activation patterns could thus result not only from changes in the values of a given set of motor parameters but also because of a more fundamental change in the relationship of a neuron to motor output. This could result from task-dependent changes in the movement parameters signaled by synaptic inputs, as well as from changes in how the neuron's output signal is processed by subsequent neural circuits. This is theoretically possible in a redundant network in which there are orders of magnitude more neurons than output variables to control.

Experimentally, the further removed a neuron is from direct sensory inputs or from the final motor output, the more difficult it is to establish a causal link between neural activity and motor behavior and thus to distinguish between these two possibilities [22,103,125]. This limits our ability to determine the origin and functional implications of adaptation-related changes in the activation patterns of cortical neurons during overt arm movements. By contrast, the BCI decoder explicitly defines the causal relationship between neural activity and effector motions, thereby making it possible to investigate directly the origin and functional implications of adaptation-related response changes.

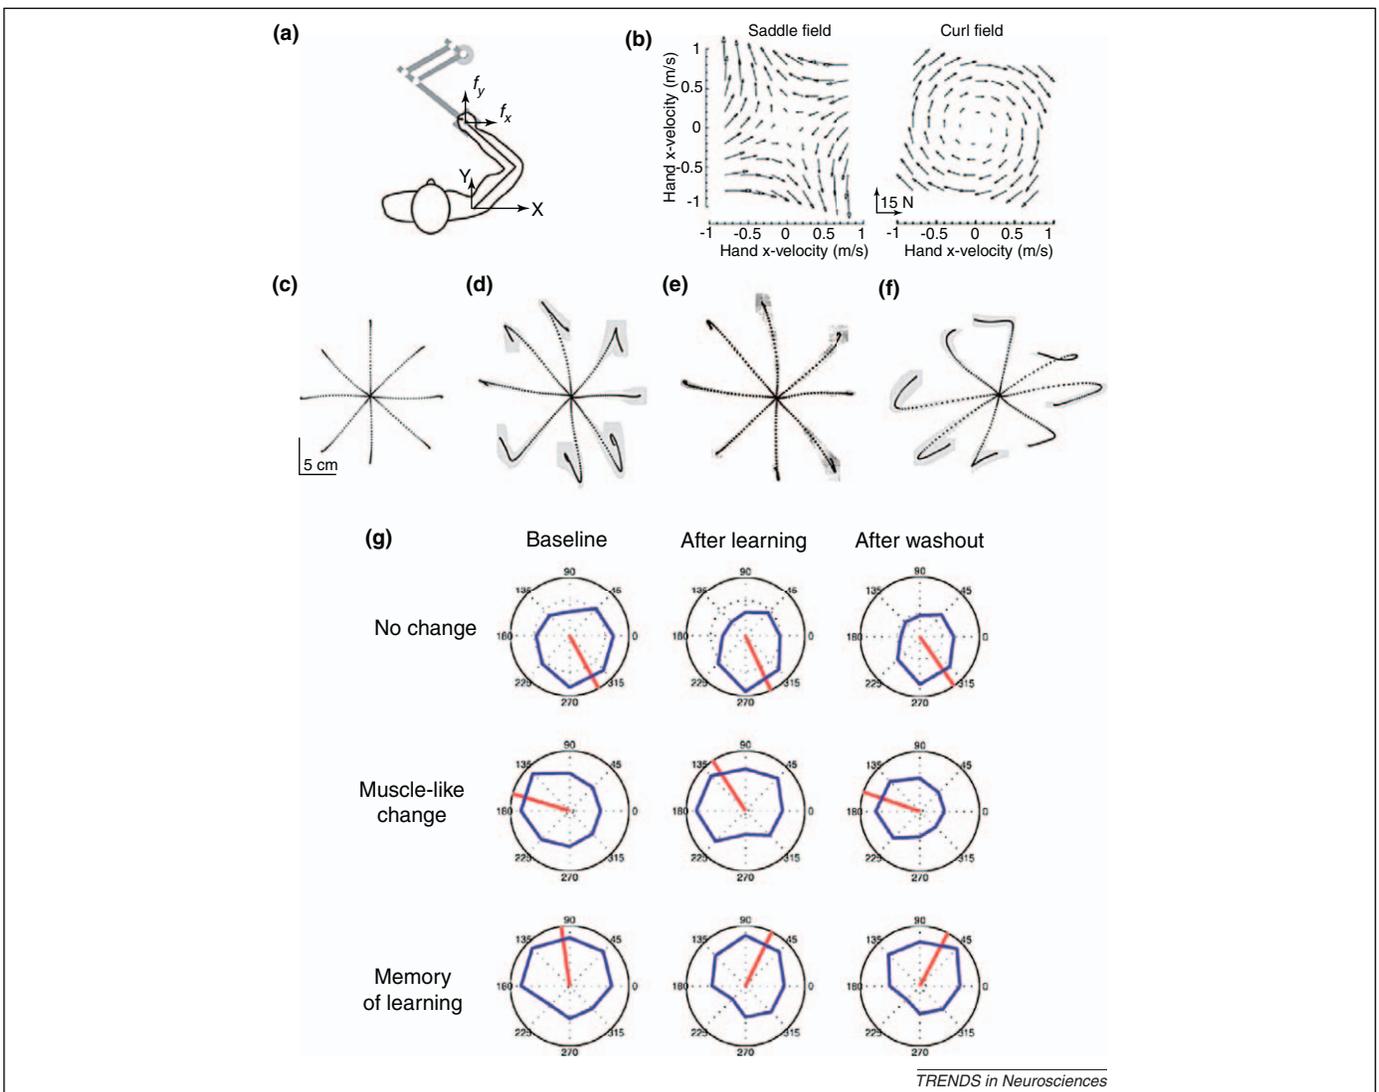


Figure 2. Behavioral [29,30] and neural [74] evidence for adaptation to a novel dynamic environment. (a) Subjects hold the handle of a robot arm and make reaching movements between targets viewed on a computer screen. The robot either imposes no force during reaching (null field) or (b) perturbs the arm trajectory by imposing forces on the hand whose direction and amplitude are related to the velocity with which the hand moves. Hand trajectories for movements in the null field (c) are typically straight but these trajectories are perturbed during initial movements made in a saddle force field (d). After continued practise in the force field the subject adapts to the novel dynamic environment and trajectories straighten again (e). After adaptation, if the saddle field is unexpectedly removed in random ‘catch’ trials (f) during a block of force-field trials, the trajectories are approximately the mirror image of the perturbed trajectories when the subject initially encounters the field (d). This indicates that the subjects have learned to generate altered patterns of muscle activity appropriate to compensate for the external force field but no longer appropriate for movements in the null-field. Neurons recorded in the primary motor cortex (M1) of monkeys while they adapt to a clockwise (CW) curl force field (g) show a range of changes in their movement direction-related activation patterns that are consistent with a role for M1 in encoding learning-related changes in motor commands. These include: (i) neurons that showed no change in activation pattern at any time during adaptation or washout (top); (ii) neurons that showed task-related changes in activation patterns similar to those of muscles, including a CW shift in the preferred direction (red line) of their activation pattern (blue lines) after adaptation to the CW curl field and then a counter-clockwise (CCW) shift back to their baseline pattern after washout (middle); and (iii) neurons that showed a CW shift of activation pattern during adaptation but did not revert back to their baseline activity during washout (bottom), or did not shift during adaptation but showed a CCW shift during washout (not shown). These latter types of neurons behaved as if they retained a memory of each learning episode. Figures reproduced, with permission, from [35] (panel a), [30] (panels b, c, d and f), [29] (panel e) and [74] (panel g).

Skill generalization is limited presumably because the set of synaptic inputs adapted during learning contribute most to movements made in directions or postures similar to those in which the learning took place, but progressively less for increasingly different directions and postures.

However, trying to learn two conflicting skills in rapid succession, such as adaptation to opposite directions of forces or visuomotor rotations, can cause interference. This results in diminished retention of the first skill (retrograde interference) and/or impeded learning of the second skill (anterograde interference) [58–64]. The interference presumably occurs because the two skills require opposite adaptive changes to similar sets of synapses, causing their

effects to cancel out. Interference decreases gradually when conflicting skills are learned in increasingly different directions or postures [29,30,65–68], presumably because of reduced overlap in the sets of synapses which adapt to learn each skill. Studies also suggest that multi-skill learning could be facilitated and interference reduced if IMs had an internal modular structure in which different subpopulations of neurons are preferentially activated and modulated in different task conditions [27,69,70]. This selection process could be determined in part by context cues that could range from motor parameters such as movement direction and limb posture to more cognitive cues such as tool identity or even arbitrary colors [50,52,64,67,71–73].

Neurophysiological studies of motor adaptation

Compared to the wealth of behavioral studies, there have been relatively few single-neuron studies of the adaptation process during reaching movements. Nevertheless, they support the link between neural response properties and motor adaptation inferred by theoretical and behavioral studies.

