



Annual Review of Control, Robotics, and Autonomous Systems

Brain–Machine Interfaces: Closed-Loop Control in an Adaptive System

Ethan Sorrell,* Michael E. Rule,* and Timothy O’Leary

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom;
email: timothy.oleary@eng.cam.ac.uk

Annu. Rev. Control Robot. Auton. Syst. 2021.
4:24.1–24.23

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at
control.annualreviews.org

<https://doi.org/10.1146/annurev-control-061720-012348>

Copyright © 2021 by Annual Reviews.
All rights reserved

*These authors contributed equally to this article

Keywords

brain–machine interfaces, neural plasticity, optimal feedback control, neural coding, adaptive decoding

Abstract

Brain–machine interfaces (BMIs) promise to restore movement and communication in people with paralysis and ultimately allow the human brain to interact seamlessly with external devices, paving the way for a new wave of medical and consumer technology. However, neural activity can adapt and change over time, presenting a substantial challenge for reliable BMI implementation. Large-scale recordings in animal studies now allow us to study how behavioral information is distributed in multiple brain areas, and state-of-the-art interfaces now incorporate models of the brain as a feedback controller. Ongoing research aims to understand the impact of neural plasticity on BMIs and find ways to leverage learning while accommodating unexpected changes in the neural code. We review the current state of experimental and clinical BMI research, focusing on what we know about the neural code, methods for optimizing decoders for closed-loop control, and emerging strategies for addressing neural plasticity.



1. INTRODUCTION

Brain–computer interfaces and brain–machine interfaces (BMIs) promise an interface between the human brain and machines, with diverse applications ranging from treating epilepsy to restoring vision. In this review, we focus on implanted brain–machine interfaces that can extract and decode kinematic motor signals and abstract representations of motor plans from populations of neurons. These signals permit control of external hardware and communication devices.

Such devices are no longer science fiction: Implantable BMIs based on multi-electrode arrays entered human clinical trials in 2004 (1). State-of-the-art BMIs have demonstrated cortical control of robotic arms (2–5) as a proof of principle. BMIs for controlling computer cursors and typing are more advanced. For example, Nuyujukian et al. (6) have demonstrated full control of a tablet computer through a BMI, including web browsing and email. Using the on-screen keyboard and predictive word completion, users achieved typing speeds between 13 and 31 characters per minute. For text entry, novel approaches that decode handwriting gestures directly are even faster, achieving up to 90 characters per minute (7).

The possibility of extracting the electrophysiological activity of neurons in the cerebral cortex and using it to control an external interface was first demonstrated in the mid-twentieth century. In 1969, Fetz (8) showed that animals could learn to volitionally modulate the firing frequency of single neurons in the motor cortex to obtain a reward. This seminal study also presaged the benefits and challenges of neural adaptation in closed-loop interfaces. If a single neuron among a population of many millions can rapidly and reliably alter its firing activity, then BMI design must take this into account, and can potentially exploit it.

Neural engineering has advanced rapidly in the intervening decades, providing new technologies for high-bandwidth recording from the nervous system and a better understanding of BMI architecture based on closed-loop control. These technological advances have revealed new and pressing open challenges. While the first BMIs focused on brain areas specialized for movement, we now know that neural representations are multi-modal—fusing sensory, motor, and contextual information—as well as distributed, with motor-related information represented in many parts of the brain. The brain is also plastic and adaptive. The frontier of BMI research therefore lies in decoding and incorporating these distributed representations and designing decoding algorithms that act synergistically with neural plasticity and learning.

We organize this review around three central challenges in BMI engineering: (a) obtaining and decoding neural signals, (b) understanding the human–BMI system as a closed-loop controller, and (c) handling long-term learning, adaptation, and neural plasticity. First, we address the issue of interfacing with the brain itself, outlining different neural representations that are useful for BMI control and how these representations can be extracted. We also discuss emerging technologies that will shape BMI development in the coming decades. Second, we cover the design and architecture of BMI systems from the standpoint of closed-loop control. Third, we address the implications of neural plasticity and learning for BMI design and cover emerging technologies to address these challenges.

2. INTERFACING WITH THE BRAIN

At minimum, all BMIs must extract signals from the brain to use as control commands. Throughout this review, we interpret BMI design through the lens of the plant–controller framework of closed-loop control (**Figure 1**). In this nomenclature, the brain is the controller, and the computer or machine effector (e.g., a robot arm) is the plant. In control engineering, one typically identifies the plant and designs the controller. With BMIs, the brain is a preexisting controller, so the aim is to design and calibrate an interface and plant that can be controlled easily and reliably.



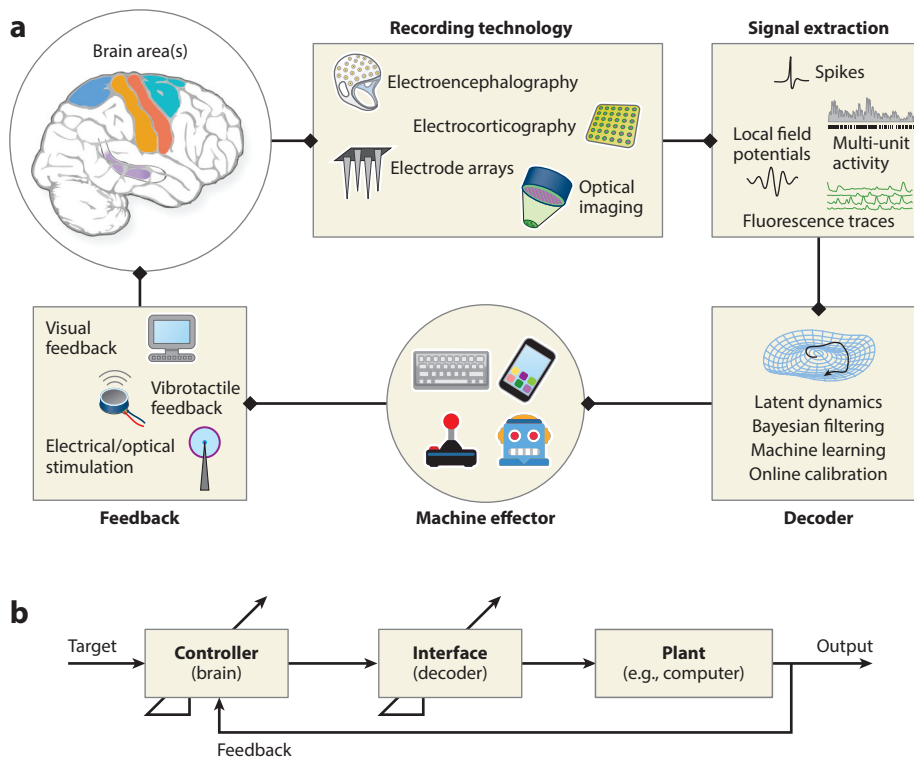


Figure 1

(a) A BMI consists of many distinct components, each of which must be selected based on the design considerations of the system. Generically, one identifies brain regions that contain relevant signals and records these signals with electrical or optical methods. Desired neural signals are extracted and processed, and a decoder translates these signals into control commands. These commands are passed to a machine effector, such as a robotic arm or computer. Information about the state of the controlled effector is provided to the user, usually through visual feedback, although specialized hardware or neural stimulation can also be used. (b) These components form a closed-loop controller. Here, the brain is a high-level controller that generates a target action. This target can be externally provided during calibration or training of the device (e.g., a visual cue) or volitionally generated by the user during use. The BMI decoder acts as an interface between the brain and the plant being controlled. Sensory feedback or (in some cases) additional stimulation interfaces close the loop. Crucially, both the controller (brain) and the BMI can adapt over time, denoted by diagonal arrows. Abbreviation: BMI, brain-machine interface.

High-bandwidth neural recording technologies are essential for this, but they are not the only challenge to be solved. One must also decide where to record from and understand how relevant information is encoded in neural activity. The prevailing view is that the relevant signals are encoded in the instantaneous firing rates of individual neurons, with the overall population firing rate vector representing the state of external or internal variables.

Areas of the brain are specialized, and some neural circuits are more involved in volitional movement than others. These circuits comprise many millions of individual cells whose combined population activity correlates with external stimuli and motor actions. Most BMI studies decode from the motor cortex. However, there is growing appreciation that representations for sensorimotor control can be found throughout the brain (Figure 2). This has motivated BMI studies targeting other brain areas that contain more abstract representations of motor tasks.

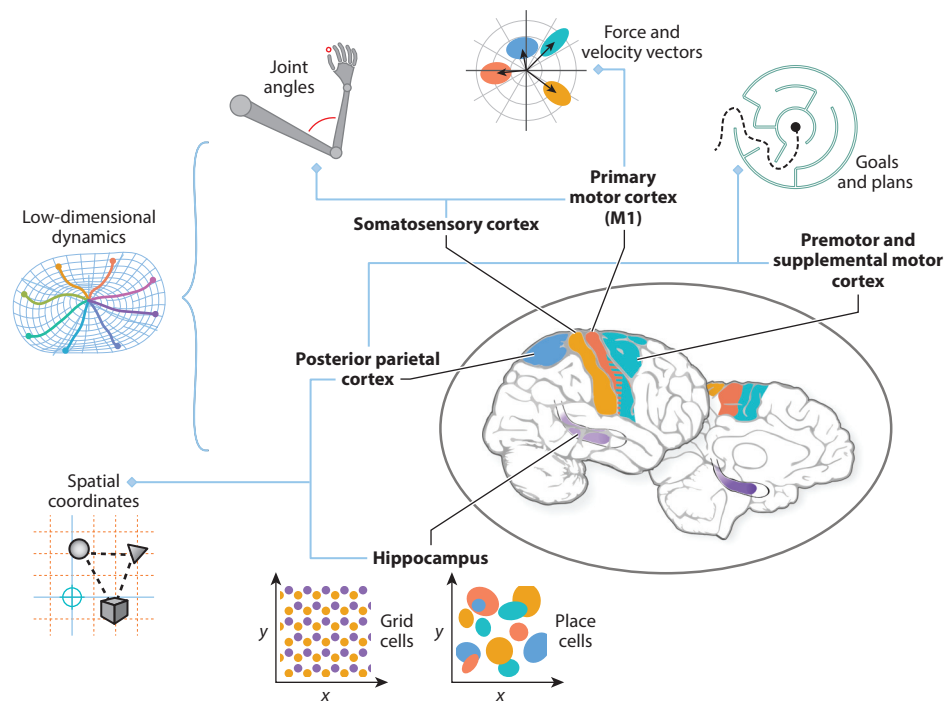


Figure 2

Sensorimotor representations are distributed throughout the brain. The primary motor cortex (M1) is considered the canonical output area, and limb movements and force and velocity commands can be decoded from it. The premotor cortex also sends connections to the spinal cord and encodes plans and goals as well as movement. The somatosensory cortex also contains information about limb state, in this case provided by sensory feedback. It is not used as a source of motor commands but is a target for stimulation to provide feedback in BMIs. Association areas of the cortex, such as the posterior parietal cortex, integrate sensory and motor variables and support planning and decision-making. Sensory, association, and motor areas collectively implement control policies during closed-loop motor control. Algorithms that attempt to identify a hidden or latent low-dimensional state space that describes how neural population activity evolves (sometimes called low-dimensional manifold dynamics) invariably find variables correlated to a specific task distributed throughout the brain. Representations of navigation in real and virtual environments can be found in the hippocampus and entorhinal cortex. Abbreviation: BMI, brain-machine interface.

