CellPress

Trends in Cognitive Sciences

Review

Brain–computer interfaces as a causal probe for scientific inquiry

Asma Motiwala[1,4,5,6], Joana Soldado-Magraner[1,4,5,6], Aaron P. Batista[3,4,7], Matthew A. Smith[2,4,5,7], and Byron M. Yu[1,2,4,5,7,*]

Establishing causal relationships between neural activity and brain function requires experimental perturbations of neural activity. Many existing perturbation methods modify activity by directly applying external signals to the brain. We review an alternative approach where brain–computer interfaces (BCIs) leverage volitional control of neural activity to manipulate and causally perturb it. We highlight the potential of BCIs to manipulate neural activity in ways that are flexible, accurate, and adhere to intrinsic biophysical and network-level constraints to investigate the consequences of configuring neural population activity in specified ways. We discuss the advantages and disadvantages of using BCIs as a perturbation tool compared with other perturbation methods and how BCIs can expand the scope of questions that can be addressed about brain function.

Highlights

Causal manipulations of neural activity can go beyond correlational studies in establishing how neural activity relates to neural computations and behavior.

Multiple neural perturbation tools exist, which directly manipulate neural activity by using forces external to the brain.

Brain–computer interfaces (BCIs) provide a causal tool whereby animals perturb their own neural activity using volition – a force internal to the brain.

We can assess perturbation tools, including BCIs, in terms of the flexibility, accuracy, and naturalness of population activity patterns that they elicit.

Causal manipulations of neural activity link activity to behavior

How neural activity implements neural computations and gives rise to behavior is a central question in neuroscience. We often observe neural activity while animals perform different behaviors, which can uncover **correlational relationships** (see Glossary) between the neural activity and behavior. When neural activity is controlled by the experimenter while an animal performs behaviors of interest, we can establish **causal relationships** between neural activity and behavior [1].

A wide range of tools are available to directly manipulate neural activity. These include intracortical electrical microstimulation [2,3]; transcranial electrical stimulation (TES) [4,5]; transcranial magnetic stimulation (TMS) [6,7]; optogenetics [8,9]; pharmacology, such as cortical inactivation using muscimol [10,11]; pharmacogenetics, such as designer receptors exclusively activated by designer drugs (DREADDs) [12,13]; ultrasound [14,15]; lesions [16,17]; and cortical cooling [18,19]. These tools differ in many aspects, such as the spatial and temporal scale at which they manipulate neural activity, the biological mechanism through which they act, and the specificity with which they can target different circuits. A key characteristic shared by all these perturbation methods is that they exert external control over neural activity, relying on forces or interventions originating outside the brain itself.

In this review, we highlight an alternative way to causally perturb neural activity: using a BCI to leverage an animal's control of their own neural activity. In this way, neural activity is controlled by volitional forces internal to the brain. Here we discuss how volitional control can modulate neural activity in ways that are complementary to other perturbation tools. By considering BCIs as an additional tool to manipulate neural activity, we can open new avenues to probe neural circuits and advance our understanding of the brain.

BCIs for manipulating neural activity to investigate neural circuit function

A BCI – also known as a brain–machine interface (BMI) – is a system that translates neural activity into movements of an external **effector**, such as a computer cursor or a robotic limb. BCIs

[1]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[2]Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[3]Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA 15213, USA
[4]Center for the Neural Basis of Cognition, Carnegie Mellon University and University of Pittsburgh, Pittsburgh, PA 15213, USA
[5]Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[6]These authors contributed equally to this work
[7]These authors contributed equally to this work

*Correspondence: byronyu@cmu.edu (B.M. Yu).

consist of three primary components: (i) measurement of neural activity; (ii) a mathematical mapping of the neural activity to desired movements (referred to as a 'BCI mapping'); and (iii) an external effector whose movement is specified by the neural activity (Figure 1A) [20]. Neural activity used for BCIs can range from intracortical measurements of neuronal firing (e.g., using surgically implanted multi-electrode arrays) to non-invasive measurements of neural activity [e.g., using electroencephalography (EEG)]. The BCI mapping then converts the recorded activity into movements of the external effector, which provides the user with moment-by-moment sensory feedback of their neural activity. This **closed-loop system** allows users to volitionally modulate their neural activity to achieve task goals (Figure 1B). Laboratory proof-of-concept BCIs [21–25] have informed the development of BCIs used in clinical settings to restore motor and communication capabilities in individuals with paralysis or other disorders [26–34]. BCIs are also referred to as neurofeedback systems when they are used in settings where self-regulation of brain signals, typically measured non-invasively, is leveraged for therapeutic purposes [35–41].