For instance, neurons in the primary motor cortex (M1) and premotor cortex (PM) were recorded in a few studies while monkeys adapted to a viscous–curl force field (Glossary) and then re-adapted ('washout') to baseline conditions (Figure 2g) [74–79]. Some neurons showed apparent rotations in their directional activation patterns during adaptation and washout that resembled the changes in muscle activation patterns required to compensate for the force field (Figure 2g, 'muscle-like change') [74–77,80]. This implicated those neurons in representing the altered inverse-dynamics transformation underlying curl-field adaptation. Other neurons showed little activity change during the adaptation–washout cycle (Figure 2g, 'no change'). Still others showed activation changes that, unlike muscles, occurred only during adaptation or washout, but not both (Figure 2g, 'memory of learning') [74]. This diversity of response changes suggested that single cortical neurons could reflect different aspects of the force-field adaptation process to varying degrees [74–77]. However, it is unclear whether this diversity indicates the existence of separate neural subpopulations with distinct functional roles or simply reflects local ranges of responses within a larger continuum of activation patterns.

Neuronal activity was also recorded in M1 as monkeys learned to reach in a direction rotated at an angle away from the spatial location of a visual 'target' stimulus [81,82]. Unlike during curl-field adaptation, the direction of the peak movement-related activation of M1 neurons did not rotate. However, as the monkeys learned the visuo-motor dissociation, many M1 neurons whose preferred movement direction matched the new reach direction signaled by the visual stimulus began to discharge more strongly during an instructed-delay period between visual stimulus presentation and movement onset [81,82]. This suggested that the adaptive changes primarily involved a remapping of visual input about target location onto the neurons whose preferred movement direction matched the required rotated movement [81,82].

These findings confirm that motor adaptation involves changes in the activation patterns of neurons in cortical motor areas, and that the nature of the changes can reflect differences in the sensorimotor transformations that must adapt. We now examine whether BCI systems require similar adaptive mechanisms, fundamentally different processes, or no adaptive processes at all.

Brain–computer interfaces

BCI systems comprise four main components [1–7]: (i) multi-electrode arrays to record brain activity, (ii) a 'decoder' algorithm that processes the activity to extract control signals about the presumed motor intentions of the BCI user, (iii) an effector to implement the desired motor action extracted by the decoder, and (iv) sensory

feedback about the resulting effector action that closes the control loop (Figure 1a).

Decoding the motor intentions of subjects from brain activity

A crucial component of the BCI is the decoder – the computer algorithm that uses recorded brain activity to infer what the subject wants the effector to do. Different sources of signals about brain activity can be recorded with progressively more invasive methods. These include magnetoencephalograms (MEG) and electroencephalograms (EEG) recorded on the scalp, electrocorticograms (ECoG) recorded on the dura, and local field potentials (LFP) and single- or multi-neuron spiking discharges recorded with intracortical microelectrodes. Each approach has its advantages and limitations for long-term BCI use (for details, please refer to [1–4,6,7,83,84]). Field potentials including MEG, EEG, ECoG and LFP are extracellular signals that reflect the spatially-summed postsynaptic currents and spiking discharge of large neural populations. Subjects can learn to control BCI systems with field potentials [83,84], and paralyzed subjects have used them successfully to restore computer-based communication and in other applications [2,10–13,83,84]. However, the relationship between intended motor output and field potentials remains poorly characterized [85–88] compared to that for neural spiking activity [22–26]. Because this review examines the causal neural mechanisms by which subjects use BCI systems and compares them to the mechanisms underlying overt arm movements, we will focus the rest of our discussions on BCI systems that use neural spiking activity.

Before subjects can use a BCI system to control an effector (BCI-control mode), the decoder must be calibrated to recognize the specific pattern of neural activity that signals each desired effector action. In experimental animals, the activity used for calibration is usually recorded while they perform overt arm movements to control an effector in a learned motor task (arm-control mode) [17–20,89–97]. However, paralyzed patients cannot move. Instead, activity to calibrate the decoder is evoked by asking the patients to think about moving their arm or to watch an effector as it moves under computer control and to imagine they were moving it in the same way [14–16]. Decoders have also been calibrated in monkeys using activity recorded while they watch effector motions [21,98–100]. This activity could reflect a spontaneous covert rehearsal of the arm movements the monkeys had learned to associate with the observed effector actions, because they had prior experience with overt arm control of the effectors in motor tasks [99–101].

Irrespective of how it is generated, the activity recorded during calibration is used to parameterize algorithms that describe how each neuron's discharge varies with selected motor parameters such as spatial position, direction and speed [17–20,89–97]. Ideally, the decoder once calibrated should be able to combine the movement-related information extracted from neural activity to generate an output signal about how the subject wants the effector to move.

The BCI as a novel motor-control environment

The human motor system is designed to control arm movements via spinal motoneuron activity. When subjects try to

use a BCI system, however, they face several unique challenges. The decoder and effector replace much of the motor system and their arm. The subjects must control effector motion via the activity of the cortical neurons recorded by the decoder, and not via spinal motoneuron activity (Figure 1). This could also reduce the effectiveness of many cortical and subcortical circuits that normally contribute to motor planning and to rapid corrections for errors during arm movements. Although currently under development, BCI systems cannot yet provide proprioceptive feedback of effector motions [1–4]. Subjects can only make visual feedback-mediated voluntary corrections for effector motion errors. Furthermore, the physical properties of the BCI effector are very different from the arm, the effector that the motor system has evolved to control [2,17,20,21,93].

The input–output characteristics of the decoder that generates the control signals for the effector present further challenges. Even after calibration, the decoder's output signal never perfectly replicates the overt arm movements, observed effector motions or presumed motor imagery that were used to calibrate it [14–17,19–21,95,98]. There is always a residual error in the decoder output that can cause mistranslation of the subject's motor intentions into effector actions. This error results from several further factors.

The first is a sampling bias in the input signals into the decoder. Millions of neurons throughout the motor system control different aspects of an intended movement. The BCI electrodes sample a minuscule fraction of that population at random, often in only one cortical area. The decoder's effector control improves as the neural sample size increases, but it always plateaus at a level well short of perfect replication of the subject's motor intentions [16,17,20,95,98]. The decoder might not be able to interpret those intentions fully because the many unobserved neurons that contribute to arm control could express motor-related information that is not signaled by the recorded neurons.

Another factor is the choice of the motor parameters the decoders extract from neural activity [16,90,95]. Most decoders extract a signal about the instantaneous spatial position or velocity of desired effector motion. However, although the recorded activity could be statistically correlated with those parameters, the neurons might actually be processing combinations of sensory, motor and cognitive information that do not correspond to any single definable parameter space [3,6,7,22,102–104]. Any component of the activity of even strongly task-related neurons that is not directly related to the parameters the decoder extracts could cause the effector to deviate from its intended motion.

Finally, neural activation patterns could change from decoder calibration to BCI control for reasons other than adaptation. For instance, when the decoder is calibrated during overt arm movements, the recorded activity is a combination of feed-forward signals about intended actions and sensory feedback about the actual state of the limb. However, during BCI control, proprioceptive feedback is absent if the limb does not move. Consequently, the decoder calibration established during arm movements might be inappropriate in BCI-control mode.

To what extent must subjects learn to use a BCI system? An ideal BCI system would allow users to control effector motions by generating patterns of neural activity about desired effector motions with little or no training, despite the unique control environment it presents. However, there is a wide diversity of reports of how readily subjects acquire BCI control and to what degree effector control improves with practice.

To control effector motions with MEG- or EEG-based BCI systems, subjects must impose volitional control over the amplitude of specific frequency bandwidths of their MEG/EEG spectrum by modulating the activity of large neural populations in different brain regions. Human subjects usually need extensive training before they can master this skill [2,10–13,83,84]. ECoG-based BCI control is acquired more quickly, possibly because this requires control of more localized neural populations [12].

Compared to spatially-summated field potentials, the spiking discharge patterns of single neurons provide a more temporally and spatially precise expression of a subject's motor intentions and are a rich source of signals about many parameters of intended movements [22–26]. Studies in which subjects used spike-based decoders have usually reported some degree of effector control shortly after decoder calibration [1–8,14–21,89–94]. Nevertheless, initial control was often awkward [17–21]. For instance, paralyzed human subjects can use M1 neural activity to control effectors immediately after decoder calibration [4,6,14–16]. The control was often somewhat clumsy and improvement was limited across training sessions. However, the conditions in those studies might not have been ideal for revealing the full potential for BCI control improvement because practice in each task was usually short (e.g. 40–80 trials) and repeated sessions were often separated by many days during which the subjects performed other BCI tasks [14–16].

Experimental animals likewise often appear to switch fairly seamlessly and quickly from arm-control to BCI-control of effector motions when first presented with this condition [1–8,17–21,89–94]. The significance of this seemingly rapid acquisition of BCI control is difficult to interpret in some studies [17–20,89,94] because, for a period of time after the experimenter switched from arm-control to BCI-control mode, the animals continued to produce arm movements and associated neural activity resembling those used to calibrate the decoder. Despite the ongoing arm movements, effector control often deteriorated after the initial switch. BCI control improved progressively within and across sessions, during which time the animals gradually stopped making overt limb movements [1–8,17–21,89–94]. This suggested that the animals slowly acquired the ability to control the effector by covert brain activity alone (Box 2).