2.1. Brain Areas and What They Encode

The motor cortex is directly involved in generating voluntary movement. Fetz (8) revealed that animals could volitionally control the output from single cells in the motor cortex (**Figure 3a**). Subsequent studies revealed that neuronal firing rates in this area correlate with the velocities of hand movements (11), and many cortical BMIs exploit this to decode hand position and velocity (1, 12, 13). The motor cortex consists of the primary motor cortex (M1) and several premotor and supplemental motor areas (**Figure 2**).

Neurons in the primary motor cortex encode kinematic variables or combinations thereof (14) as well as signals related to muscle contraction (15). Limb coordination and appropriate control of stiffness and compliance all depend on the state of the musculoskeletal system. Accordingly, the kinematic tuning of neurons in the motor cortex depends on the posture or the phase of movement being performed (16, 17). Although seemingly disparate, these various coding roles

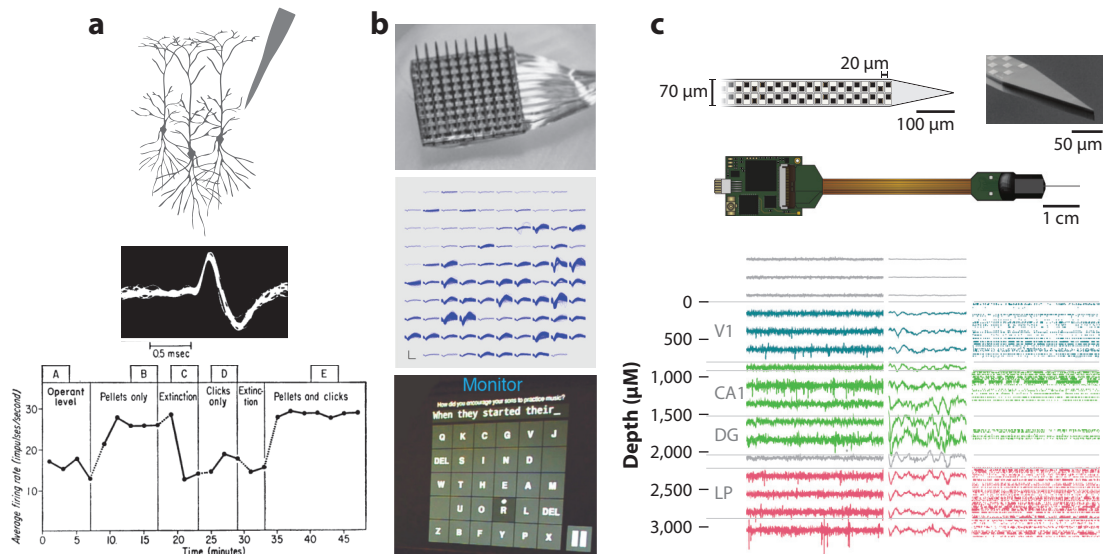


Figure 3

(a) One of the earliest precursors to a BMI (*top*) was demonstrated in 1969 by Fetz (8). In this work, monkeys learned to increase the firing rate of individual motor cortex neurons (*middle* and *bottom*) to receive food. Middle and bottom subpanels adapted with permission from Reference 8. (b) Developments of recording hardware and decoding software have led to clinical BMIs that record the activity of tens to hundreds of neurons using multi-electrode arrays (*top*) and use this activity (*middle*), for example, to control cursors for typing (*bottom*). Panel adapted from Reference 9 under the Creative Commons CC0 Public Domain Dedication license. (c) Next-generation multi-electrode arrays (*top*) can record activity from thousands of neurons (*bottom*) and will likely support higher-performance BMI systems in the future. Panel adapted with permission from Reference 10. Abbreviation: BMI, brain–machine interface.

are consistent with the view that the motor cortex generates movement and implements policies for musculoskeletal control (18).

Area M1 contains the majority of projections to the neurons that drive muscles, but it may not always be ideal for extracting BMI signals. For example, these output neurons can die in patients with motor neuron disease. Moreover, motor planning must to some extent convert distal, goal-oriented coordinates into raw muscle contraction signals. Such abstract motor plans are necessarily less dependent on the specifics of the musculoskeletal system and can be more useful for controlling a robotic arm or pointing device. Better BMI control may thus be possible from brain areas higher up in the processing hierarchy to the primary motor cortex.

In primates, the premotor cortex encodes discrete and abstract variables useful to BMI, such as motor plans (19) and reaching goals (e.g., 20). These variables complement decoding of continuous velocity commands to improve BMI usability (e.g., 21). The dorsal premotor cortex especially is used in BMI applications, either in conjunction with M1 (e.g., 22, 23) or in isolation (24).

The motor cortex translates goals and intentions into movement, but what if we could decode these intentions directly? This is the philosophy behind experimental BMIs that target an area of the brain called the posterior parietal cortex (PPC), which integrates sensory information for learning, planning, and decision-making (25). The PPC has a key role in visually guided movement and navigation (26–28), encoding locations, choices (29), and intentions (e.g., 30). It also contains predictive forward models for sensorimotor control, which plan future actions and anticipate their consequences, and therefore may be exceptionally useful to BMIs (31). In the context of BMIs, it has been used to decode goals and continuous trajectories for reaching and grasping (e.g., 32, 33). Animals can learn novel BMI readouts by volitionally controlling PPC activity (34, 35).

PPC-based BMIs have been tested in humans (36), and primate studies show decoding performance competitive with the motor cortex (37).

Motor and parietal cortices are a natural BMI target for visually guided movement, but variables relevant to other tasks, such as navigation, are encoded elsewhere. The hippocampus, located in the temporal lobe, is involved in memory and spatial navigation. An animal's location within an environment can be decoded from the hippocampus and the surrounding entorhinal cortex (e.g., 38, 39). In the hippocampus, place cells show high activity when the animal visits a specific area (40). In the entorhinal cortex, spatial modulation in the activity of grid cells shows a hexagonal pattern (**Figure 2**). These representations track location in real or virtual space (41) and enable navigation toward goals. One might therefore imagine that the hippocampus could be useful for BMIs that involve navigation. Existing recordings in the hippocampus are focused on treating epilepsy (42). Experiments in these patients have found that some people can volitionally control activity in the hippocampus and surrounding areas (43). Offline open-loop decoding of navigation has also been well demonstrated in the hippocampus, but how this brain area might be incorporated into a closed-loop BMI remains an exciting open question.

While the motor cortex has received the bulk of attention as a target for BMIs, ongoing research exploring abstract and associative areas is promising. Other brain areas, such as the hippocampus, remain largely unexplored for BMI applications, but studying them in the context of a closed-loop BMI may reveal new possibilities and clarify their function.

2.2. Recording Technology

The previous section described quantities that can be decoded from different brain areas. Here, we discuss how neural signals can be recorded and how control-related quantities can be extracted from neural population activity.

To extract signals for complex and precise motor control from the brain, one must record the simultaneous activity of large populations of neurons. Noninvasive methods such as electroencephalography necessarily lack the required bandwidth and resolution for many applications. Implanted recording devices in (or on) the surface of the brain are therefore preferable, despite their invasiveness. High-performance BMIs used in clinical trials extract neural population activity using implanted multi-electrode arrays (**Figure 3b**). Next-generation electrode arrays (**Figure 3c**) that can record many thousands of neurons (10, 44, 45) are likely to move from animal to human trials in the coming years.

Multi-electrode arrays detect a rich variety of neural signals that can be used to control a BMI (**Figure 4**). These include brief, <5-ms spiking events from single neurons, superimposed background spiking activity of multiple nearby cells (multi-unit activity), and lower-frequency (below ~100–300 Hz) collective electrical potentials known as local field potentials that arise from coordinated activity in large populations. In electrical recordings, individual cells can be identified based on their spike waveform. Some BMIs leverage this by including spike sorting in the signal-processing pipeline, a source separation method that assigns each spiking event to a specific neuron. Decoding from sorted spikes is common (e.g., 13, 47). However, cortical BMIs based on isolated single neurons suffer from recording instability (e.g., 48, 49). The shape of a spike waveform depends on the position of the electrode relative to the cell. Movement of electrodes and changes in tissue properties therefore degrade, distort, or lead to loss of unambiguous signals from individual cells.

A pragmatic and effective way to address this signal instability is to avoid relying on signals from identified cells altogether. Signals based on the average activity of many neurons are less sensitive to mechanical and physiological changes. Successful BMIs have been demonstrated based on






Figure 4

Implanted electrical recording methods include multi-electrode arrays (*top*), which are embedded in the cortex, or grids of electrodes placed over the surface of the brain (electrocorticography) (*middle*). Electrocorticography arrays can detect electrical local field potentials arising from neural population activity on scales ranging from centimeters down to fractions of a millimeter (e.g., 46). Modulation in various local field potential frequency bands correlates with motor and cognitive phenomena. Implanted multi-electrode arrays can detect spiking activity from isolated single neurons as well as local field potentials and have been the primary recording approach used in human clinical trials. Optical recording methods (*bottom*) are used in experimental research and allow simultaneous recording of activity from thousands of cells. All of these signals contain neural correlates of motor commands, which can be decoded and used to control a BMI. Abbreviation: BMI, brain-machine interface.

multi-unit activity (e.g., 2, 9, 50, 51) and local field potentials (e.g., 52–54). Despite the lower bandwidth of these population signals, performance competitive with BMIs based on isolated single neurons is possible (55, 56), and one can gain both the accuracy benefits of single neurons and the robustness of aggregate signals by combining the two (22).