In addition to their use for clinical applications, BCIs have also been used as a tool to investigate basic scientific questions about neural circuit function [42–48]. BCIs allow us to simplify the output interface of the brain, while preserving the cognitive processes involved in sensorimotor control [46]. Because experimenters define the BCI mapping, the causal relationship between the

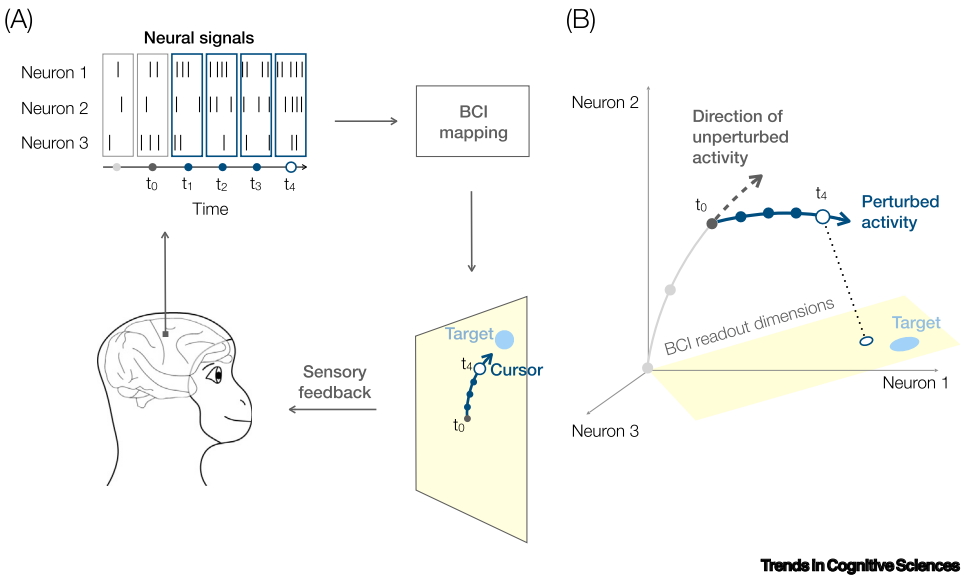

Figure 1. Brain–computer interfaces (BCIs) as a tool to manipulate neural activity. (A) Schematic of a BCI in which neural activity is recorded from a population of neurons and spike counts are taken in sequential time bins (illustrated here using three neurons; top left). A BCI mapping translates the spike counts into movements of a computer cursor (bottom right, 'BCI workspace'). Moment-by-moment visual feedback is provided to the animal. It is also possible to use other sensory feedback modalities, such as auditory, proprioceptive or tactile feedback, or a combination [53,114]. The feedback enables the animal to modify its neural activity in real time to guide the computer cursor to an on-screen target (light blue circle). (B) We define a population activity space, in which each axis represents the activity of one recorded neuron. Neural activity at different times constitutes a sequence of points in the population activity space, forming a neural trajectory. The solid gray trajectory represents neural activity prior to BCI control. The blue trajectory represents volitional modulation of neural activity to drive the BCI cursor to the target (light blue circle) under moment-by-moment sensory feedback. The perturbed activity (blue trajectory) deviates from the activity that would have been observed if left unperturbed (broken gray trajectory). The yellow plane represents the dimensions along which neural activity is read out by the BCI mapping. The current population activity state (hollow circle along the blue trajectory) is mapped (along the broken line) to cursor kinematics (hollow circle in the yellow plane).

Glossary

Causal relationship: statistical dependency between variables that describes which variables directly influence others. This is typically revealed by an experimenter-induced manipulation.

Closed-loop system: a system that receives control inputs at each time step that are informed by recent observations (i.e., feedback) of the state of the system.

Correlational relationship: statistical dependency between variables for which it is not known which variables directly influence others. This is typically revealed by only observing a system without intervening in it.

Effector: a body part (e.g., a limb) or a mechanical device (e.g., a robotic arm or a computer cursor) that moves in response to control commands, such as neural activity.

Targeted neural population: a neural population whose activity is directly influenced by a perturbation. In the case of external perturbation methods, the targeted population refers to the set of neurons that the external signals impinge upon. In the case of BCIs, the targeted population refers to the recorded neurons that drive BCI readouts.

recorded neural activity and effector (e.g., computer cursor) movements is known. This fundamental property of BCIs has enabled insights into the extent and limits of volitional modulation of neural activity [49–54], how those limits can change with repeated practice [55–57] and sleep [58], how volitional modulation relies on sensory feedback [59,60], and how volitional modulation relies on multiple brain areas [61,62].