Finally, in the most demanding effector-control task to date, monkeys gradually acquired BCI control of an arm-like robot to reach out to food, grasp it and bring it to their mouth [21]. Even though the monkeys had prior experience with related tasks, they were unable to impose BCI control over the robot when it was first presented. Instead, they had to be guided through a series of intermediate steps with computer assistance of robot motions over many

Box 2. How do experimental subjects stop moving their arm during BCI control?

When BCI decoders are calibrated while subjects perform overt arm movements, studies often report that the subjects gradually and spontaneously reduce or stop moving their arm after switching to BCI-control mode, although some residual motions and muscle contractions can still occur [4,17,18,20,21,89,94,105]. This phenomenon has been frequently reported but not well documented to date, nor has the possible contribution of residual movements to the BCI-related activity of recorded neurons been studied in detail. This might be insightful, because the residual movements or muscle contractions could be external signs of part of the subject's solution for BCI control.

How a dissociation between cortical activity and overt movement develops during BCI control is not understood [4]. Visual feedback about discrepancies between observed effector motions and their own intended and sensed arm movements could inform the subjects that effector motions are no longer directly coupled to their arm movements. This could be accelerated if the neural activation patterns required to produce the desired effector motions correspond to unusual or physically impossible contortions of the arm. This could also reflect a gradual recognition by the experimental animals that they are in a highly unnatural situation in which they can control the actions of an effector by generating covert brain activity about motor intentions, but without making overt arm movements or interacting

physically with the effector. By contrast, this knowledge can be quickly transmitted to human subjects by simple verbal instructions.

As a result, as the subjects acquire BCI control, they might gradually stop generating the cortical activity responsible for the descending motor commands that evoke movements and only generate central neural activity implicated in putative higher-order planning processes, similar to the activity recorded during the delay periods of instructed-delay tasks. One might expect that this transition to covert mental rehearsal of motor intentions would lead to decreased activity in neurons that are normally associated with movement execution, such as those in M1. However, although some studies have reported a net decrease in M1 activity during BCI control [17,89] others have reported the opposite trend [20,94]. Alternatively, subjects could continue to generate cortical activity that is normally associated with overt arm movements. This would require a mechanism to actively suppress the arm movements that would otherwise result, as occurs during rapid eye movement (REM) sleep [89,99,126,127]. Further understanding of how subjects stop making overt arm movements while exerting covert BCI control of effectors could also yield new insight into how subjects normally initiate voluntary movements or transition from a state of motor readiness to overt action, processes that are still not fully understood.

training sessions before they could perform this neuroprosthetic feeding task with a reasonable degree of success.

Several factors probably contribute to the diversity of findings about the speed of acquisition and degree of improvement of BCI control across studies. These include different neural sample sizes and cortical areas in which neurons were recorded [17–20,90,92,94], differences in decoder extraction algorithms [14–21,89–98], task demands [21], degrees of prior training and familiarity of experimental animals with behavioral tasks, and even differences in the criteria used to assess performance. Another key issue in animal studies is the degree to which improved BCI control reflects adaptation to the properties of the BCI system versus non-motor processes associated with the animals' progressive recognition that they are in a

highly unnatural situation in which they can cause and control movements of the effector, but without the need to interact physically with it or even make overt movements.

Practice-related improvements in BCI performance display several properties in common with adaptation of overt reaching movements. For example, improvements are progressive and incremental within and across daily training sessions [14–21,89,90,92,93,105] (Figure 3b). This appears to reflect adaptation to both the input–output properties of the decoder and to the physical properties of the effectors [17,20,21,93]. Subjects show limited generalization of newly-acquired BCI skills to untrained conditions. After practising cursor movements to one set of targets, subjects can then move the cursor to new targets in orthogonal directions [20]. However, improved performance after practice

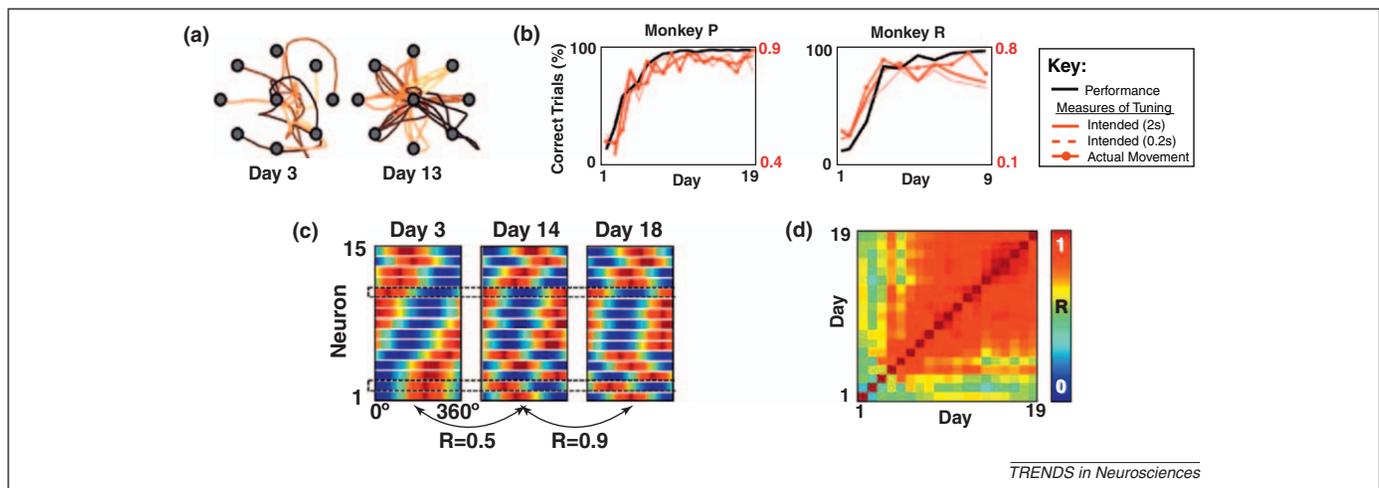


Figure 3. Behavioral performance improvements and neural response changes observed in monkeys during BCI control [105]. (a) Example of single-trial cursor trajectories during initial (left) and late (right) stages of BCI control learning. Early in training (day 3), cursor paths were highly variable, and the subject often failed to displace the cursor to the targets (targets are depicted as grey circles; the failed trials are not shown). After further practice (day 13), cursor paths were significantly straighter, aimed more directly at the targets, and most trials were successful. (b) The success rate across daily training sessions (solid black) closely resembles the rate of change in ensemble neural tuning in M1 (red curves) for two monkeys (monkey P, left panel, and monkey R, right panel). Red curves show ensemble directional tuning calculated either relative to the intended target locations (2 s window, solid; 0.2 s window, dotted) or relative to actual cursor movement directions (solid with dots). (c) Color maps illustrating the shift in the ensemble of activation patterns of 15 stable neurons recorded in one monkey across 18 successive daily training sessions. Red colors indicate the peak of the directional activation pattern for each neuron. As the monkey became more skilled at BCI control, the activation patterns of most neurons changed during BCI-control mode. The ensemble of BCI-control mode activation patterns became more constant as training progressed, increasing from a correlation value (R) of 0.5 between days 3 and 14, to 0.9 for days 14 and 18. (d) Pairwise correlations (R) of ensemble tuning across all practice days. As training progressed, the ensemble of neural activation patterns during BCI-control mode became increasingly similar from day to day. Figures reproduced, with permission, from [105].

with a decoder that controls cursor position does not result in improved control with a cursor-velocity decoder [17].

BCI control in experimental animals is also an example of multi-skill learning. The animals spend most of the day outside of the lab using their arm in normal behavior. In the lab, they become increasingly adept at switching between arm-control and BCI-control modes, with progressively less retrograde and anterograde interference at each transition [17,20,89,105]. This could reflect both improved control of the activity of the neurons recorded by the decoder, as well as greater familiarity of the animals with the unnatural context of covert BCI control of effectors via cortical neural activity without overt arm movement.

Neural response changes during BCI control

Studies in experimental animals have reported widely diverse changes in the activation patterns of neurons between arm-control and BCI-control modes. Although not always well documented, the time-course of response changes reportedly parallels the improvement in BCI control, suggesting that subjects must activate the neurons in new ways to produce the desired effector actions [16,17,20,89,90,105,106] (Figures 3,4). This includes changes in task-related variance of activity [94] and in mean discharge rate and depth of modulation of activity, to the extreme that some neurons eventually discharge only during arm-control or BCI-control mode but not both [16,17,89] (Figure 4).

Some of the most striking changes involve the directional tuning of neural activation patterns [14–17,19–21,89,90,92,94,105]. The reported changes vary substantially across studies, ranging from uncorrelated shifts between neurons [20] (Figure 4d) to systematic rotations across populations [106], or even convergence of neural activation patterns onto similar directional preferences [17,89] (Figure 4f). This perplexing diversity of results could reflect many factors, including differences in tasks, size and location of recorded neural populations, decoder algorithms and the motion parameters they extract, and even the methods used to define directional tuning [14–23,89,90] and the time range over which the neural activity is sampled [104,107].