In animal models, new recording methods based on single-cell optical imaging have been developed. These methods circumvent many of the difficulties with multi-electrode arrays because source separation can be achieved based on a cell's location in an imaging frame, permitting reliable signal extraction from many thousands of neurons (57). These methods employ dyes or proteins that fluoresce with an intensity that depends on the concentration of intracellular ions and in some cases the membrane potential of a cell, thus providing an optical readout of activity. For example, fluorescence of the protein GCaMP6 (58) depends on the concentration of calcium ions inside a cell, which in turn depends on neuronal firing rate.

Successful BMIs based on calcium imaging have been demonstrated in rodents (35, 59, 60) and nonhuman primates (61). New innovations in implanted microendoscopes now enable optical imaging in freely behaving nonhuman primates (62). The next challenge will be to translate these methods to humans. Calcium imaging requires genetic modifications or viral injection to express fluorescent proteins. Viral delivery of genes in humans is already being explored to restore sight in genetic disorders (63). Although substantial work is needed to develop a safe and effective way to deliver the genes required for calcium imaging in humans, this is a key first step toward developing optical BMIs in humans.

3. NEURAL DECODING FOR THE BRAIN–MACHINE INTERFACE

High-density recording methods allow us to observe the activity of large populations of neurons. Even though our understanding of this activity remains incomplete, neural correlates of movement can be extracted using standard signal-processing methods, which can then be used to control a BMI. State-of-the-art BMIs incorporate more advanced methods of signal filtering and decoding. Additionally, recently developed approaches allow BMIs to be calibrated online and to optimize their behavior based on models of closed-loop feedback control.

3.1. Decoding Algorithms

The millions of neurons in the motor cortex encode a high-dimensional signal that ultimately controls the contraction of all voluntary muscle groups in the body. Given this, one might assume that a simple relationship between, say, the velocity of a computer cursor and neural activity in the motor cortex might be hard to identify. However, in the context of a fixed, repeated task, the firing rates of single cells are correlated with quantities relevant for a BMI, such as velocity in a specific direction (11). This allows control signals to be decoded from neural population activity using a simple weighted sum (e.g., 1, 4, 64). Accurate decoding is possible even without a complete model of what each neuron encodes.

Neural population codes thus display another critical feature that makes BMI decoding possible: The codes are redundant. Because correlations between neurons constrain the space of possible activity patterns (e.g., 65, 66), the dimensionality of neural population activity is often much lower than the number of cells being recorded. Indeed, the structure of neuronal variability closely resembles the structure of the task being solved (e.g., 67). It is therefore typically possible to identify a hidden, or latent, low-dimensional state space that describes how activity in a large neural population evolves in time. The latent variables in this state-space description can then be mapped to behavioral or kinematic variables relevant to BMIs. Accordingly, after extracting the activity of single neurons, BMI signal-processing pipelines often include a linear dimensionality reduction step. This redundancy allows a multi-electrode array to extract control signals from a small and random subsample of neurons, and also implies that many different decoders can work well. While redundancy reduces the number of neurons required to extract control signals, it may also contribute to instability in BMI readouts. We return to this issue in Section 4.

These features of neural population codes enable basic BMI decoding using simple linear readouts. Both continuous commands like limb velocity (via linear regression) and discrete events like mouse clicks (via linear discriminant analysis) can be decoded this way. Contemporary high-performance BMIs often combine both, e.g., mouse movement and clicks (68) or reaching and grasping (2).

Incorporating recursive Bayesian filtering can bring further improvements. In some cases, good control can be achieved by modeling both neural dynamics and their relation to movement as linear with Gaussian noise. The Kalman filter is optimal under these assumptions and has proven effective (69, 70). However, some variables, such as movement speed, are poorly decoded under linear assumptions (71). There have been efforts to extend the Kalman filter to nonlinear dynamics (72) and non-Gaussian observations (73, 74). Recurrent neural networks and related machine learning approaches have also been explored (e.g., 75, 76) and led to more naturalistic BMI control in primate studies (77).

In nonhuman primate studies, BMIs can be calibrated by training the decoder to predict the subject's movements. For people with paralysis, this is not possible. However, initial calibration can be achieved by recording neural activity while a patient watches, imagines, or attempts to perform an instructed action (e.g., 1). These tasks evoke neural activity similar to movement



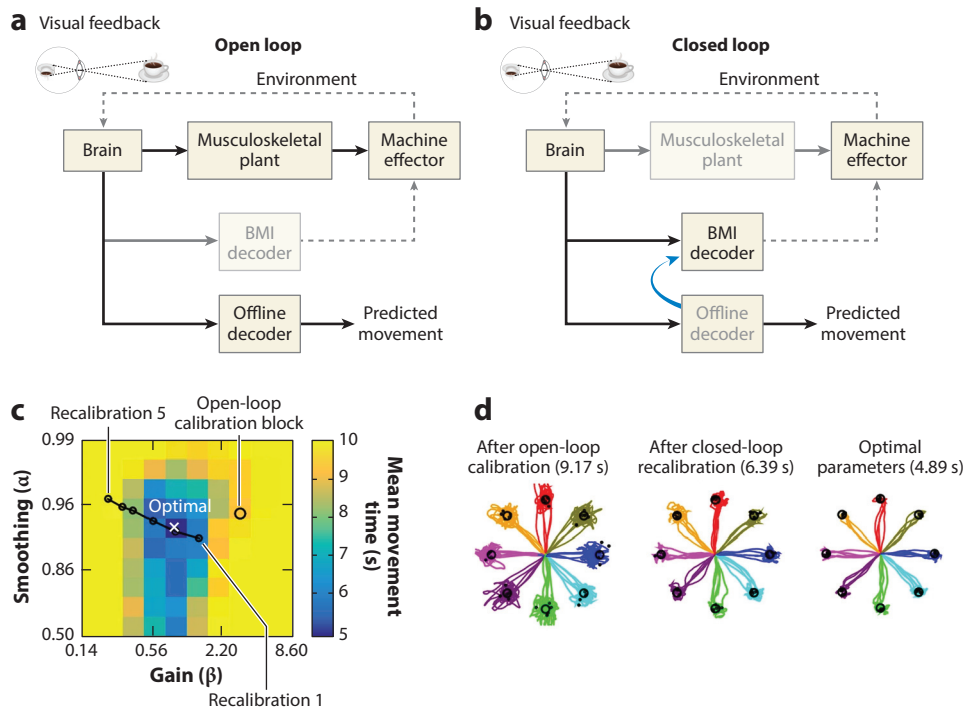


Figure 5

(a) Open-loop decoding entails predicting, e.g., arm trajectories from neural activity, but these decoded results are not used for control in real time. (b) A decoder trained in open loop can be used to control a BMI. Due to feedback, however, this decoder is rarely optimal. Recalibrating the decoder based on closed-loop activity leads to improvements. Further improvements are possible by identifying the user's control policy. (c) Willett et al. (81) identified users' control policy from initial training data, which allowed them to predict the optimal decoder for closed-loop control. This plot shows the predicted decoder performance as a function of the decoder gain and the amount of temporal smoothing. Panel adapted from Reference 81 under a Creative Commons Attribution 4.0 International license (<https://creativecommons.org/licenses/by/4.0>). (d) Center-out reaching trajectories simulated from their model recapitulate the differences between open-loop and closed-loop calibration and predict better control using the optimized parameters. Panel adapted from Reference 81 under a Creative Commons Attribution 4.0 International license (<https://creativecommons.org/licenses/by/4.0>). Abbreviation: BMI, brain-machine interface.

(78, 79) and can therefore be used to calibrate a decoder offline. This initial calibration suffices for rudimentary control and can be refined during BMI use, as we will see in the next section.

3.2. Open Versus Closed Loop

The earliest experiments to decode movement from the brain were open loop (e.g., 80): Neural activity was recorded during movement execution, and these movements were reconstructed after the fact using a decoder trained offline (Figure 5a). Open-loop experiments are necessary to understand how the brain encodes movement (11) and are a prerequisite for building a BMI decoder. Closed-loop experiments require at minimum a rudimentary decoder in order for the control loop to be closed at all. Many challenges in BMI design stem from the differences between open-loop calibration and closed-loop control. In closed-loop control, neural activity must be decoded in real time, and the resulting output is sensed by the user (Figure 5b).

Early experiments revealed differences in neuronal properties and activity between open-loop decoding and closed-loop control. Because users can detect and correct mistakes, closed-loop control can be better than open-loop decoding (82, 83). Practice in closed-loop BMI control can even lead to apparent changes in neuronal tuning as the user's motor system adapts (84). Consequently, optimal parameters for a decoder in closed loop can differ from those in open loop (e.g., 85). For example, providing rapid error feedback is more important than getting the BMI to filter out errors (86).

These observations have led to techniques for improving closed-loop performance, broadly known as closed-loop decoder adaptation. These techniques include the recalibrated feedback intention-trained (ReFIT) Kalman filter (6, 9, 13, 87) and other similar methods (88). Such methods refine a decoder based on data recorded during closed-loop BMI use. They use the goal or ultimate outcome of a movement (e.g., moving to a target) to better infer the meaning of past neural signals in retrospect (89). Closed-loop decoder adaptation improves on open-loop calibration but does not converge on parameters optimal for control (88).

None of the differences between the open- and closed-loop behavior of BMI devices will surprise control engineers. In engineering, principled design of controllers is achieved with a model of the plant. Analogously, principled design of a BMI may benefit from a model of the brain as an adaptive controller (81, 90). Such modeling begins with the assumption that the brain itself approximates optimal feedback control (91). When combined with a model of how control outputs are encoded in neural activity (e.g., 92), optimal feedback control models provide a principled route to optimizing closed-loop BMIs. Even simple models of the brain as a linear controller can explain the changes observed during closed-loop BMI use. For example, optimal feedback control predicts that neuronal tuning should change during closed-loop control in order to compensate for biases in the decoder (93) and that rapid feedback is important (81, 94) because it is better to have the user quickly correct mistakes than to have the BMI try to filter them out.