This body of work shows how BCIs can be used as a tool to causally manipulate neural activity. The goal in many of the aforementioned studies has been to investigate changes in neural population activity during BCI control. These studies demonstrate how the specifications of different BCI components and task goals directly influence the requirements for volitional modulation of neural activity to achieve task success, thereby exerting a causal influence on neural activity (Figure 1B and Box 1).

Insights into how neural activity can be manipulated using BCIs have laid the groundwork for using BCIs to investigate the consequences of specific perturbations of neural activity. This shifts the focus of investigation from the neural mechanisms underlying activity changes during BCI control to the subsequent consequences of having intentionally changed neural activity. These two uses for BCIs for basic science – for discovering the limits of volitional activity modulation and for perturbing neural activity – are synergistic. Understanding the capabilities of volitional activity modulation informs the manner in which BCIs can be used for neural activity perturbations. The use of BCIs to perturb brain signals has been extensively explored at the scale of brain-wide signals such as those measured using fMRI, EEG, and other non-invasive approaches [63]. However, at the scale of populations of individual neurons, the use of BCIs as tools to perturb neural activity has been less common. Studies that have used BCIs in this way have advanced our

Box 1. Design choices for manipulating neural activity using a BCI

The design of each BCI component can be tailored to require specific patterns of neural activity for successful BCI control. We explain this in the context of two key BCI components (highlighted in Figure IA).

(i) Readout dimensions. The BCI mapping defines what aspects of the recorded neural activity drive BCI effector movements and what aspects are ignored. When this mapping is linear, the readout operations can be represented as projections of neural activity onto specific dimensions of the population activity space. Typically, one dimension is specified for each degree of freedom of the effector to be controlled. For example, to control a computer cursor on a 2D screen, two readout dimensions are specified. Changes in the neural activity along these two dimensions map to the movement of the cursor along the horizontal and vertical coordinates of the screen. By varying the readout dimensions, one can constrain the elicited population activity patterns along specific dimensions. However, the evolution of population activity outside the readout dimensions is left unconstrained (compare Figure IB and C).

(ii) Control of cursor kinematics. The kinematic variables to which neural activity is mapped can be chosen by the experimenter. Each choice imposes distinct requirements on the neural activity that would yield successful BCI control. For example, the BCI mapping might drive the cursor's position, velocity, acceleration, or a combination of these. A center-out cursor movement using a position mapping sets a requirement on the direction in which population activity needs to be modulated and how far from the baseline level that neural activity needs to be at the time of target acquisition (Figure ID). By contrast, mapping neural activity to cursor velocity sets requirements only on the direction in which population activity must be modulated, but does not tightly constrain how neural activity changes from time step to time step. The degree to which neural activity is modulated away from the baseline level controls cursor speed, and cursor position is determined by integrating cursor velocity over time. In the same center-out task, using a velocity mapping, successful cursor control may be possible even if, at the time of target acquisition, neural activity has returned close to the baseline level (Figure IE).

Among other BCI components, the chosen sensory feedback modality (e.g., visual, auditory, proprioceptive, or tactile) should be appropriate for the known properties of the targeted neural population and the scientific question being asked. For example, if the targeted population responds directly to visual stimuli, then non-visual feedback should be used, or visual feedback should be designed carefully to avoid confounding aspects of the neural activity that arise due to volition versus being driven by a sensory stimulus [53].