This lack of consensus about response changes between arm-control and BCI-control modes must be resolved before we can gain more insight into the neural mechanisms by which subjects use BCI systems. Why do some studies report seemingly random directionality changes [16,20,98] whereas others find highly correlated convergent directional changes across neurons [17,89]? How can such disparate changes both apparently result in improved BCI-control performance? Do they reflect similar or different adaptive processes, or even different cognitive strategies to impose volitional BCI control over effector motions? The answers to these questions will help to clarify the extent to which the activity changes are causal for BCI performance or are merely secondary consequences of BCI control.

The BCI decoder provides a unique tool to study motor control and adaptation

The neural mechanisms underlying the control and adaptation of overt arm movements are also not fully understood, in part because the causal relationship between cortical neural activity and the resulting movement is

almost always indirect and difficult to establish [5,22,23,102,103]. By contrast, BCI decoders exactly define the relationship between the activity of recorded neurons and effector motions, and permit direct experimental manipulation of that relationship. Two recent studies [105,106] provide striking demonstrations of the potential to exploit these features of BCI decoders to study motor adaptation.

BCI studies typically recalibrate the decoder at the start of each session because they assume that the population of recorded neurons will tend to change over time due to electrode isolation instability [1–7,14–21]. This confounds long-term studies of skill acquisition. To avoid this problem, a recent study [105] selected a small subset of 10–15 M1 neurons that could be stably recorded over many days, and used only their activity to calibrate a decoder while monkeys made overt arm movements to control cursor motions on a monitor. The decoder parameters were then fixed for the rest of the multi-day experiment. When the task was switched to BCI-control mode, performance was initially poor, in part because of the fairly demanding nature of the task [105]. However, it improved progressively within and across the first three to eight daily sessions with the fixed decoder (Figure 3a,b). In parallel, differences in the neural activation patterns between arm-control and BCI-control modes increased gradually over the first sessions and stabilized once BCI control reached a plateau, until the end of the experiment (9–19 days) [105] (Figure 3c,d). At that point, the monkeys could switch between arm-control and BCI-control modes in one to two trials with no interference.

To further study multi-skill learning, monkeys were presented with both the original decoder, with which they now had extensive practise, and a new decoder in which there were significant changes in the mapping between neural activity and effector motions [105]. The two decoders were presented in alternating blocks of trials signaled by different color-context cues on the computer monitor. The animals became increasingly skilled at effector control with the new decoder across days of practice while retaining good control with the original decoder (Figure 5). In parallel, the monkeys developed a new set of activation patterns of the same stably-isolated neurons when using the new decoder that differed both from that for the original decoder and that during overt arm control. Those 'prosthetic motor maps' [105] became stable, readily recalled, and resistant to interference from each other when the subjects switched between decoders or between arm-control and BCI-control modes. The emergence of distinct stable prosthetic motor maps during decoder adaptation might be one neural correlate of the acquisition and consolidation of multiple motor skills [27,69,70]. Crucially, the stable maps only emerged if the parameters in each decoder and the set of recorded neurons remained fixed, allowing the subjects to retain the maximum amount of acquired skill across days.

These findings in carefully controlled conditions provide strong evidence that BCI control can improve with practice and involves changes in neural activation patterns. They also suggest that long-term BCI users could become skilled at generating distinct sets of activation patterns of the