State-of-the-art calibration approaches now incorporate models of the brain as an optimal controller. Willett et al. (81) identified users' control policies using a piecewise linear model and found that, while closed-loop recalibration improves on open-loop initialization, it does not converge to the optimal decoder from a control perspective (**Figure 5c**). In a simulated center-out reaching task, this calibrated model recapitulated the improvements between open- and closed-loop calibration and predicted parameters that might further improve control (**Figure 5d**). The ability to simulate and study closed-loop control without running new experiments allowed the authors to design a new decoder offline, which then showed clear improvements in human testing.

Modeling the brain as a closed-loop optimal feedback controller allows researchers to optimize BMI decoding in a principled way. However, even with these optimizations, a BMI cannot be expected to achieve performance comparable to musculoskeletal control, because the control loop in a BMI is fundamentally different. In an intact musculoskeletal system, the spinal cord aids in stabilizing and generating movement, and proprioceptors provide rapid state feedback by sensing the contraction state of muscles and joints. This low-latency, hierarchical control is absent in a BMI. Users must instead rely on visual or haptic (95) feedback, which has higher latency and lower bandwidth.

BMIs must find a way to replace these low-level control and feedback functions or provide solutions that do not require them. Invasive approaches that electrically stimulate the cortex promise high-bandwidth and low-latency feedback (e.g., 96). Successful bidirectional BMIs using this technology have already been demonstrated (97, 98), opening up new experiments to explore how to use feedback to improve control. Optogenetic stimulation methods have also been developed in rodents (99, 100), and work is ongoing to translate them to primates (101, 102). Alternatively, one may sidestep the need for rapid feedback by delegating aspects of low-level control that require



rapid feedback to the BMI itself. Some approaches incorporate robotic sensors in a shared control system (e.g., 3, 5) and sensor fusion that can be used to improve object grasping and manipulation (103).

4. NEURAL PLASTICITY AND THE BRAIN–MACHINE INTERFACE

In addition to adapting rapidly to biases in closed-loop control, the brain changes over slower timescales. Learning can increase BMI proficiency over time, but new research reveals that the brain continues to change even without overt learning. While adaptability has benefits, ever-evolving neural population codes make long-term stable decoding challenging. In this section, we address how BMIs can handle nonstationarity in neural signals. We cover adaptive processes that allow users to improve BMI control, as well as the limitations of this adaptation. We address the phenomenon of drift, wherein neural representations gradually reconfigure, and discuss decoding algorithms that adapt to changing neural codes or robustly ignore irrelevant changes. We then close with a discussion of the implications of distributed and multi-modal neural codes for drift and adaptive BMIs.

4.1. Adaptation and Learning

Fetz's (8) original experiments highlighted that single neurons can adapt rapidly to control a BMI. Indeed, subjects can control motor cortex neuron activity to use a novel and arbitrary BMI readout (104). Learning volitional control of arbitrary neurons is possible not only in the motor cortex but also in other brain regions, such as the PPC (35) or visual cortex (105). Subjects can also learn to modulate populations of cells together to achieve a more robust interface (106). The extent to which neurons in various parts of the brain can be conditioned to events in the external world remains an open question, and one that may challenge a strict anatomical classification of brain areas as being predominantly motor or sensory. However, as we learn from studies in this section, there are constraints on the adaptability of neural activity, and these ultimately constrain BMI design.

Animal studies have revealed that learning consists of multiple phases governed by different mechanisms (107). Initial practice with a novel BMI entails exploratory and apparently random neural variability (e.g., 35, 108), while long-term learning sees the emergence of new internal models that incorporate the device (109, 110). BMIs can take advantage of this longer-term learning, and indeed subjects report a transition from deliberate to automatic control of a BMI with practice (111).

In light of these studies, one might optimistically conclude that nearly any BMI readout could be learned. However, the configuration of the BMI readout affects learning difficulty. Neural population dynamics arise from the collective activity of many neurons, which generates low-dimensional dynamics (65). This constrains the forms of learning and adaptation that are possible, because low-dimensional dynamics reflect correlations between single neurons, some of which likely arise from rigid constraints on anatomical connectivity. Widespread changes in synaptic connections are therefore required for the brain to generate activity with a very different correlation structure.

Learning to use a novel interface is easier if the readout is aligned with the preexisting subspace of neural activity in which signals can fluctuate (112) and difficult if it is misaligned (113). This is sometimes attributed to the ability to use cognitive strategies (such as reaiming) in learning to use the decoder (114). Rapid adaptation to novel BMI readouts entails harnessing the existing behavioral repertoire in new ways (115, 116), rather than learning new and arbitrary outputs (34, 117). Learning novel mappings is slower, taking several days of active practice (118).



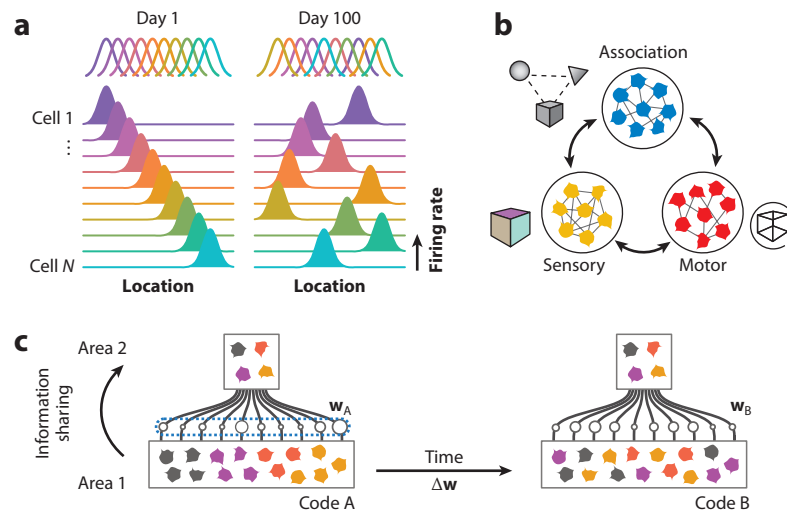


Figure 6

Neural representations reconfigure over time. (a) Shown here is the type of drift observed by Driscoll et al. (124), in which the location tuning of cells in the parietal cortex changes over days to weeks. (b) Different brain areas share distributed representations of the same information. Neural computation requires that these areas communicate, despite plasticity in some populations. (c) For example, reconfiguration of the neural code in one region (area 1) would require compensatory changes in a second, interconnected area. Here, we illustrate a scenario where the code in the second area is stable, and this area therefore adjusts synaptic weights (Δw) to track drift in area 1.

Although it is possible for the brain to use a fixed representation, the extent to which a fixed decoder might influence neural plasticity is currently unclear. Long-term practice with a fixed BMI readout leads to steady improvement and stabilization of the neural activity used for control (23, 119, 120), and dictating a new, fixed mapping changes the roles of single neurons in a BMI application (121). Theoretical studies (122) predict that better control can be achieved by requiring the user to learn a new readout that improves controllability compared with native neuronal tunings. Enforcing stability in a BMI readout could therefore train neural populations to provide more useful control signals (123).

4.2. Neural Representations Drift Over Time

Recent research has revealed that neural representations change even in learned, stable behaviors. For example, neuronal tuning to spatial location gradually reconfigures over time (Figure 6a) in both the PPC (124) and the hippocampus (125). While this reconfiguration may be slower in motor areas (126), any BMI that decodes from abstract brain areas must contend with it. This effect has been termed representational drift (127) and could contribute to the long-term instability seen in BMIs. Drift suggests that continual recalibration might be necessary even with stable recordings and after extensive learning.

Although the roles of single neurons can change, overall the brain maintains stable sensorimotor skills, which implies that communication and coordination between brain areas is unaffected by drift (Figure 6b). Drift must therefore be compensated for by additional plasticity or ignored by downstream areas (Figure 6c), which has two implications for BMIs. First, decoders must compensate for drift to achieve a stable readout. Second, ongoing BMI use may affect

endogenous plasticity mechanisms; for example, a fixed BMI readout from an association area might limit drift. This may lead to consolidation of skillful BMI use (23, 119, 120, 128), but the extent to which drift can be stabilized remains unclear.

Interpreting drift at the population level is complicated by the fact that low-dimensional dynamics permit many valid readouts. Consider, for example, a BMI that extracts velocity as a linear projection of population activity. While this readout may work during calibration, there is no guarantee that it is unique. The readout learned by the BMI may not even be directly related to motor output. For example, monkeys can learn to control a BMI without moving (e.g., 129, 130), even when the BMI was initially calibrated to decode arm movements, a finding that may be explained by orthogonality between the subspaces for arm and brain control (131). Given that it is difficult to identify a projection uniquely related to movement, a natural question is whether drift is an artifact of a failure to identify directions in neural activity stably related to movement.

Although the directions of activity used by a decoder may be unstable, other stable representations may exist (132). There is evidence that stable representations can be identified by recording over long time periods (133). However, even seemingly stable representations may gradually change (134). Additionally, identifying a fixed readout may not be optimal for BMI usability. The motor system contains a mixture of stable representations and volatile ones that might be used for learning (108). These volatile representations are nevertheless informative, and it may be better to have an adaptive BMI that can coadapt along with a neural code.

4.3. Adaptive and Robust Brain–Machine Interfaces

Much of the progress to translate proof-of-concept devices to the high-performance devices used in clinical trials can be attributed to improved decoding strategies. Automatically recalibrating and adapting decoders have been essential in achieving stable and robust long-term use. These adaptive decoders help BMIs tolerate recording instability and drift and reduce the amount of ongoing learning required to maintain proficiency.

In adaptive BMIs, the user and the decoder adapt together to improve performance. To do so, methods developed to calibrate decoders during closed-loop operation (Section 3.2) can be run periodically (e.g., 135) or continuously (e.g., 136). Adaptive and online filtering methods can therefore address instability and plasticity (e.g., 89, 135, 136). A nonhuman primate study by Orsborn et al. (135) (**Figure 7**) found that closed-loop decoder adaptation led to improved performance over the course of two weeks. They also found that user learning continued to improve above and beyond the improvements achieved by closed-loop decoder adaptation.