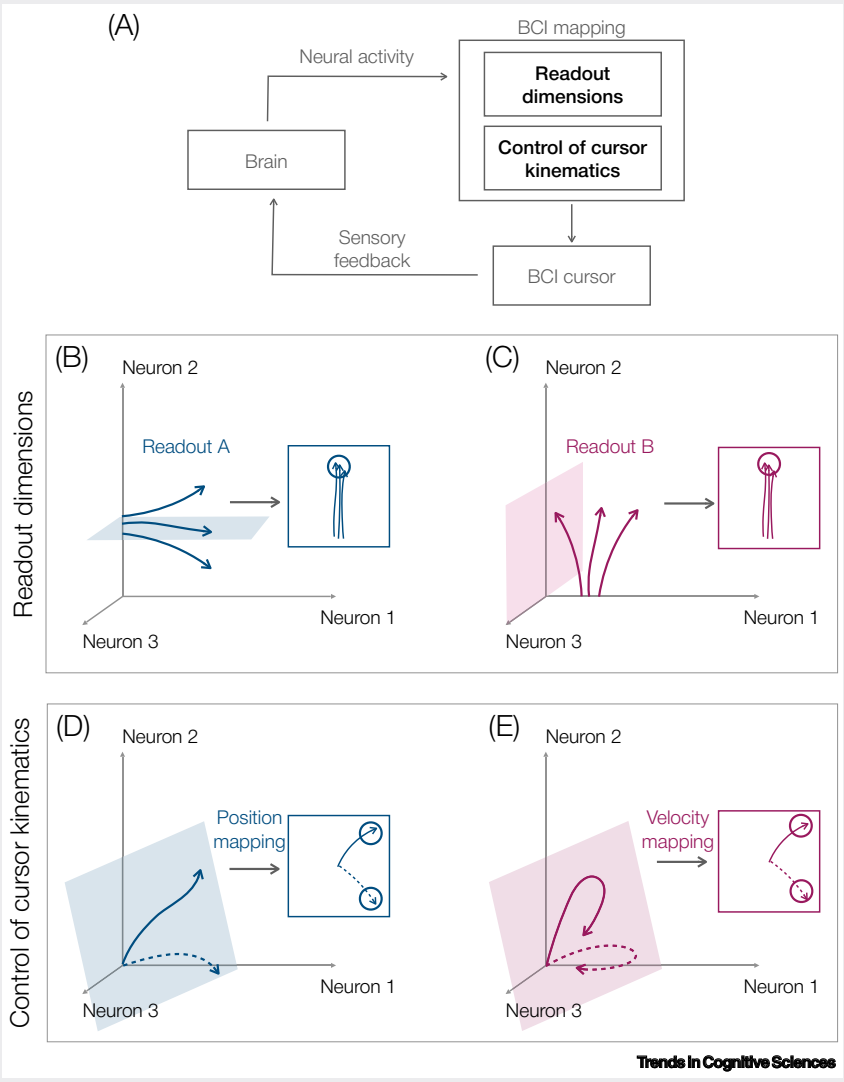

Figure I. Implications of brain–computer interface (BCI) design choices on neural activity perturbations. (A) Schematic of a closed-loop BCI system. We illustrate how different design choices for the readout dimensions and the control of cursor kinematics impose different requirements on neural activity. (B–C) Readout dimensions. (B) Readout A places requirements on the neural activity along the blue readout dimensions, but allows activity to vary in dimensions orthogonal to them. Each neural trajectory in the population activity space represents an experimental trial. Three trials for the same BCI target are shown. The shaded plane represents the readout dimensions of the BCI mapping. The neural activity is converted by the BCI mapping into cursor trajectories (shown in the square inset, representing the BCI workspace). For each cursor trajectory, there is a corresponding neural trajectory. Circle represents the BCI target; (C) Readout B places requirements on the neural activity along the red readout dimensions (orthogonal to readout A). The requirements placed on the BCI cursor are the same for the two readouts. Same conventions as panel B. (D–E) Control of cursor kinematics. (D) A position-based mapping directly translates activity along readout dimensions to cursor position. This allows experimenters to specify requirements on the path of the cursor on the screen to set requirements on the path of neural activity in the readout dimensions. Same conventions as panel B, but here we show two BCI targets and one trajectory (solid and dashed lines) for each BCI target. (E) A velocity-based mapping translates activity along the readout dimensions to cursor velocity. This allows experimenters to elicit neural activity consistent with moving the cursor towards the target, whatever the speed of the cursor. Same conventions as panel D.

understanding of cognitive and motor functions, including spatial attention [64], motor learning transfer [65], memory [58,66], and motor planning [67].

Compared with other perturbation methods, BCIs stand out as a distinct perturbation tool due to a fundamental difference in the mechanism by which they induce changes in neural activity. Any perturbation of neural activity can be viewed as a controlled force applied on a **targeted neural population**. When using perturbation methods such as electrical or optogenetic stimulation, the targeted population comprises the neurons that respond directly to the external stimulation applied. In this case, the force of change that triggers neural activity perturbations in the targeted population is generated externally and controlled by the experimenter (Figure 2A, blue arrow). By contrast, when using BCIs as a perturbation tool, the population that is read out for BCI control is considered the targeted population, because task success is a function of the activity in this population. As a result, the BCI task requirements influence how activity in this population is modulated. These activity changes are initiated by recruiting upstream circuits that drive the targeted population. The force of change that alters neural activity is thus generated internally within the brain by leveraging an animal's own volitional control (Figure 2B, blue arrows). The reliance on the animal's volition underlies the unique properties of BCIs as tools to perturb neural activity. Using either external perturbation methods or BCIs, subsequent changes in activity would be observed in circuits downstream to and/or recurrently connected with the targeted population.