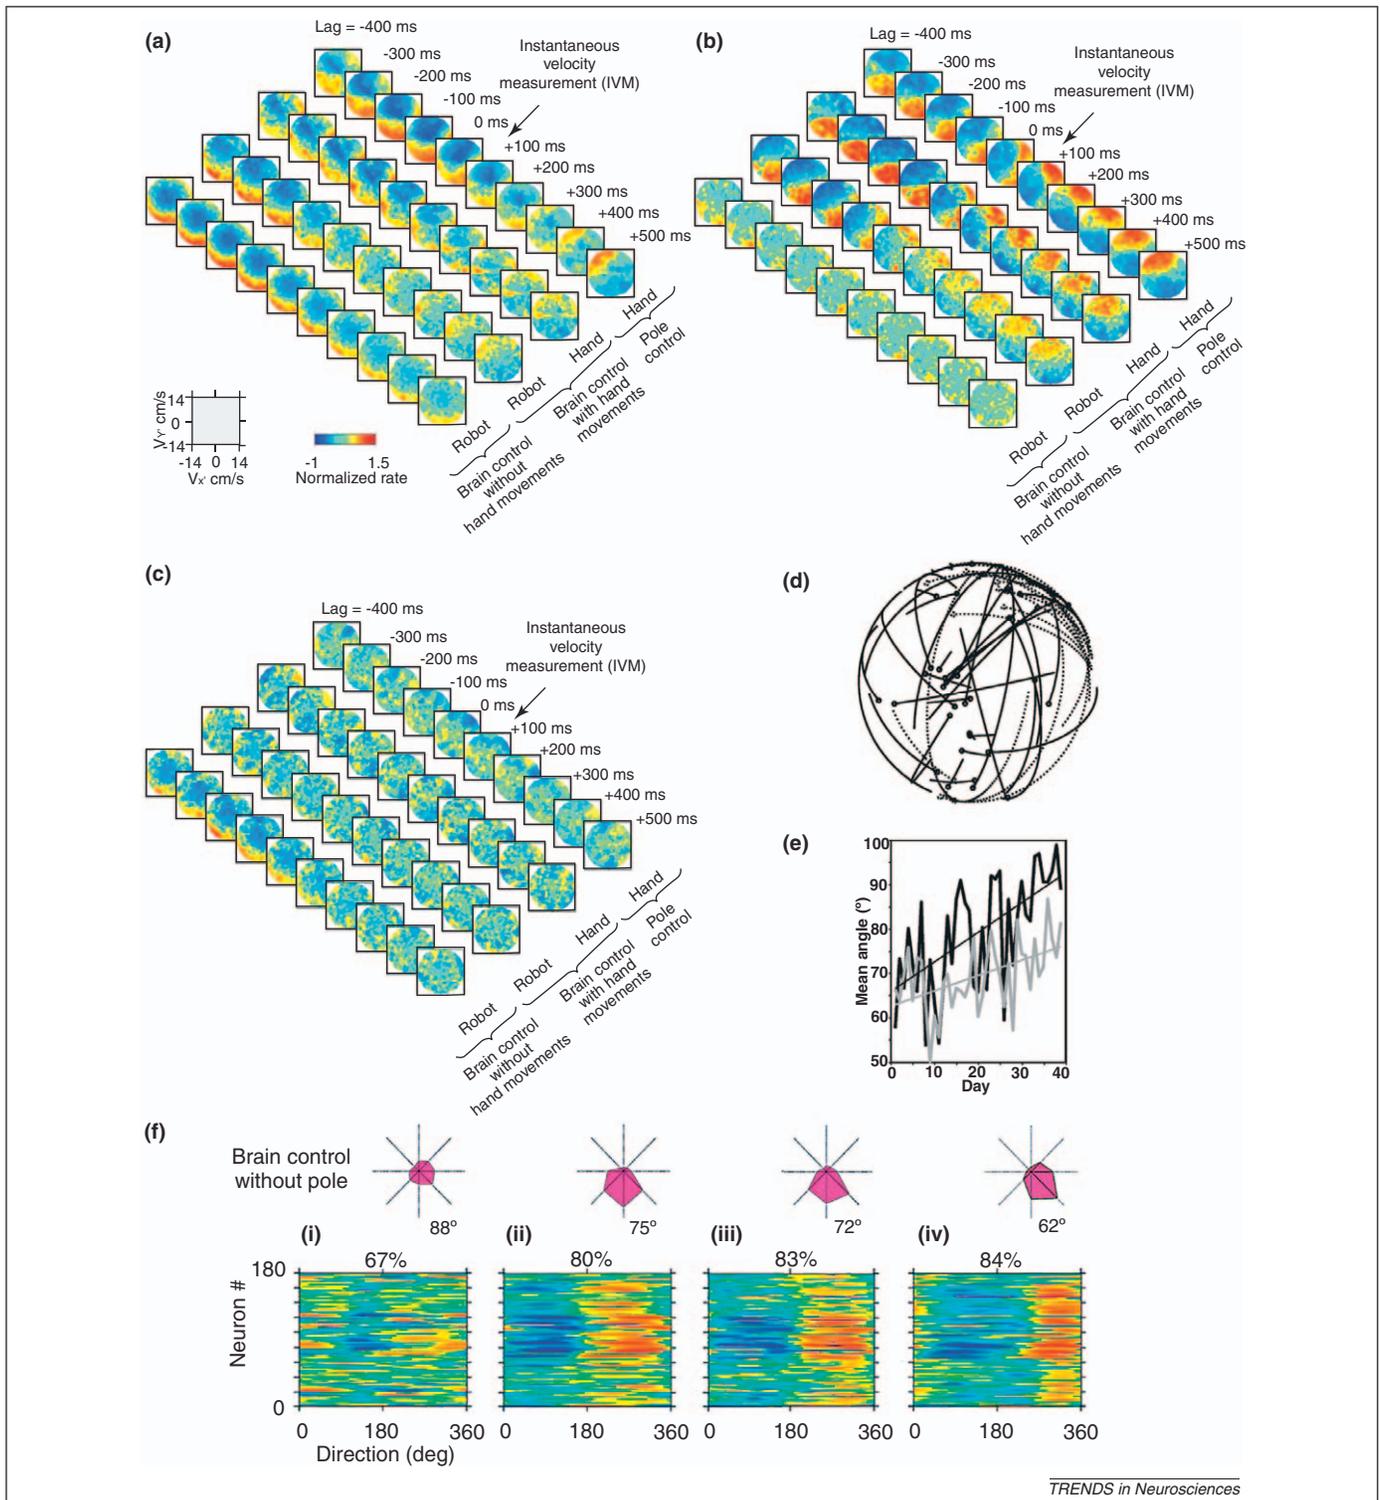


Figure 4. Neural responses exhibit a broad range of directional tuning changes during brain- versus arm-control. (a,b,c) Examples of the activities of three M1 cells recorded in a monkey when movement of a robotic arm was controlled using either arm ('pole') control, brain control accompanied by arm movements, or by pure brain control [89]. Polar plots illustrate neural activation patterns as a function of direction and magnitude of velocity with warm colors representing higher activity. Neural activity is plotted at different times before (negative lags) and after (positive lags) the time at which an instantaneous velocity measurement (IVM) was made. Tuning during brain control with arm movements was calculated relative to either robot (cursor) or hand movement velocity. (a) A neuron whose activity preceding the current IVM shows a fairly constant activation pattern between arm-control and brain-control modes but whose activity recorded after the current IVM shows a shift in activation pattern across control modes. (b) A neuron that was preferentially activated during arm-control but not brain-control mode. (c) A neuron that became preferentially modulated during brain-control compared to arm-control mode. (d) Illustration of large and random variability in the changes in directional tuning of the activation patterns of neurons recorded from a monkey during arm versus brain control [20]. Each line denotes the shift in the preferred direction of activation for an M1 neuron during arm-control (unmarked line end) versus brain-control mode (circles). Activation changes are projected onto a unit sphere. (e) Change across daily training sessions in the mean difference in preferred direction during arm- versus brain-control (mean angle; thick lines) for all M1 cells significantly tuned during both arm- and brain-control modes [20]. Black, contralateral; grey, ipsilateral M1 relative to the arm that moved during the hand-control task. (f) Distributions of directional tuning curves of a population of monkey motor cortex neurons at different stages in the acquisition of BCI control without arm movements [17]. Each color-coded horizontal line represents the tuning curve of a single neuron, with its preferred movement direction indicated by warm colors. Initially, directional preferences were widely distributed across the population (i) but as the animal continued to practice BCI control and performance improved (ii-iv) the tuning functions of the neurons shifted to converge into the quadrant between 270-360°. Panels a-c reproduced with permission from [89]. Panels d-e from [20]; reproduced with permission from the American Association for the Advancement of Science. Panel f reproduced with permission from [17].

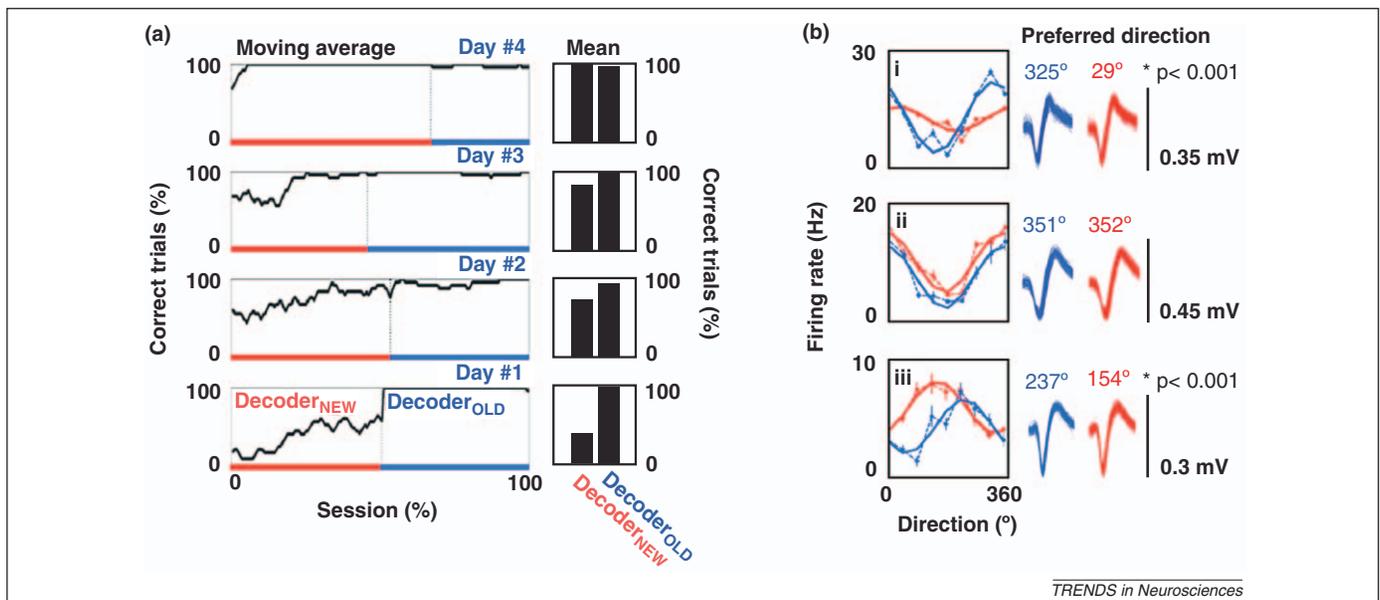


Figure 5. Acquisition and simultaneous retention of two BCI-control ‘prosthetic motor maps’ for two different neural decoders [105]. **(a)** Changes in performance of a monkey were measured by the percentage of trials in which the monkey was able to control cursor motion well enough to move it to a target within 10 s. Performance changes are illustrated over four days when a new decoder (Decoder_{NEW}) was introduced on day one after BCI control had already been learned over the course of 19 days with Decoder_{OLD}. By the fourth day, performance with Decoder_{NEW} was as good as that with Decoder_{OLD}, illustrating the ability of the animal to successfully use – and switch quickly between – two distinct decoders. Left, moving average of performance during each session; right, mean session performance. **(b)** Directional activation patterns (left) and neural spike waveforms (right) from three different stably-isolated neurons (i–iii) while performing brain control of cursor motions with the two decoders (Decoder_{NEW}, red; Decoder_{OLD}, blue). Directional activation patterns during BCI control often changed when switching between the two decoders, illustrating that each decoder invoked a unique ensemble of directional activation patterns (‘prosthetic motor maps’). Figure reproduced with permission from [105].

same neurons that would allow them to switch readily between motor tasks and to use different effectors that require unique decoders. This is crucial for the flexibility and ease of use of BCI systems in a clinical setting.

Another recent study also used decoders as tools to study adaptation by creating perturbations at the single-neuron level [106]. Monkeys first learned to use M1 activity to move a cursor to targets in a virtual 3D visual space in BCI-control mode. They were then confronted with a novel visuomotor dissociation created by choosing a random subset of the recorded neurons and rotating the decoder’s mapping of their directional activation patterns onto 3D cursor motions by 90°. When the monkeys first used the altered decoder the cursor moved at an angle to the intended direction. However, after many trials of practice the monkeys could compensate partially for the rotation so that the cursor moved more directly towards the target. As adaptation progressed, neural activity showed a gradual shift from delayed feedback-mediated responses (to correct for observed cursor motion errors) to more predictive feed-forward response changes before cursor motion onset.

The directionality and depth of modulation of single-neuron activity changed after practice with the altered decoder [106]. There was a small net counter-rotation of the directional activation patterns of the entire population, consistent with a global strategy of re-aiming the intended cursor motion towards a virtual target in the direction opposite to the perturbed cursor motion. However, the most novel findings were that the degree of counter-rotation was larger and the depth of discharge modulation was reduced for those neurons whose decoder mapping had been altered, compared to the non-rotated neurons. Those differential trends suggested a credit assignment process that identified and selectively altered how the motor system activated the

neurons whose decoder mapping had been rotated [106]. It preferentially enhanced the compensatory counter-rotation of their activation pattern (re-mapping) while simultaneously reducing their perturbing impact on cursor motions by decreasing their overall activation (re-weighting).

One limitation of this study [106] was that the monkeys only compensated for about 25% of the applied rotation during each session, and the neural changes were correspondingly small. Nevertheless, the insights provided by both of these studies [105,106] were possible because the decoder provided exact knowledge about the effect of each neuron’s activity on cursor motions before, during, and after adaptation. Thus, this approach shows great promise for studying neural mechanisms of motor learning.

Some other unresolved issues

Several issues must still be resolved before we can understand how subjects impose BCI control of effectors. By providing the ability to define and control the relationship between neural activity and motor output, BCI technology provides unique tools to examine these issues and to provide novel insights into motor-skill acquisition during overt arm movements.

Which observed changes in neural activation patterns are adaptive and learned?