Adaptive decoders can compensate for instability, but the details of how they interact with learning in the long term remain unclear. Theoretical results (137) suggest that an adapting BMI system cannot do better than a fixed one if the subject's learning is optimal. However, learning is almost certainly not optimal. For example, Feulner & Clopath (113) showed that the brain's ability to understand the neural causes of control errors is critical, and that misconfigured BMI decoders could make this exceptionally difficult. At minimum, then, one might expect that BMI decoders should adapt to facilitate learning.

Rather than having a decoder adapt to changes, it may also be possible to ignore these changes altogether. New machine learning strategies to design BMI readouts that are robust to drift are promising. These strategies include using large amounts of training data and simulated perturbations to train a neural network to be robust to drift (133), using reinforcement learning to track drift (138), and using transfer learning to realign decoders from previous days based on a few sample trials (139). There have also been promising human trials of decoders based on Gaussian processes that are inherently robust (73) and support rapid calibration (140).



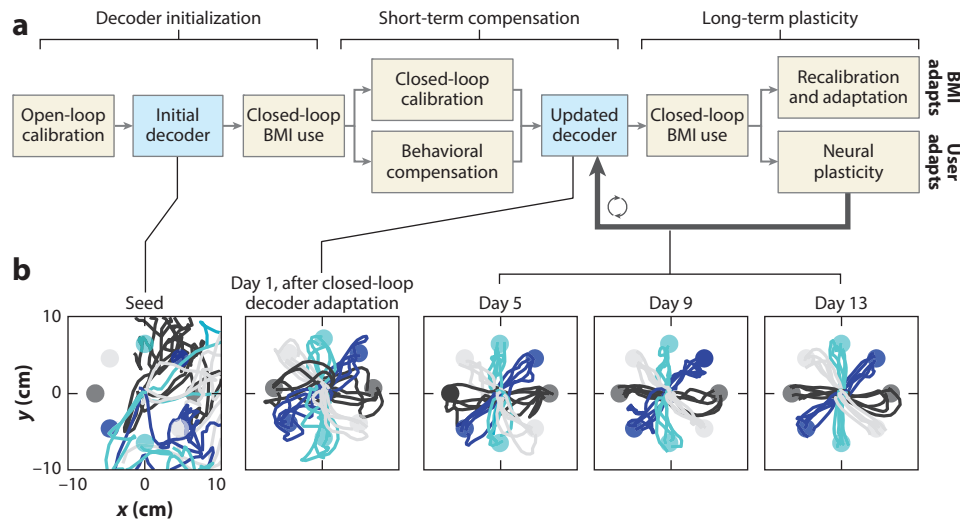


Figure 7

Through coadaptation in closed-loop control, learning performance improves as both the user and decoder adapt over time. (a) Learning and adaptation proceeds in multiple phases. The initial decoder calibrated in open loop is rarely optimal. Closed-loop decoder adaptation can improve closed-loop control. Users also adopt behavioral strategies to compensate for bias in the decoder. In the long term, adaptation and recalibration of the decoder combined with ongoing neural plasticity and learning can lead to continued improvement of BMI control. (b) Example trajectories from a center-out reaching task. User and BMI adaptation in closed loop leads to initial rapid improvement. Accuracy continues to improve with ongoing practice and periodic closed-loop recalibration. Panel adapted with permission from Reference 135. Abbreviation: BMI, brain-machine interface.

Even if neural population representations drift, the latent low-dimensional dynamics that they encode are remarkably consistent and continue to encode relevant task variables (66). These latent structures can be extracted with unsupervised learning (141) and used to control a BMI (142). Creating a stable BMI readout by registering the latent dynamics over time and even across individuals is a promising area of active research (e.g., 143–145). Latent population dynamics can be extracted even without spike sorting (146), conferring robustness to recording instability.

4.4. Multi-Area, Multi-Modal, and Context-Aware Decoders

During movement, multiple brain areas become engaged with a task and act in unison. For example, the premotor cortex in monkeys is involved in planning and intention and represents planned movements before they occur (20). During movement, however, motor and premotor cortices become engaged as part of the same dynamical system. Could targeting multiple brain areas improve decoding? Philip et al. (37) compared decoding using the M1 and PPC areas individually and together and found no improvement when using a combined decoder (37). This result suggests that these regions have redundant information, at least during movement. While using a combined decoder may offer only limited benefits when decoding brain activity on short timescales, on longer timescales it may provide a means of contending with drift and reconfiguration in the neural code. The fact that distinct brain areas pool and exchange information during behavior means that encoding of information needs to coevolve and adapt consistently between them. A decoder with access to multiple independent codes may thus be able to identify a consensus code that is more stable over time.

Multi-area decoding shows more promise when using higher-order planning and intention information to guide the interpretation of motor signals. Decoding both reaching commands and goals simultaneously can improve performance (147, 148), either from different brain areas (80) or from the same areas (149). Neural activity also contains representations of errors that reflect the discrepancy between intended movements and observed results (150). New research is exploring ways to use these error signals to help a BMI detect decoding errors and adapt its readout to compensate (151).

In addition to understanding how the brain encodes nonmotor variables, BMIs benefit from understanding how behavioral context affects how movement is encoded. For example, neural signals used in BMI are influenced by sensory feedback (152, 153), reward value (154, 155), and the proximity to a target in a reach-to-grasp task (156). Even similar movements, such as pedaling forward and backward on a bike (157), reflect different contexts and can be encoded in separable subspaces of population activity. Multiple neural codes can even exist for the same context (158).

Better understanding of contextual encoding may allow us to construct decoders that are stable across contexts. For example, Downey et al. (156) developed a decoder that can cancel out changes in neural firing rates that occur as a user approaches a target, and Zhao et al. (155) successfully removed the influence of an expected reward in an offline decoding study. Some representations are also inherently context invariant: A component of low-dimensional manifold dynamics (159) and the neural signatures demarcating different task phases (160) can be similar across tasks. Classifying context from brain signals, including the intention to use a BMI or not (161), and using separate decoders according to the context (162) are other potential solutions.

New tools and methods of data analysis are allowing us to extract more information about how the brain encodes actions, how disparate brain areas interact, and how the neural code evolves in time. Instability in BMI performance may thus simply reflect our own limited understanding of the neural code. While some instability may be due to recording methods, other changes are related to learning and processes that we do not yet fully understand. Future BMIs will likely benefit from better algorithms to identify stable representations, track and use evolving population codes, and account for the influence of cognitive variables and context in neural coding.

5. CONCLUSION

Many of the hardware and software challenges involved in creating high-performance BMIs have been solved, and work is ongoing to make motor neuroprosthetics available in clinical practice. Further improvements are likely to come from studying how distributed, multi-modal population representations evolve under routine BMI use. High-volume recording methods using optogenetics now enable such experiments in animal models. These methods will allow us to build a deeper understanding of neural plasticity in closed-loop BMI control, which in turn will guide further refinement in motor neuroprosthetics. In the long term, translating state-of-the-art high-volume recording technology to humans may lead to higher-bandwidth BMIs and new ways to provide feedback to the brain. Likewise, ongoing advances in machine learning and robotics will improve the intelligence of BMI effectors.

Some of the most important advances may come not from optimizing the BMI per se but from optimizing the usefulness of the computer or machine being controlled. For example, enabling patients to use a tablet computer (6) or making BMIs more portable (47, 163–165) can have as much (or more) impact on quality of life as increasing decoder performance by a few percent. For some tasks, BMIs already perform similarly to other human–computer interfaces: Controlling a cursor to type on a screen resembles using a video-game controller to do the same. Naturalistic typing, however, might require an interface more similar to touch-typing and decoding neural signals, such as finger movements (166) or handwriting gestures (7). Future BMIs may even decode



speech directly (167). Solving the design challenges of user interfaces with limited bandwidth can be addressed independently of BMI hardware trials and can also benefit users ineligible for current clinical BMIs.

Modeling the brain as an optimal feedback controller has proved tremendously useful. Even so, we are still learning how neural activity relates to optimal feedback control during the execution of a task, especially in more abstract brain areas. This limits our ability to decode and dissociate control-related neural representations and to predict how these representations change in closed loop and under long-term learning. Constructing more detailed simulations of the sensorimotor system as a controller will allow us to better understand neural signals and how to use them in a BMI. Additionally, bidirectional optogenetic BMIs that both record and stimulate enable experiments to reveal how the brain functions. Closed-loop optogenetic BMIs in rodents have been used to manipulate neural circuit activity during behavior (168), understand how the sensory causes of movement affect plasticity (99), and interrogate the functional properties of different types of neurons (169).

The emerging consensus is that, in order to achieve stable BMI performance, one must also understand neural plasticity during closed-loop sensorimotor control. Rather than viewing neurons as having a fixed tuning, or simply allowing for a fixed but context-dependent control policy, one should view neural populations as implementing learning strategies for achieving and maintaining good control. This can help us understand neural plasticity and help us design better adaptive decoders. Ultimately, better BMI technology requires us to understand and exploit the same kinds of adaptive processes that are at work in the brain. The work we have surveyed here shows that this understanding is well underway. Formidable challenges remain, and overcoming them will require coordinated advances and collaboration among fields in neuroscience, control engineering, robotics, and automation.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

E.S. was supported by the Cambridge Commonwealth, European, and International Trust and the Cambridge Institute for Neuroscience PhD studentship. M.E.R. and T.O. were supported by Human Frontier Science Program grant RGY 0069/2017. T.O. was supported by European Research Council grant FLEXNEURO (716643).