Properties of BCIs as a perturbation tool

In the following sections, we identify three desirable properties of perturbation tools and describe the benefits of using BCIs to perturb neural activity in terms of these properties. To do so, we define a multi-dimensional population activity space (also referred to as a 'state space'), in which each axis represents the activity of a recorded neuron (Figure 1B). Within this state space, we can characterize how the activity of recorded neurons covaries and how neurons change their activity together over time [68–71]. Any perturbation method can be conceptualized as a means of configuring neural activity within this state space [1]. The key properties are: (i) the degree of

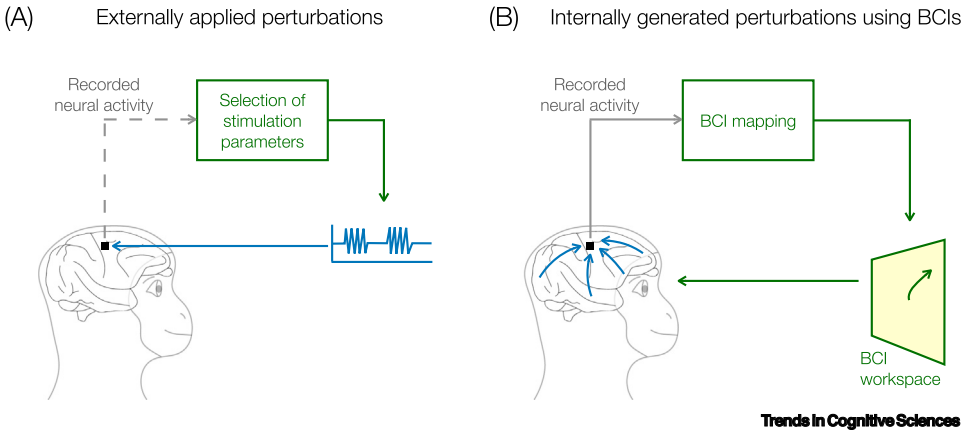

Figure 2. Approaches for neural activity perturbations. (A) Perturbation of neural activity using an external signal or force (blue arrow) directly applied to the targeted neural population (black square). Such approaches involve the selection of stimulation parameters, such as the amplitude and frequency of the external signal. In some cases, neural activity is also recorded (broken gray arrow) and can be used in a closed-loop fashion to guide the selection of stimulation parameters (green arrow). (B) Perturbation of neural activity via volitional control of the targeted population using brain–computer interfaces (BCIs). This exerts an internally generated force to modulate neural activity, which originates in upstream circuits (blue arrows). In the case of BCIs, the targeted population (black square) comprises the neurons whose activity is read out by the BCI mapping to move the BCI cursor. The animal is provided with moment-by-moment sensory feedback to move the BCI cursor to a target.

flexibility in eliciting different population activity patterns ('flexibility'); (ii) the accuracy in eliciting desired neural population activity patterns ('accuracy'); and (iii) the extent to which elicited population activity conforms to intrinsic single-neuron and network constraints ('naturalness'). We also discuss how different BCI design choices can bring about these properties and be tailored to investigate basic science questions.

Flexibility in bringing neural activity to different states

We operationalize the 'flexibility' of a perturbation tool as the degree to which it can elicit a diverse range of neural population activity patterns. This is an important capability for mapping causal relationships between various neural activity states and behavioral outcomes. The flexibility of a perturbation tool can be characterized by its ability to guide neural activity throughout state space. This includes bringing neural activity to diverse targets in state space, whether that target is a specific state or a neural trajectory (i.e., a particular sequence of activity states) (Figure 3A). For a perturbation tool to exhibit high flexibility, there should exist perturbation parameters that elicit population activity that can achieve diverse targets.

To flexibly manipulate neural activity using external perturbation methods, we need to understand how neural activity varies as a function of the perturbation parameters. For instance, when using electrical microstimulation or optogenetic stimulation, one needs to characterize how the choice of stimulation parameters, such as the location, frequency, and amplitude of the stimulation waveforms [72,73], influences neural activity. An explicit mathematical model can help capture these relationships. However, estimating such models can be challenging because it requires measurements of neural activity for numerous stimulation parameter combinations. Novel approaches are necessary to overcome data limitations to properly constrain models with such a large parameter space [73–78]. Even if one can properly constrain such a model, it is possible that no combination of parameters exists that attains certain activity targets. This can limit the flexibility to move neural activity to different regions in state space (Figure 3B).