One cannot assume that all observed changes in neural activation patterns during BCI use contribute equally to BCI control or reflect learning processes. For instance, some BCI studies used adaptive decoders whose parameters were updated periodically [16,19,20,108,109]. To what degree BCI control in those studies resulted from changes in neural activity versus changes in how the decoder transformed that activity into effector actions is unclear.

Furthermore, some decoders assign both tuning parameters and a weight to each neuron, and these determine how much its activity contributes to effector motion [17,90,95,98]. As a result, small activity changes in a heavily-weighted neuron could make a larger contribution to BCI control than larger changes in a weakly-weighted neuron, and some changes might even be detrimental [90,95,98]. The extent to which a given neuron's activity change contributes to changes in effector control is readily available from the decoder but, with a few exceptions [98,105,106], studies have largely not exploited this information.

Reported activity changes could also result from undetected changes in the population of recorded neurons across time because of isolation instability, or a time-dependent stochastic drift in neural response properties unrelated to any learning process [3,108–112]. Changes in the cognitive state of the subject – such as a transition from passive observation or covert motor imagery during initial decoder calibration to attempts at volitional BCI control of effector motions – could also alter neural responses [5,6].

Finally, activation patterns could change as a result of a decrease or loss of proprioceptive feedback about arm posture and movement after switching from arm- to BCI-control mode. Intriguingly, the loss of joint/muscle-centred proprioceptive signals might actually facilitate BCI control. Centrally-generated motor activity might provide information primarily about extrinsic spatial parameters of intended effector actions [113–115] that is more compatible with the motor-output parameters extracted by most decoders.

How do subjects adapt to a BCI system?

An issue that has rarely been examined in detail [106] is the degree to which BCI-control improvement results from adaptive changes in feed-forward processes versus more efficient use of visual feedback. Equally unclear is whether the motor system treats BCI-control mode primarily as a change in the kinematic transformation from visual inputs to desired motions, or treats the decoder and effector as a surrogate arm or a novel environment with new dynamic properties that require a new dynamics transformation between desired effector actions and motor commands. The apparent rotations of activation patterns of many neurons between arm-control and BCI-control modes [14–17,19–21,89,90,92,94] are reminiscent of the rotations seen during adaptation of arm movements to external curl-force fields [74–77,79,112] but not during adaptation to visuomotor rotations [81,82]. This could indicate that the motor system's response to the BCI system more closely resembles that due to a change in the physical properties of the arm or the environment than to a change in the mapping between visual input and effector motions. This could occur because, during BCI control, the activity of the recorded neurons is directly causal to effector motions via the decoder, in the same way as forces, muscle contractions and associated neural activity are causal to overt arm movements. However, this highly speculative interpretation presupposes that the motor system implements BCI control by performing the same basic computations as during overt arm control [1–8,14–21]. Other very different processes might also contribute, including mechanisms that

permit direct volitional control of neural activity [5,8,116,117].

A related question is to what degree BCI control involves volitional processes to seek a solution versus more subconscious error-driven mechanisms [106]. Subjects learning operant control of the activity of single M1 neurons often start by searching through their natural motor repertoire to find an action that produces the required neural activity [8,116,117]. Subjects could use a similar approach when they first switch from arm-control to BCI-control mode. Aided by visual feedback about effector motion errors, subjects could generate diverse patterns of neural activity associated with different arm movements or combinations of muscle contractions to find one which results in the desired effector motions, not unlike a golfer who contorts his body to try to will the flight of a badly slicing ball back onto the fairway.

However, the incremental nature of improved BCI control also implicates error-driven adaptive processes such as those associated with adaptation of overt arm movements. This might involve a credit-assignment process to determine the nature and source of the performance errors and to identify the neural populations whose activation patterns should be altered [50,52,106]. Different patterns of activity changes could also indicate the existence of distinct functional subpopulations with specific roles in adaptation that might be subject to different adaptive mechanisms [74–77].

Alternatively, however, recent modeling studies suggest that the diverse adaptation-related activation changes could result from a single learning mechanism applied uniformly across a highly redundant neural population [112,118]. Whenever a change is encountered, such as a force field or a BCI system, the learning mechanism drives the neural population towards a new functional state that is expressed as a new task-dependent set of neural activation patterns, by progressive adjustment of synaptic connection strengths. Different patterns of activation changes suggestive of different functional subclasses of neurons [74–77] or different adaptive processes [106] could simply be epiphenomena that emerge from the unique history of stochastic trial-to-trial variability of each neuron's discharge and error-driven versus noise-driven changes in synaptic weights during adaptation [112,118]. These intriguing ideas await further experimental verification. Even the degree of inherent stability of the M1 motor map while interacting with a stable world is still controversial [108–112]. BCI technology offers a unique experimental tool to examine these important issues.

What is the contribution of all the neurons not used by the BCI decoder?

When the activity of a local population of neurons in M1 [14–17,19–21,89,90,94,105] or other cortical areas [17,89,90,92,94,96] is used by a decoder to control an effector, their activation patterns often appear to change. Normally, however, neurons throughout the distributed cortical and subcortical motor system all contribute to the control of overt arm movements [27,39,69,70,74,75,77,81,119,120]. To what degree do the activation patterns of the many neurons not used by the decoder also change and contribute indirectly to BCI control?

For example, visual feedback is crucial for BCI control. Preliminary evidence shows that when humans begin to use BCI systems, M1 neural activity used by a decoder signals both the intended effector motion and the current visual error between desired and actual motion [14–16]. However, M1 does not receive direct visual inputs. The visual input must be processed elsewhere, including the parietal cortex, which influences M1 activity via cortico-cortical projections [24,25,91,121,122]. This suggests that visual error-driven activation changes could also occur in the parietal cortex [39,91] and contribute to the altered activation patterns of the M1 neurons used by the decoder. Similarly, concurrent activation changes in the premotor cortex and cerebellum could contribute to improved control of effector spatial motion [113–115] and to adaptation to its dynamic properties [27,69,70], respectively.

BCI technology could be used to study this issue. For example, electrode arrays could be implanted in M1 and other cortical and subcortical structures. Neurons in only one structure such as M1 would be used for BCI control. Neural activity could then be recorded in all the structures while subjects practice BCI control of a variety of effectors with different kinematic or dynamic properties. Studying whether and how the activation patterns of neurons not directly used by the decoder change could provide unique insights into the contribution of each structure to BCI control, as well as to overt motor control.

Can BCI control ever be as good as overt arm control?

No BCI user to date has consistently achieved the level of precision, speed, and flexibility of effector control seen during overt arm movements. The reduced nature of the BCI neural control circuitry, neural sampling biases, non-optimal decoder algorithms, frequent decoder recalibration and other factors could all impose limits on the level of skill that can be acquired and retained with current BCI technology.

Furthermore, most decoders generate a continuous time-varying output signal to control the moment-to-moment details of the desired effector actions. However, centrally-generated feed-forward signals about intended arm or effector movements could normally provide only fairly general information about the desired movement such as its overall direction and endpoint [96,123]. The signals that control the details of the trajectory might emerge primarily during the movement through interactions between these feed-forward processes and feedback about current limb state, and could involve neural populations that are synaptically downstream of the recorded neurons.

Finally, the recorded neurons are embedded in a cortical network whose functional architecture evolved to control the subject's arm. This synaptic architecture could impose constraints that limit the ability of each neuron to assume any random activation pattern and so limit the ability of subjects to acquire the perfect BCI prosthetic motor map.

Nevertheless, a recent preliminary report described BCI control that approached the speed and accuracy of able-bodied subjects after only minutes of practice by using a novel correction algorithm to identify errors in decoder output signals and redefine how the decoder mapped neural activation patterns into effector motions [124]. This gives

reason to be optimistic that BCI technology can be refined to the point where subjects can use it efficiently to produce well-controlled effector motions for extended periods of time with minimal concentration, fatigue or frustration.

Do human subjects and experimental animals acquire BCI control differently?

Experimental animals gradually acquire covert BCI-control of remote effectors by trial and error, usually after extensive training in similar tasks using overt arm-control. In contrast, the brain activity needed to calibrate a decoder for paralyzed human subjects can be generated simply by asking them to observe and imagine making effector movements. They can then be asked to use similar motor imagery to impose volitional BCI-control on the effector's motions. In theory, the subjects should be able to immediately control the effector by reproducing the same neural activation patterns during BCI-control as during decoder calibration, without any learning. Pilot clinical studies have indeed shown that human subjects can exert a reasonable degree of effector control right after decoder calibration [14–16]. However, the ability to consistently recall and replicate the same sets of neural activation patterns is itself a skill that should improve with practice. Furthermore, the initial imagery-evoked activity might not be adequate to produce the desired effector motions, for many of the reasons discussed here, so human subjects may have to learn how to use visual feedback efficiently to improve closed-loop BCI-control. More studies are required to understand whether and how human subjects might have to learn efficient BCI-control, how this might differ from experimental animals, and what impact this has on causal neural mechanisms.