LITERATURE CITED

1. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, et al. 2006. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164–71
2. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, et al. 2012. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485:372–75
3. Vogel J, Haddadin S, Jarosiewicz B, Simeral J, Bacher D, et al. 2015. An assistive decision-and-control architecture for force-sensitive hand–arm systems driven by human–machine interfaces. *Int. J. Robot. Res.* 34:763–80
4. Wodlinger B, Downey JE, Tyler-Kabara EC, Schwartz AB, Boninger ML, Collinger JL. 2015. Ten-dimensional anthropomorphic arm control in a human brain-machine interface: difficulties, solutions, and limitations. *J. Neural Eng.* 12:016011



5. Downey JE, Weiss JM, Muelling K, Venkatraman A, Valois JS, et al. 2016. Blending of brain-machine interface and vision-guided autonomous robotics improves neuroprosthetic arm performance during grasping. *J. NeuroEng. Rehabil.* 13:28
6. Nuyujukian P, Albites Sanabria J, Saab J, Pandarinath C, Jarosiewicz B, et al. 2018. Cortical control of a tablet computer by people with paralysis. *PLOS ONE* 13:e0204566
7. Willett FR, Avansino DT, Hochberg LR, Henderson JM, Shenoy KV. 2020. High-performance brain-to-text communication via imagined handwriting. bioRxiv 2020.07.01.183384. <https://doi.org/10.1101/2020.07.01.183384>
8. Fetz EE. 1969. Operant conditioning of cortical unit activity. *Science* 163:955–58
9. Pandarinath C, Nuyujukian P, Blabe CH, Sorice BL, Saab J, et al. 2017. High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife* 6:e18554
10. Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, et al. 2017. Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551:232–36
11. Georgopoulos A, Schwartz A, Kettner R. 1986. Neuronal population coding of movement direction. *Science* 233:1416–19
12. Truccolo W, Friehs GM, Donoghue JP, Hochberg LR. 2008. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *J. Neurosci.* 28:1163–78
13. Gilja V, Pandarinath C, Blabe CH, Nuyujukian P, Simeral JD, et al. 2015. Clinical translation of a high-performance neural prosthesis. *Nat. Med.* 21:1142–45
14. Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP. 2004. Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *J. Neurophysiol.* 91:515–32
15. Bullard AJ, Nason SR, Irwin ZT, Nu CS, Smith B, et al. 2019. Design and testing of a 96-channel neural interface module for the Networked Neuroprosthesis system. *Bioelectron. Med.* 5:3
16. Hatsopoulos NG, Xu Q, Amit Y. 2007. Encoding of movement fragments in the motor cortex. *J. Neurosci.* 27:5105–14
17. Kadmon Harpaz N, Ungarish D, Hatsopoulos NG, Flash T. 2019. Movement decomposition in the primary motor cortex. *Cereb. Cortex* 29:1619–33
18. Todorov E. 2003. On the role of primary motor cortex in arm movement control. In *Progress in Motor Control*, Vol. 3: *Effects of Age, Disorder, and Rehabilitation*, ed. ML Latash, MF Levin, pp. 125–66. Champaign, IL: Hum. Kinet.
19. Hatsopoulos N, Mukand J, Polykoff G, Friehs G, Donoghue J. 2005. Cortically controlled brain-machine interface. In *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*, pp. 7660–63. Piscataway, NJ: IEEE
20. Vargas-Irwin CE, Franquemont L, Black MJ, Donoghue JP. 2015. Linking objects to actions: encoding of target object and grasping strategy in primate ventral premotor cortex. *J. Neurosci.* 35:10888–97
21. Achtman N, Afshar A, Santhanam G, Yu BM, Ryu SI, Shenoy KV. 2007. Free-paced high-performance brain-computer interfaces. *J. Neural Eng.* 4:336–47
22. Stavisky SD, Kao JC, Nuyujukian P, Ryu SI, Shenoy KV. 2015. A high performing brain-machine interface driven by low-frequency local field potentials alone and together with spikes. *J. Neural Eng.* 12:036009
23. Athalye VR, Ganguly K, Costa RM, Carmena JM. 2017. Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* 93:955–970.e5
24. Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. 2006. A high-performance brain-computer interface. *Nature* 442:195–98
25. Andersen R, Musallam S, Burdick J, Cham J. 2005. Cognitive based neural prosthetics. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pp. 1908–13. Piscataway, NJ: IEEE
26. Buneo CA, Andersen RA. 2006. The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia* 44:2594–606
27. Calton JL, Taube JS. 2009. Where am I and how will I get there from here? A role for posterior parietal cortex in the integration of spatial information and route planning. *Neurobiol. Learn. Mem.* 91:186–96
28. Minderer M, Brown KD, Harvey CD. 2019. The spatial structure of neural encoding in mouse posterior cortex during navigation. *Neuron* 102:232–48.e11



29. Harvey CD, Coen P, Tank DW. 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68
30. Snyder LH, Batista AP, Andersen RA. 1997. Coding of intention in the posterior parietal cortex. *Nature* 386:167–70
31. Cui H. 2016. Forward prediction in the posterior parietal cortex and dynamic brain-machine interface. *Front. Integr. Neurosci.* 10:35
32. Bosco A, Breveglieri R, Filippini M, Galletti C, Fattori P. 2019. Reduced neural representation of arm/hand actions in the medial posterior parietal cortex. *Sci. Rep.* 9:936
33. Filippini M, Breveglieri R, Akhras MA, Bosco A, Chinellato E, Fattori P. 2017. Decoding information for grasping from the macaque dorsomedial visual stream. *J. Neurosci.* 37:4311–22
34. Hwang E, Bailey P, Andersen R. 2013. Volitional control of neural activity relies on the natural motor repertoire. *Curr. Biol.* 23:353–61
35. Clancy KB, Mrsic-Flogel TD. 2019. The sensory representation of causally controlled objects. bioRxiv 786467. <https://doi.org/10.1101/786467>
36. Aflalo T, Kellis S, Klaes C, Lee B, Shi Y, et al. 2015. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science* 348:906–10
37. Philip BA, Rao N, Donoghue JP. 2013. Simultaneous reconstruction of continuous hand movements from primary motor and posterior parietal cortex. *Exp. Brain Res.* 225:361–75
38. Tampuu A, Mattisen T, Ólafsdóttir HF, Barry C, Vicente R. 2019. Efficient neural decoding of self-location with a deep recurrent network. *PLoS Comput. Biol.* 15:e1006822
39. Tu M, Zhao R, Adler A, Gan WB, Chen ZS. 2020. Efficient position decoding methods based on fluorescence calcium imaging in the mouse hippocampus. *Neural Comput.* 32:1144–67
40. O’Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34:171–75
41. Garvert MM, Dolan RJ, Behrens TE. 2017. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife* 6:e17086
42. Schevon CA, Tobochnik S, Eissa T, Merricks E, Gill B, et al. 2019. Multiscale recordings reveal the dynamic spatial structure of human seizures. *Neurobiol. Dis.* 127:303–11
43. Patel KV, Katz CN, Kalia SK, Popovic MR, Valiante TA. 2020. Volitional control of individual neurons in the human brain. bioRxiv 2020.05.05.079038. <https://doi.org/10.1101/2020.05.05.079038>
44. Chung JE, Joo HR, Fan JL, Liu DF, Barnett AH, et al. 2019. High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays. *Neuron* 101:21–31
45. Musk E, Neuralink. 2019. An integrated brain-machine interface platform with thousands of channels. *J. Med. Internet Res.* 21:e16194
46. Chiang CH, Won SM, Orsborn AL, Yu KJ, Trumpis M, et al. 2020. Development of a neural interface for high-definition, long-term recording in rodents and nonhuman primates. *Sci. Transl. Med.* 12:eaay4682
47. Simeral JD, Hosman T, Saab J, Flesher SN, Vilela M, et al. 2019. Home use of a wireless intracortical brain-computer interface by individuals with tetraplegia. medRxiv 2019.12.27.19015727. <https://doi.org/10.1101/2019.12.27.19015727>
48. Perge JA, Homer ML, Malik WQ, Cash S, Eskandar E, et al. 2013. Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *J. Neural Eng.* 10:036004
49. Downey JE, Schwed N, Chase SM, Schwartz AB, Collinger JL. 2018. Intracortical recording stability in human brain–computer interface users. *J. Neural Eng.* 15:046016
50. Perel S, Sadtler PT, Oby ER, Ryu SI, Tyler-Kabara EC, et al. 2015. Single-unit activity, threshold crossings, and local field potentials in motor cortex differentially encode reach kinematics. *J. Neurophysiol.* 114:1500–12
51. Christie BP, Tat DM, Irwin ZT, Gilja V, Nuyujukian P, et al. 2015. Comparison of spike sorting and thresholding of voltage waveforms for intracortical brain–machine interface performance. *J. Neural Eng.* 12:016009
52. Flint RD, Scheid MR, Wright ZA, Solla SA, Slutzky MW. 2016. Long-term stability of motor cortical activity: implications for brain machine interfaces and optimal feedback control. *J. Neurosci.* 36:3623–32
53. Jackson A, Hall TM. 2017. Decoding local field potentials for neural interfaces. *IEEE Trans. Neural Syst. Rehabil. Eng.* 25:1705–14