BCIs allow us to specify the desired outcome of a perturbation by determining how neural activity must be modulated to achieve BCI task goals, such as moving a computer cursor to different target locations on the screen. We can then observe whether the animals are able to flexibly modulate activity in the target population to achieve the BCI task goals (Figure 1B). In addition to BCI task goals, other BCI design choices also allow experimenters to influence the range of population activity patterns that must be elicited for successful BCI control. For example, by changing the BCI readout (a component of the BCI mapping; see Figure IA in Box 1), one can specify the dimensions along which volitional changes in population activity result in changes in the cursor's movement (see Figure IB,C in Box 1) [49,51,55,57,79–81]. However, not all desired population activity patterns and neural trajectories are volitionally achievable. In particular, there are constraints on how neurons covary ('neural manifold') [52], on the population activity patterns that are readily expressed within the neural manifold ('neural repertoire') [79,82], and on the time courses of neural activity within the neural repertoire [54], among other constraints (Box 2).

Accuracy in guiding neural activity to a specified state

The 'accuracy' of a perturbation tool refers to how reliably it can elicit a desired population activity pattern. To assess its accuracy, we ask first, how close can it bring neural activity to a target region in neural population activity space, and second, how consistently it can do this across repeated trials. To attain accurate control over neural activity, it may be necessary to account for the neural state at the onset of the perturbation because this can influence the outcome (Figure 3C,D). There are other senses in which we can consider a perturbation technique to exhibit accuracy.

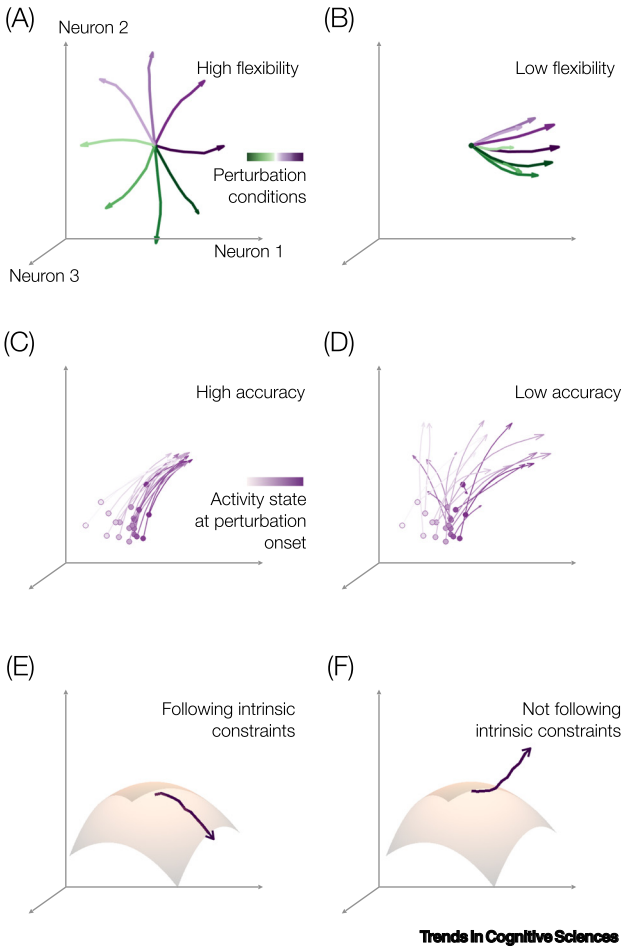

Figure 3. The properties of neural activity perturbations. (A) High flexibility. Neural activity in response to different combinations of perturbation parameters (colored trajectories) evolve in different directions. Each trajectory represents the trial-averaged activity for one particular combination of perturbation parameters. (B) Low flexibility. The direction in which neural activity trajectories evolve is similar for different combinations of perturbation parameters. As a consequence, large regions of state space cannot be achieved using such a perturbation method. (C) High accuracy. Repeated application of the same combination of perturbation parameters drives neural activity to a similar state. Each trajectory represents neural activity on a single experimental trial. The trial-to-trial variability in the neural activity at perturbation onset (shaded dots) does not influence the outcome of the perturbation. (D) Low accuracy. Population activity patterns elicited by the perturbation are dispersed because the perturbation does not compensate for the trial-to-trial variability in neural activity at perturbation onset (shaded dots, the same locations as in panel C). The neural activity after perturbation can be even more variable from trial-to-trial than the neural activity at perturbation onset. (E) The neural activity in response to a perturbation adheres to intrinsic single-neuron and network-level constraints (beige surface). (F) The neural activity in response to a perturbation deviates from the intrinsic neural constraints.