Conclusions

The role of the motor system is to accomplish behavioral goals of benefit to the individual. The motor system normally achieves these goals by controlling the skeletomuscular system, including the arm and hand. During adaptation of arm movements to visuomotor dissociations or force fields, the nature of the motor output needed to achieve the goal changes, and this leads to changes in the neural signals generated in cerebral cortical motor areas. By contrast, the means of implementing the desired goal also changes when subjects use a BCI system to control a remote effector. The challenges confronting a subject when they begin to use a BCI system are complex and multiple. They include neural sampling biases, potentially suboptimal decoder algorithms, differences in the dynamic properties of the arm versus effectors, alteration or loss of sensory feedback signals, changes in the nature of the sensorimotor transformations required to generate appropriate motor outputs, and even changes in the definition of the desired goal state when using a remote effector instead of the arm. Further studies with BCI technology will provide a deeper understanding of how the motor system responds to these challenges, leading to improvements in BCI design and performance. At the same time, BCI technology offers unique opportunities to study the general principles and neural mechanisms underlying skill learning during overt arm movements.

Acknowledgments

The preparation of this manuscript was supported by a Canadian Institutes of Health Research operating grant (J.K. and A.G.), a salary support grant from the Fonds de la Recherche en Santé du Québec (FRSQ) (A.G.) and an infrastructure grant from the FRSQ.

References

- Schwartz, A.B. *et al.* (2006) Brain-controlled interfaces: movement restoration with neural prosthetics. *Neuron* 52, 205–220
- Wolpaw, J.R. (2007) Brain–computer interfaces as new brain output pathways. *J. Physiol.* 579, 613–619
- Nicolelis, M.A. and Lebedev, M.A. (2009) Principles of neural ensemble physiology underlying the operation of brain–machine interfaces. *Nat. Rev. Neurosci.* 10, 530–540
- Hatsopoulos, N.G. and Donoghue, J.P. (2009) The science of neural interface systems. *Annu. Rev. Neurosci.* 32, 249–266
- Fetz, E.E. (2007) Volitional control of neural activity: implications for brain–computer interfaces. *J. Physiol.* 579, 571–579
- Donoghue, J.P. (2008) Bridging the brain to the world: a perspective on neural interface systems. *Neuron* 60, 511–521
- Andersen, R.A. *et al.* (2010) Cognitive neural prosthetics. *Annu. Rev. Psychol.* 61, 169–190 C161–163
- Moritz, C.T. *et al.* (2008) Direct control of paralysed muscles by cortical neurons. *Nature* 456, 639–642
- Pohlmeier, E.A. *et al.* (2009) Toward the restoration of hand use to a paralyzed monkey: brain-controlled functional electrical stimulation of forearm muscles. *PLoS One* 4, e5924
- Kennedy, P.R. and Bakay, R.A. (1998) Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport* 9, 1707–1711
- Wolpaw, J.R. and McFarland, D.J. (2004) Control of a two-dimensional movement signal by a noninvasive brain–computer interface in humans. *Proc. Natl. Acad. Sci. U. S. A.* 101, 17849–17854
- Schalk, G. *et al.* (2008) Two-dimensional movement control using electrocorticographic signals in humans. *J. Neural. Eng.* 5, 75–84
- Buch, E. *et al.* (2008) Think to move: a neuromagnetic brain–computer interface (BCI) system for chronic stroke. *Stroke* 39, 910–917
- Hochberg, L.R. *et al.* (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442, 164–171
- Truccolo, W. *et al.* (2008) Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *J. Neurosci.* 28, 1163–1178
- Kim, S.P. *et al.* (2008) Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J. Neural. Eng.* 5, 455–476
- Carmena, J.M. *et al.* (2003) Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS Biol.* 1, E42
- Chapin, J.K. *et al.* (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* 2, 664–670
- Serruya, M.D. *et al.* (2002) Instant neural control of a movement signal. *Nature* 416, 141–142
- Taylor, D.M. *et al.* (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296, 1829–1832
- Velliste, M. *et al.* (2008) Cortical control of a prosthetic arm for self-feeding. *Nature* 453, 1098–1101
- Kalaska, J.F. (2009) From intention to action: motor cortex and the control of reaching movements. *Adv. Exp. Med. Biol.* 629, 139–178
- Kalaska, J.F. *et al.* (1997) Cortical control of reaching movements. *Curr. Opin. Neurobiol.* 7, 849–859
- Rizzolatti, G. and Luppino, G. (2001) The cortical motor system. *Neuron* 31, 889–901
- Buneo, C.A. and Andersen, R.A. (2006) The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia* 44, 2594–2606
- Takei, S. *et al.* (2003) Sensorimotor transformations in cortical motor areas. *Neurosci. Res.* 46, 1–10
- Wolpert, D.M. and Kawato, M. (1998) Multiple paired forward and inverse models for motor control. *Neural. Netw.* 11, 1317–1329
- Kawato, M. (1999) Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9, 718–727
- 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. (2004) Generalization as a behavioral window to the neural mechanisms of learning internal models. *Hum. Mov. Sci.* 23, 543–568
- Zipsper, D. and Andersen, R.A. (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331, 679–684
- Pouget, A. and Snyder, L.H. (2000) Computational approaches to sensorimotor transformations. *Nat. Neurosci.* 3 (Suppl), 1192–1198
- Poggio, T. and Bizzi, E. (2004) Generalization in vision and motor control. *Nature* 431, 768–774
- Salinas, E. and Abbott, L.F. (1995) Transfer of coded information from sensory to motor networks. *J. Neurosci.* 15, 6461–6474
- Donchin, O. *et al.* (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
- Thoroughman, K.A. and Shadmehr, R. (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407, 742–747
- Lackner, J.R. and DiZio, P. (2005) Motor control and learning in altered dynamic environments. *Curr. Opin. Neurobiol.* 15, 653–659
- Hinder, M.R. and Milner, T.E. (2007) Rapid adaptation to scaled changes of the mechanical environment. *J. Neurophysiol.* 98, 3072–3080
- Clower, D.M. *et al.* (1996) Role of posterior parietal cortex in the recalibration of visually guided reaching. *Nature* 383, 618–621
- Kitazawa, S. *et al.* (1997) Prism adaptation of reaching movements: specificity for the velocity of reaching. *J. Neurosci.* 17, 1481–1492
- Krakauer, J.W. *et al.* (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* 20, 8916–8924
- Wang, J. and Sainburg, R.L. (2005) Adaptation to visuomotor rotations remaps movement vectors, not final positions. *J. Neurosci.* 25, 4024–4030
- Mazzoni, P. and Krakauer, J.W. (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J. Neurosci.* 26, 3642–3645
- Scheidt, R.A. *et al.* (2001) Learning to move amid uncertainty. *J. Neurophysiol.* 86, 971–985
- Smith, M.A. *et al.* (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 4, e179
- Kording, K.P. *et al.* (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* 10, 779–786
- Fine, M.S. and Thoroughman, K.A. (2007) Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *J. Neurophysiol.* 98, 1392–1404
- Izawa, J. *et al.* (2008) Motor adaptation as a process of reoptimization. *J. Neurosci.* 28, 2883–2891
- Berniker, M. and Kording, K. (2008) Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci.* 11, 1454–1461
- Kluzik, J. *et al.* (2008) Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J. Neurophysiol.* 100, 1455–1464
- Wei, K. and Kording, K. (2009) Relevance of error: what drives motor adaptation? *J. Neurophysiol.* 101, 655–664
- Cothros, N. *et al.* (2006) Are there distinct neural representations of object and limb dynamics? *Exp. Brain Res.* 173, 689–697
- van Beers, R.J. (2009) Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63, 406–417
- Ghez, C. *et al.* (2007) Different learned coordinate frames for planning trajectories and final positions in reaching. *J. Neurophysiol.* 98, 3614–3626
- Malfait, N. *et al.* (2005) Generalization of motor learning based on multiple field exposures and local adaptation. *J. Neurophysiol.* 93, 3327–3338
- Malfait, N. *et al.* (2002) Transfer of motor learning across arm configurations. *J. Neurosci.* 22, 9656–9660
- Shadmehr, R. and Moussavi, Z.M. (2000) Spatial generalization from learning dynamics of reaching movements. *J. Neurosci.* 20, 7807–7815