54. Milekovic T, Sarma AA, Bacher D, Simeral JD, Saab J, et al. 2018. Stable long-term BCI-enabled communication in ALS and locked-in syndrome using LFP signals. *J. Neurophysiol.* 120:343–60
55. Flint RD, Lindberg EW, Jordan LR, Miller LE, Slutzky MW. 2012. Accurate decoding of reaching movements from field potentials in the absence of spikes. *J. Neural Eng.* 9:046006
56. Nason SR, Vaskov AK, Willsey MS, Welle EJ, An H, et al. 2020. A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain-machine interfaces. *Nat. Biomed. Eng.* In press. <https://doi.org/10.1101/2019.12.27.19015727>
57. Pachitariu M, Stringer C, Dipoppa M, Schröder S, Rossi LF, et al. 2017. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. bioRxiv 061507. <https://doi.org/10.1101/061507>
58. Chen TW, Wardill TJ, Sun Y, Pulver SR, Renninger SL, et al. 2013. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499:295–300
59. Clancy KB, Koralek AC, Costa RM, Feldman DE, Carmena JM. 2014. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* 17:807–9
60. Liberti W, Gong XL, Roseberry T, Costa RM, Carmena JM. 2019. Local network coordination supports neuroprosthetic control. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*, pp. 558–61. Piscataway, NJ: IEEE
61. Trautmann EM, O’Shea DJ, Sun X, Marshel JH, Crow A, et al. 2019. Dendritic calcium signals in rhesus macaque motor cortex drive an optical brain-computer interface. bioRxiv 780486. <https://doi.org/10.1101/780486>
62. Bollimunta A, Santacruz SR, Eaton RW, Xu PS, Morrison JH, et al. 2020. Head-mounted microendoscopic calcium imaging in dorsal premotor cortex of behaving rhesus macaque. bioRxiv 2020.04.10.996116. <https://doi.org/10.1101/2020.04.10.996116>
63. Kleinlogel S, Vogl C, Jeschke M, Neef J, Moser T. 2020. Emerging approaches for restoration of hearing and vision. *Physiol. Rev.* 100:1467–525
64. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. 2008. Cortical control of a prosthetic arm for self-feeding. *Nature* 453:1098–101
65. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, et al. 2012. Neural population dynamics during reaching. *Nature* 487:51–56
66. Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE. 2020. Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.* 23:260–70
67. Rubin A, Sheintuch L, Brande-Eilat N, Pinchasof O, Rechavi Y, et al. 2019. Revealing neural correlates of behavior without behavioral measurements. *Nat. Commun.* 10:4745
68. Bacher D, Jarosiewicz B, Masse NY, Stavisky SD, Simeral JD, et al. 2015. Neural point-and-click communication by a person with incomplete locked-in syndrome. *Neurorehabil. Neural Repair* 29:462–71
69. Kim SP, Simeral JD, Hochberg LR, Donoghue JP, Black MJ. 2008. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J. Neural Eng.* 5:455–76
70. Malik WQ, Truccolo W, Brown EN, Hochberg LR. 2011. Efficient decoding with steady-state kalman filter in neural interface systems. *IEEE Trans. Neural Syst. Rehabil. Eng.* 19:25–34
71. Inoue Y, Mao H, Suway SB, Orellana J, Schwartz AB. 2018. Decoding arm speed during reaching. *Nat. Commun.* 9:5243
72. Li S, Li J, Li Z. 2016. An improved unscented Kalman filter based decoder for cortical brain-machine interfaces. *Front. Neurosci.* 10:587
73. Brandman DM, Burkhart MC, Kelemen J, Franco B, Harrison MT, Hochberg LR. 2018. Robust closed-loop control of a cursor in a person with tetraplegia using Gaussian process regression. *Neural Comput.* 30:2986–3008
74. Burkhart MC, Brandman DM, Franco B, Hochberg LR, Harrison MT. 2020. The discriminative Kalman filter for Bayesian filtering with nonlinear and nongaussian observation models. *Neural Comput.* 32:969–1017
75. Hosman T, Vilela M, Milstein D, Kelemen JN, Brandman DM, et al. 2019. BCI decoder performance comparison of an LSTM recurrent neural network and a Kalman filter in retrospective simulation. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*, pp. 1066–71. Piscataway, NJ: IEEE
76. Tseng PH, Urpi NA, Lebedev M, Nicolelis M. 2019. Decoding movements from cortical ensemble activity using a long short-term memory recurrent network. *Neural Comput.* 31:1085–113



77. Sussillo D, Nuyujukian P, Fan JM, Kao JC, Stavisky SD, et al. 2012. A recurrent neural network for closed-loop intracortical brain-machine interface decoders. *J. Neural Eng.* 9:026027
78. Vargas-Irwin CE, Feldman JM, King B, Simeral JD, Sorice BL, et al. 2018. Watch, imagine, attempt: Motor cortex single-unit activity reveals context-dependent movement encoding in humans with tetraplegia. *Front. Hum. Neurosci.* 12:450
79. Rastogi A, Vargas-Irwin CE, Willett FR, Abreu J, Crowder DC, et al. 2020. Neural representation of observed, imagined, and attempted grasping force in motor cortex of individuals with chronic tetraplegia. *Sci. Rep.* 10:1429
80. Hatsopoulos N, Joshi J, O'Leary JG. 2004. Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *J. Neurophysiol.* 92:1165–74
81. Willett FR, Young DR, Murphy BA, Memberg WD, Blabe CH, et al. 2019. Principled BCI decoder design and parameter selection using a feedback control model. *Sci. Rep.* 9:8881
82. Chase SM, Schwartz AB, Kass RE. 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.* 22:1203–13
83. Koyama S, Chase SM, Whitford AS, Velliste M, Schwartz AB, Kass RE. 2010. Comparison of brain-computer interface decoding algorithms in open-loop and closed-loop control. *J. Comput. Neurosci.* 29:73–87
84. Taylor DM, Helms Tillery SI, Schwartz AB. 2002. Direct cortical control of 3D neuroprosthetic devices. *Science* 296:1829–32
85. Gowda S, Orsborn AL, Overduin SA, Moorman HG, Carmena JM. 2014. Designing dynamical properties of brain-machine interfaces to optimize task-specific performance. *IEEE Trans. Neural Syst. Rehabil. Eng.* 22:911–20
86. Shanechi MM, Orsborn AL, Moorman HG, Gowda S, Dangi S, Carmena JM. 2017. Rapid control and feedback rates enhance neuroprosthetic control. *Nat. Commun.* 8:13825
87. Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, et al. 2012. A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* 15:1752–57
88. Willett FR, Murphy BA, Young D, Memberg WD, Blabe CH, et al. 2018. A comparison of intention estimation methods for decoder calibration in intracortical brain-computer interfaces. *IEEE Trans. Biomed. Eng.* 65:2066–78
89. Jarosiewicz B, Sarma AA, Bacher D, Masse NY, Simeral JD, et al. 2015. Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Sci. Transl. Med.* 7:313ra179
90. Willett FR, Pandarinath C, Jarosiewicz B, Murphy BA, Memberg WD, et al. 2017. Feedback control policies employed by people using intracortical brain-computer interfaces. *J. Neural Eng.* 14:016001
91. Todorov E. 2004. Optimality principles in sensorimotor control. *Nat. Neurosci.* 7:907–15
92. Shanechi MM, Orsborn AL, Carmena JM. 2016. Robust brain-machine interface design using optimal feedback control modeling and adaptive point process filtering. *PLOS Comput. Biol.* 12:e1004730
93. Benyamini M, Zacksenhouse M. 2015. Optimal feedback control successfully explains changes in neural modulations during experiments with brain-machine interfaces. *Front. Syst. Neurosci.* 9:71
94. Lagang M, Srinivasan L. 2013. Stochastic optimal control as a theory of brain-machine interface operation. *Neural Comput.* 25:374–417
95. Godlove JM, Whaithe EO, Batista AP. 2014. Comparing temporal aspects of visual, tactile, and microstimulation feedback for motor control. *J. Neural Eng.* 11:046025
96. Sombeck JT, Miller LE. 2019. Short reaction times in response to multi-electrode intracortical microstimulation may provide a basis for rapid movement-related feedback. *J. Neural Eng.* 17:016013
97. Weiss JM, Flesher SN, Franklin R, Collinger JL, Gaunt RA. 2018. Artifact-free recordings in human bidirectional brain-computer interfaces. *J. Neural Eng.* 16:016002
98. Ganzer PD, Colachis SC, Schwemmer MA, Friedenber DA, Dunlap CF, et al. 2020. Restoring the sense of touch using a sensorimotor demultiplexing neural interface. *Cell* 181:763–73.e12
99. Prsa M, Galiñanes GL, Huber D. 2017. Rapid integration of artificial sensory feedback during operant conditioning of motor cortex neurons. *Neuron* 93:929–39.e6
100. Abbasi A, Goueytes D, Shulz DE, Ego-Stengel V, Estebanez L. 2018. A fast intracortical brain-machine interface with patterned optogenetic feedback. *J. Neural Eng.* 15:046011



101. Lu Y, Truccolo W, Wagner FB, Vargas-Irwin CE, Ozden I, et al. 2015. Optogenetically induced spatiotemporal gamma oscillations and neuronal spiking activity in primate motor cortex. *J. Neurophysiol.* 113:3574–87
102. Yazdan-Shahmorad A, Silversmith DB, Kharazia V, Sabes PN. 2018. Targeted cortical reorganization using optogenetics in non-human primates. *eLife* 7:e31034
103. Hotson G, Smith RJ, Rouse AG, Schieber MH, Thakor NV, Wester BA. 2016. High precision neural decoding of complex movement trajectories using recursive bayesian estimation with dynamic movement primitives. *IEEE Robot. Autom. Lett.* 1:676–83
104. Moritz CT, Fetz EE. 2011. Volitional control of single cortical neurons in a brain-machine interface. *J. Neural Eng.* 8:025017
105. Neely RM, Koralek AC, Athalye VR, Costa RM, Carmena JM. 2018. Volitional modulation of primary visual cortex activity requires the basal ganglia. *Neuron* 97:1356–68.e4
106. Law AJ, Rivlis G, Schieber MH. 2014. Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. *J. Neurophysiol.* 112:1528–48
107. Zhou X, Tien RN, Ravikumar S, Chase SM. 2019. Distinct types of neural reorganization during long-term learning. *J. Neurophysiol.* 121:1329–41
108. You AK, Liu B, Singhal A, Gowda S, Moorman H, et al. 2019. Flexible modulation of neural variance facilitates neuroprosthetic skill learning. bioRxiv 817346. <https://doi.org/10.1101/817346>
109. Pierella C, Casadio M, Mussa-Ivaldi FA, Solla SA. 2019. The dynamics of motor learning through the formation of internal models. *PLOS Comput. Biol.* 15:e1007118
110. Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DV, et al. 2003. Learning to control a brain-machine interface for reaching and grasping by primates. *PLOS Biol.* 1:e42
111. Wander JD, Blakely T, Miller KJ, Weaver KE, Johnson LA, et al. 2013. Distributed cortical adaptation during learning of a brain-computer interface task. *PNAS* 110:10818–23
112. Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, et al. 2014. Neural constraints on learning. *Nature* 512:423–26
113. Feulner B, Clopath C. 2020. Neural manifold under plasticity in a goal driven learning behaviour. bioRxiv 2020.02.21.959163. <https://doi.org/10.1101/2020.02.21.959163>
114. Sakellaridi S, Christopoulos VN, Aflalo T, Pejsa KW, Rosario ER, et al. 2019. Intrinsic variable learning for brain-machine interface control by human anterior intraparietal cortex. *Neuron* 102:694–705.e3
115. Perich MG, Miller LE. 2017. Altered tuning in primary motor cortex does not account for behavioral adaptation during force field learning. *Exp. Brain Res.* 235:2689–704
116. Perich MG, Gallego JA, Miller LE. 2018. A neural population mechanism for rapid learning. *Neuron* 100:964–976.e7
117. Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, et al. 2018. Learning by neural reassociation. *Nat. Neurosci.* 21:607–16
118. Oby ER, Golub MD, Hennig JA, Degenhart AD, Tyler-Kabara EC, et al. 2019. New neural activity patterns emerge with long-term learning. *PNAS* 116:15210–15
119. Ganguly K, Carmena JM. 2009. Emergence of a stable cortical map for neuroprosthetic control. *PLOS Biol.* 7:e1000153
120. You A, Zippi EL, Carmena JM. 2019. Large-scale neural consolidation in BMI learning. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*, pp. 603–6. Piscataway, NJ: IEEE
121. Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. 2008. Functional network reorganization during learning in a brain-computer interface paradigm. *PNAS* 105:19486–91
122. Zhang Y, Chase SM. 2018. Optimizing the usability of brain-computer interfaces. *Neural Comput.* 30:1323–58
123. Oweiss KG, Badreldin IS. 2015. Neuroplasticity subserving the operation of brain-machine interfaces. *Neurobiol. Dis.* 83:161–71
124. Driscoll LN, Pettit NL, Minderer M, Chettih SN, Harvey CD. 2017. Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170:986–999.e16
125. Rubin A, Geva N, Sheintuch L, Ziv Y. 2015. Hippocampal ensemble dynamics timestamp events in long-term memory. *eLife* 4:e12247