For example, it may provide the ability to define the exact location in the brain where activity changes originate, the timing of activity changes elicited, and the knowledge or specification of the mechanisms by which the elicited changes in neural activity are achieved (Table 1).

External perturbation methods are limited in their accuracy due to the impact of ongoing activity fluctuations on the perturbation outcome, which results in trial-to-trial variability. A way to improve the accuracy of neural activity control is to increase the granularity of the external perturbation method itself – for example, optical methods can selectively target specific cells and deliver

Box 2. Considerations and limitations of BCIs as a perturbation tool

Using BCIs for neural perturbations requires subjects to volitionally modulate their neural activity based on task goals and sensory feedback. This involves several key considerations. First, some physiologically possible population activity patterns may be inaccessible under volitional control. For example, animals might be unable to volitionally elicit population activity patterns that drive muscle activations without actually moving the arm. As another example, brief disruptions in excitation/inhibition (E/I) balance may be caused by the transient onset of a sensory stimulus propagating through the brain via intrinsic biophysical and network-level mechanisms [115], but might not be achievable through volitional control. Second, intrinsic circuit constraints can set limits on the timescales of perturbations that are possible using BCI-based approaches. Desired activity changes that are very fast (e.g., millisecond timescale) or that need to be triggered with very short latencies might not be achievable using volitional modulation [25]. Finally, the BCI design and the training protocols used by experimenters may also impact an animal's ability to identify appropriate control strategies to elicit certain population activity patterns using volition. For example, even when it is possible for an animal to learn to volitionally elicit new population activity patterns, how the animal learns to control a BCI can lead them to a local optimum in the solution space of neural activity [116,117]. In this case, they might require an incremental training strategy to overcome such cognitive constraints and produce the instructed patterns [56].

A limitation of BCIs as a perturbation tool is that there may be contexts in which it would be desirable to override intrinsic constraints on neural activity. This might be the case for the purpose of system identification. For example, prefrontal [118], premotor [119], and motor [98] circuits have been shown to exhibit robustness to large-scale, non-natural activity perturbations, potentially through mechanisms that involve neural population-level properties [118], inter-hemispheric communication [119], or maintaining E/I balance [98]. Another context in which we might require activity perturbations that do not adhere to intrinsic neural constraints would be for treatments of neurological disorders, where we may require activity perturbations that can override these constraints in order to create therapeutic interventions.

External perturbation methods can, in principle, allow for direct control over the location and timing of the change in neural activity. By contrast, BCI perturbations rely on upstream populations to trigger changes in the targeted population. Since BCI perturbations do not specify how neural activity should change beyond the targeted population, there can be multiple possible solutions for upstream activity changes and recurrent interactions that ultimately lead to similar activity in the targeted population. These upstream activity changes could influence the neural computation or behavior being studied, even if the computation or behavior does not rely on the targeted population.

specific amounts of stimulation to each of the selected cells [8,9]. Even still, the effect of stimulation can depend on the ongoing state of the population activity. This dependence might be overridden if the stimulation is strong (e.g., optical stimulation using high laser intensities, or electrical microstimulation using large currents). However, if the stimulation parameters are adjusted to modulate rather than override activity, then the outcome of the stimulation would depend on the ongoing activity [83]. To further improve control, a closed-loop design can be used where the stimulation parameters are adjusted in real time based on measurements of the current neural state. Implementing this approach requires models that can accurately predict how external perturbations will affect neural activity (as described in the section 'Flexibility in bringing neural activity to different states') and it is still a challenge to deploy these models in real time to account for ongoing fluctuations in neural activity (but see [75,84–87]).

By contrast, a closed-loop design is inherent to BCI paradigms. Fluctuations in neural activity (e.g., due to a slow drift in neural activity [88] or changes in neural engagement [89]) that are read out by the BCI mapping have a direct impact on the movement of the external effector and, therefore, the sensory feedback received by animals. By leveraging this feedback, animals can counteract the effect of these fluctuations using their own volitional control of neural activity to achieve successful BCI control. Thus, BCIs are a particularly powerful tool to accurately produce targeted activity that is consistent across trials.