- 58 Brashers-Krug, T. *et al.* (1996) Consolidation in human motor memory. *Nature* 382, 252–255
- 59 Shadmehr, R. and Brashers-Krug, T. (1997) Functional stages in the formation of human long-term motor memory. *J. Neurosci.* 17, 409–419
- 60 Caithness, G. *et al.* (2004) Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J. Neurosci.* 24, 8662–8671
- 61 Hinder, M.R. *et al.* (2007) The interference effects of non-rotated versus counter-rotated trials in visuomotor adaptation. *Exp. Brain Res.* 180, 629–640
- 62 Krakauer, J.W. *et al.* (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *J. Neurosci.* 25, 473–478
- 63 Miall, R.C. *et al.* (2004) Adaptation to rotated visual feedback: a re-examination of motor interference. *Exp. Brain Res.* 154, 201–210
- 64 Imamizu, H. *et al.* (2007) Explicit contextual information selectively contributes to predictive switching of internal models. *Exp. Brain Res.* 181, 395–408
- 65 Hwang, E.J. *et al.* (2003) A gain-field encoding of limb position and velocity in the internal model of arm dynamics. *PLoS Biol.* 1, E25
- 66 Hwang, E.J. and Shadmehr, R. (2005) Internal models of limb dynamics and the encoding of limb state. *J. Neural. Eng.* 2, S266–278
- 67 Gandolfo, F. *et al.* (1996) Motor learning by field approximation. *Proc. Natl. Acad. Sci. U. S. A.* 93, 3843–3846
- 68 Woolley, D.G. *et al.* (2007) Dual adaptation to two opposing visuomotor rotations when each is associated with different regions of workspace. *Exp. Brain Res.* 179, 155–165
- 69 Haruno, M. *et al.* (2001) Mosaic model for sensorimotor learning and control. *Neural. Comput.* 13, 2201–2220
- 70 Imamizu, H. *et al.* (2004) Functional magnetic resonance imaging examination of two modular architectures for switching multiple internal models. *J. Neurosci.* 24, 1173–1181
- 71 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
- 72 Osu, R. *et al.* (2004) Random presentation enables subjects to adapt to two opposing forces on the hand. *Nat. Neurosci.* 7, 111–112
- 73 Wada, Y. *et al.* (2003) Acquisition and contextual switching of multiple internal models for different viscous force fields. *Neurosci. Res.* 46, 319–331
- 74 Li, C.S. *et al.* (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
- 75 Padoa-Schioppa, C. *et al.* (2004) Neuronal activity in the supplementary motor area of monkeys adapting to a new dynamic environment. *J. Neurophysiol.* 91, 449–473
- 76 Padoa-Schioppa, C. *et al.* (2002) Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. *Neuron* 36, 751–765
- 77 Xiao, J. *et al.* (2006) Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. *Exp. Brain Res.* 168, 106–119
- 78 Arce, F. *et al.* (2010) Neuronal correlates of memory formation in motor cortex after adaptation to force field. *J. Neurosci.* 30, 9189–9198
- 79 Arce, F. *et al.* (2010) Combined adaptiveness of specific motor cortical ensembles underlies learning. *J. Neurosci.* 30, 5415–5425
- 80 Thoroughman, K.A. and Shadmehr, R. (1999) Electromyographic correlates of learning an internal model of reaching movements. *J. Neurosci.* 19, 8573–8588
- 81 Paz, R. *et al.* (2003) Preparatory activity in motor cortex reflects learning of local visuomotor skills. *Nat. Neurosci.* 6, 882–890
- 82 Paz, R. and Vaadia, E. (2004) Learning-induced improvement in encoding and decoding of specific movement directions by neurons in the primary motor cortex. *PLoS Biol.* 2, E45
- 83 Wolpaw, J.R. *et al.* (2002) Brain–computer interfaces for communication and control. *Clin. Neurophysiol.* 113, 767–791
- 84 Birbaumer, N. (2006) Breaking the silence: brain–computer interfaces (BCI) for communication and motor control. *Psychophysiology* 43, 517–532
- 85 Hwang, E.J. and Andersen, R.A. (2009) Brain control of movement execution onset using local field potentials in posterior parietal cortex. *J. Neurosci.* 29, 14363–14370
- 86 Scherberger, H. *et al.* (2005) Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron* 46, 347–354
- 87 Wang, W. *et al.* (2009) Human motor cortical activity recorded with Micro-ECoG electrodes, during individual finger movements. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* 2009, 586–589
- 88 Zhuang, J. *et al.* (2010) Decoding 3D reach and grasp kinematics from high-frequency local field potentials in primate primary motor cortex. *IEEE Trans. Biomed. Eng.* 57, 1774–1784
- 89 Lebedev, M.A. *et al.* (2005) Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain–machine interface. *J. Neurosci.* 25, 4681–4693
- 90 Mulliken, G.H. *et al.* (2008) Decoding trajectories from posterior parietal cortex ensembles. *J. Neurosci.* 28, 12913–12926
- 91 Mulliken, G.H. *et al.* (2008) Forward estimation of movement state in posterior parietal cortex. *Proc. Natl. Acad. Sci. U. S. A.* 105, 8170–8177
- 92 Musallam, S. *et al.* (2004) Cognitive control signals for neural prosthetics. *Science* 305, 258–262
- 93 Taylor, D.M. *et al.* (2003) Information conveyed through brain-control: cursor versus robot. *IEEE. Trans. Neural. Syst. Rehabil. Eng.* 11, 195–199
- 94 Zacksenhouse, M. *et al.* (2007) Cortical modulations increase in early sessions with brain–machine interface. *PLoS ONE* 2, e619
- 95 Paninski, L. *et al.* (2004) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *J. Neurophysiol.* 91, 515–532
- 96 Santhanam, G. *et al.* (2006) A high-performance brain–computer interface. *Nature* 442, 195–198
- 97 Chase, S.M. *et al.* (2009) Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain–computer interface algorithms. *Neural. Netw.*
- 98 Wahnoun, R. *et al.* (2006) Selection and parameterization of cortical neurons for neuroprosthetic control. *J. Neural. Eng.* 3, 162–171
- 99 Tkach, D. *et al.* (2007) Congruent activity during action and action observation in motor cortex. *J. Neurosci.* 27, 13241–13250
- 100 Tkach, D. *et al.* (2008) Observation-based learning for brain–machine interfaces. *Curr. Opin. Neurobiol.* 18, 589–594
- 101 Cisek, P. and Kalaska, J.F. (2004) Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature* 431, 993–996
- 102 Fetz, E.E. (1993) Cortical mechanisms controlling limb movement. *Curr. Opin. Neurobiol.* 3, 932–939
- 103 Reimer, J. and Hatsopoulos, N.G. (2009) The problem of parametric neural coding in the motor system. *Adv. Exp. Med. Biol.* 629, 243–259
- 104 Churchland, M.M. and Shenoy, K.V. (2007) Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* 97, 4235–4257
- 105 Ganguly, K. and Carmena, J.M. (2009) Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* 7, e1000153
- 106 Jarosiewicz, B. *et al.* (2008) Functional network reorganization during learning in a brain–computer interface paradigm. *Proc. Natl. Acad. Sci. U. S. A.* 105, 19486–19491
- 107 Sergio, L.E. *et al.* (2005) Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *J. Neurophysiol.* 94, 2353–2378
- 108 Wessberg, J. *et al.* (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408, 361–365
- 109 Wessberg, J. and Nicolelis, M.A. (2004) Optimizing a linear algorithm for real-time robotic control using chronic cortical ensemble recordings in monkeys. *J. Cogn. Neurosci.* 16, 1022–1035
- 110 Carmena, J.M. *et al.* (2005) Stable ensemble performance with single-neuron variability during reaching movements in primates. *J. Neurosci.* 25, 10712–10716
- 111 Chestek, C.A. *et al.* (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J. Neurosci.* 27, 10742–10750
- 112 Rokni, U. *et al.* (2007) Motor learning with unstable neural representations. *Neuron* 54, 653–666
- 113 Shen, L. and Alexander, G.E. (1997) Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *J. Neurophysiol.* 77, 1195–1212
- 114 Ochiai, T. *et al.* (2005) Involvement of the ventral premotor cortex in controlling image motion of the hand during performance of a target-capturing task. *Cereb. Cortex* 15, 929–937

- 115 Cisek, P. *et al.* (2003) Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. *J. Neurophysiol.* 89, 922–942
- 116 Davidson, A.G. *et al.* (2007) Rapid changes in throughput from single motor cortex neurons to muscle activity. *Science* 318, 1934–1937
- 117 Fetz, E.E. and Baker, M.A. (1973) Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *J. Neurophysiol.* 36, 179–204
- 118 Legenstein, R. *et al.* (2010) A reward-modulated hebbian learning rule can explain experimentally observed network reorganization in a brain control task. *J. Neurosci.* 30, 8400–8410
- 119 Diedrichsen, J. *et al.* (2005) Neural correlates of reach errors. *J. Neurosci.* 25, 9919–9931
- 120 Tseng, Y.W. *et al.* (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J. Neurophysiol.* 98, 54–62
- 121 Desmurget, M. *et al.* (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat. Neurosci.* 2, 563–567
- 122 Grafton, S.T. *et al.* (2008) Neural substrates of visuomotor learning based on improved feedback control and prediction. *Neuroimage* 39, 1383–1395
- 123 Wu, W. and Hatsopoulos, N.G. (2007) Coordinate system representations of movement direction in the premotor cortex. *Exp. Brain Res.* 176, 652–657
- 124 Gilja, V. *et al.* (2010) High-performance continuous neural cursor control enabled by a feedback control perspective. Conference Abstract: Computational and Systems Neuroscience 2010. DOI:10.3389/conf.fnins.2010.03.00249
- 125 Chase, S.M. *et al.* (2010) Latent inputs improve estimates of neural encoding in motor cortex. *J. Neurosci.* 30, 13873–13882
- 126 Kraskov, A. *et al.* (2009) Corticospinal neurons in macaque ventral premotor cortex with mirror properties: a potential mechanism for action suppression? *Neuron* 64, 922–930
- 127 Stamos, A.V. *et al.* (2010) The spinal substrate of the suppression of action during action observation. *J. Neurosci.* 30, 11605–11611