126. Dhawale AK, Poddar R, Wolff SB, Normand VA, Kopelowitz E, Ölveczky BP. 2017. Automated long-term recording and analysis of neural activity in behaving animals. *eLife* 6:e27702
127. Rule ME, O'Leary T, Harvey CD. 2019. Causes and consequences of representational drift. *Curr. Opin. Neurobiol.* 58:141–47
128. You A, Singhal A, Moorman H, Gowda S, Carmena JM. 2019. Neural correlates of control of a kinematically redundant brain-machine interface. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*, pp. 554–57. Piscataway, NJ: IEEE
129. Ganguly K, Dimitrov DF, Wallis JD, Carmena JM. 2011. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* 14:662–67
130. Lansdell B, Milovanovic I, Mellema C, Fetz EE, Fairhall AL, Moritz CT. 2020. Reconfiguring motor circuits for a joint manual and BCI task. *IEEE Trans. Neural Syst. Rehabil. Eng.* 28:248–57
131. Lalazar H, Murray J, Abbott L, Vaadia E. 2019. Population subspaces reflect movement intention for arm and brain-machine interface control. bioRxiv 688259. <https://doi.org/10.1101/688259>
132. Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MAL. 2005. Stable ensemble performance with single-neuron variability during reaching movements in primates. *J. Neurosci.* 25:10712–16
133. Sussillo D, Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2016. Making brain-machine interfaces robust to future neural variability. *Nat. Commun.* 7:13749
134. Rule ME, Loback AR, Raman DV, Driscoll LN, Harvey CD, O'Leary T. 2020. Stable task information from an unstable neural population. *eLife* 9:e51121
135. Orsborn A, Moorman H, Overduin S, Shanechi M, Dimitrov D, Carmena J. 2014. Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron* 82:1380–93
136. Dangi S, Gowda S, Moorman HG, Orsborn AL, So K, et al. 2014. Continuous closed-loop decoder adaptation with a recursive maximum likelihood algorithm allows for rapid performance acquisition in brain-machine interfaces. *Neural Comput.* 26:1811–39
137. Merel J, Pianto DM, Cunningham JP, Paninski L. 2015. Encoder-decoder optimization for brain-computer interfaces. *PLoS Comput. Biol.* 11:e1004288
138. Pohlmeier EA, Mahmoudi B, Geng S, Prins NW, Sanchez JC. 2014. Using reinforcement learning to provide stable brain-machine interface control despite neural input reorganization. *PLoS ONE* 9:e87253
139. Zhang P, Ma X, Chen L, Zhou J, Wang C, et al. 2018. Decoder calibration with ultra small current sample set for intracortical brain-machine interface. *J. Neural Eng.* 15:026019
140. Brandman DM, Hosman T, Saab J, Burkhart MC, Shanahan BE, et al. 2018. Rapid calibration of an intracortical brain computer interface for people with tetraplegia. *J. Neural Eng.* 15:026007
141. Xing D, Aghagolzadeh M, Truccolo W, Borton D. 2019. Low-dimensional motor cortex dynamics preserve kinematics information during unconstrained locomotion in nonhuman primates. *Front. Neurosci.* 13:1046
142. Aghagolzadeh M, Truccolo W. 2016. Inference and decoding of motor cortex low-dimensional dynamics via latent state-space models. *IEEE Trans. Neural Syst. Rehabil. Eng.* 24:272–82
143. Dyer EL, Gheshlaghi Azar M, Perich MG, Fernandes HL, Naufel S, et al. 2017. A cryptography-based approach for movement decoding. *Nat. Biomed. Eng.* 1:967–76
144. Kao JC, Ryu SI, Shenoy KV. 2017. Leveraging neural dynamics to extend functional lifetime of brain-machine interfaces. *Sci. Rep.* 7:7395
145. Degenhart AD, Bishop WE, Oby ER, Tyler-Kabara EC, Chase SM, et al. 2020. Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. *Nat. Biomed. Eng.* 4:672–85
146. Trautmann EM, Stavisky SD, Lahiri S, Ames KC, Kaufman MT, et al. 2019. Accurate estimation of neural population dynamics without spike sorting. *Neuron* 103:292–308.e4
147. Milstein DJ, Pacheco JL, Hochberg LR, Simeral JD, Jarosiewicz B, Sudderth EB. 2017. Multiscale semi-Markov dynamics for intracortical brain-computer interfaces. In *Advances in Neural Information Processing Systems 30*, ed. I Guyon, UV Luxburg, S Bengio, H Wallach, R Fergus, et al., pp. 868–78. Red Hook, NY: Curran
148. Li H, Hao Y, Zhang S, Wang Y, Chen W, Zheng X. 2017. Prior knowledge of target direction and intended movement selection improves indirect reaching movement decoding. *Behav. Neurol.* 2017:2182843



149. Shانهchi MM, Williams ZM, Wornell GW, Hu RC, Powers M, Brown EN. 2013. A real-time brain-machine interface combining motor target and trajectory intent using an optimal feedback control design. *PLOS ONE* 8:e59049
150. Benyamini M, Nason SR, Chestek CA, Zacksenhouse M. 2019. Neural correlates of error processing during grasping with invasive brain-machine interfaces. In *2019 9th International IEEE/EMBS Conference on Neural Engineering*, pp. 215–18. Piscataway, NJ: IEEE
151. Even-Chen N, Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2017. Augmenting intracortical brain-machine interface with neurally driven error detectors. *J. Neural Eng.* 14:066007
152. Schroeder KE, Irwin ZT, Bullard AJ, Thompson DE, Bentley JN, et al. 2017. Robust tactile sensory responses in finger area of primate motor cortex relevant to prosthetic control. *J. Neural Eng.* 14:046016
153. Stavisky SD, Kao JC, Ryu SI, Shenoy KV. 2017. Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. *Neuron* 95:195–208.e9
154. Ramakrishnan A, Byun YW, Rand K, Pedersen CE, Lebedev MA, Nicolelis MA. 2017. Cortical neurons multiplex reward-related signals along with sensory and motor information. *PNAS* 114:E4841–50
155. Zhao Y, Hessburg JP, Asok Kumar JN, Francis JT. 2018. Paradigm shift in sensorimotor control research and brain machine interface control: the influence of context on sensorimotor representations. *Front. Neurosci.* 12:579
156. Downey JE, Brane L, Gaunt RA, Tyler-Kabara EC, Boninger ML, Collinger JL. 2017. Motor cortical activity changes during neuroprosthetic-controlled object interaction. *Sci. Rep.* 7:16947
157. Schroeder KE, Perkins SM, Wang Q, Churchland MM. 2020. Neural control of virtual ego-motion enabled by an opportunistic decoding strategy. *bioRxiv* 2019.12.13.862532. <https://doi.org/10.1101/2019.12.13.862532>
158. Sheintuch L, Geva N, Baumer H, Rechavi Y, Rubin A, Ziv Y. 2020. Multiple maps of the same spatial context can stably coexist in the mouse hippocampus. *Curr. Biol.* 30:1467–76.e6
159. Gallego JA, Perich MG, Naufel SN, Ethier C, Solla SA, Miller LE. 2018. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat. Commun.* 9:4233
160. Sumsy SL, Schieber MH, Thakor NV, Sarma SV, Santaniello S. 2017. Decoding kinematics using task-independent movement-phase-specific encoding models. *IEEE Trans. Neural Syst. Rehabil. Eng.* 25:2122–32
161. Williams JJ, Tien RN, Inoue Y, Schwartz AB. 2016. Idle state classification using spiking activity and local field potentials in a brain computer interface. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1572–75. Piscataway, NJ: IEEE
162. Sachs NA, Ruiz-Torres R, Perreault EJ, Miller LE. 2016. Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. *J. Neural Eng.* 13:016009
163. Borton DA, Yin M, Aceros J, Nurmikko A. 2013. An implantable wireless neural interface for recording cortical circuit dynamics in moving primates. *J. Neural Eng.* 10:026010
164. Weiss JM, Gaunt RA, Franklin R, Boninger ML, Collinger JL. 2019. Demonstration of a portable intracortical brain-computer interface. *Brain-Comput. Interfaces* 6:106–17
165. Gilron R, Little S, Perrone R, Wilt R, de Hemptinne C, et al. 2020. Chronic wireless streaming of invasive neural recordings at home for circuit discovery and adaptive stimulation. *bioRxiv* 2020.02.13.948349. <https://doi.org/10.1101/2020.02.13.948349>
166. Jorge A, Royston DA, Tyler-Kabara EC, Boninger ML, Collinger JL. 2020. Classification of individual finger movements using intracortical recordings in human motor cortex. *Neurosurgery*. In press. <https://doi.org/10.1093/neuros/nyaa026>
167. Wilson GH, Stavisky SD, Willett FR, Avansino DT, Kelemen JN, et al. 2020. Decoding spoken English phonemes from intracortical electrode arrays in dorsal precentral gyrus. *bioRxiv* 2020.06.30.180935. <https://doi.org/10.1101/2020.06.30.180935>
168. Zhang Z, Russell LE, Packer AM, Gauld OM, Häusser M. 2018. Closed-loop all-optical interrogation of neural circuits in vivo. *Nat. Methods* 15:1037–40
169. Mitani A, Dong M, Komiyama T. 2018. Brain-computer interface with inhibitory neurons reveals subtype-specific strategies. *Curr. Biol.* 28:77–83