The accuracy of neural activity control can depend on several BCI design choices, such as the dimensionality and modality of the sensory feedback provided to the animal and the effector kinematics that neural activity is mapped to. First, the dimensionality of sensory feedback determines the number of BCI readout dimensions, which is usually smaller than the number of dimensions of

Table 1. Comparison of BCIs as a perturbation tool with external perturbation methods[a,b,c]

Features		BCI perturbations	External perturbations
Flexibility	Eliciting diverse range of population activity patterns	(+) Achieved by specifying an appropriate BCI mapping and diverse BCI task goals	(±) Needs a model that maps various stimulation parameter combinations to activity targets
	Generating arbitrary activity patterns	(–) Activity patterns that cannot be volitionally elicited are not accessible	(+) Neural activity that is outside natural conditions but within physiological limits can be evoked
	Targeting identified neuronal subpopulations	(±) BCIs can be constructed such that the BCI mapping uses only identified subpopulations	(+) Optical methods can target genetically identified populations
Accuracy	Trial-to-trial reliability	(+) Closed-loop design of BCIs allows animals to counteract ongoing fluctuations in neural activity to reliably generate the desired activity	(±) Needs a model that considers ongoing activity fluctuations to select the appropriate perturbation parameters
	Control of location and timing of the origin of activity changes	(±) Does not allow full control over where and when neural activity changes originate that result in the desired activity modulation in the targeted population	(+) Allows full control of the location and timing of the origin of activity perturbation
	Timescale of activity changes	(–) Sensory delays and timescales of volitional modulation can constrain the timescale of changes in the targeted population	(+) The timescale of the perturbation is constrained only by the biophysical mechanisms that the perturbation recruits
	Knowledge of the mechanism producing activity changes	(±) The principles of BCI function are well understood, but the exact manner in which upstream circuits give rise to neural activity changes in the targeted population is not well known and cannot be readily controlled	(±) The nature of external signals delivered into the brain to cause the perturbation is known, although the precise mechanism of action on the circuit can be underspecified (e.g., electrical microstimulation can impact cell bodies or fibers of passage)
Naturalness	Activity achieved through intrinsic neuronal and network mechanisms	(+) Volitional modulation of neural activity inherently propagates through intrinsic neuronal and network mechanisms	(–) The action of external perturbations overrides intrinsic neuronal and network constraints to generate activity
	Activity appears in ethological conditions of the animal	(±) Requires designing a BCI mapping and task such that it generates neural activity that is observed during specific behaviors	(±) Requires a model that specifies the perturbation parameters that generate neural activity observed during specific behaviors

[a](+) Ways in which the method inherently exhibits a given feature.
[b](±) Ways in which the method can or needs to be developed to exhibit a given feature.
[c](–) Ways in which a method may be inherently limited.

population activity. Hence, many redundant population activity patterns lead to the same BCI readout and consequently elicit the same sensory feedback [79,90]. This can lead to dimensions along which ongoing activity fluctuations are not controlled (see Figure IB–C in Box 1). Second, the sensory modality of the feedback and its associated latency can influence how interpretable the feedback is about the ongoing neural activity. For instance, visual feedback might be easier to parse than proprioceptive feedback, but proprioceptive feedback is faster than visual feedback [91]. Whether the feedback is provided through the senses or via intracortical stimulation can also affect its interpretation. Visual feedback may be more interpretable than intracortical microstimulation of the sensory cortex, although with sufficient training, BCI users can learn to reliably interpret the microstimulation signals [31,92]. These aspects of the sensory feedback can

impact the accuracy with which neural activity is controlled moment-to-moment. Third, the BCI mapping might drive the effector's position, velocity, acceleration, or a combination of these. Each of these choices affects the degree to which neural activity changes are integrated over time to drive the effector state, and consequently the accuracy with which we can shape the time course of neural activity required for successful BCI control. Mapping neural activity to effector position creates the tightest correspondence between ongoing population activity state and effector state (see Figure ID–E in Box 1). A position-based mapping allows the experimenter to set requirements on the time course of the elicited neural activity in the readout dimensions [54].

Naturalness of elicited neural population activity patterns

'Naturalness' of elicited activity refers to the similarity of the elicited activity to neural activity that accompanies the sensory, cognitive, and motor processes that drive an animal's behavior in the absence of perturbations. To evaluate the naturalness of an elicited activity pattern, we can characterize the degree to which it aligns with activity observed during specific behaviors. When the elicited activity aligns with those observations, it can be confirmed to be natural. However, when the elicited activity does not match that observed during the animal's behavior, assessing its naturalness becomes more complicated. This is because the elicited activity pattern might occur during a behavior not tested in the laboratory setting. To conclude that elicited activity is not natural, in principle, would require knowledge of the full range of neural activity observed in ethological contexts.

Another way to quantify naturalness is in terms of whether elicited activity patterns are achievable through intrinsic single-neuron and network-level mechanisms, since population activity patterns that underlie sensory, cognitive, and motor processes are constrained by these intrinsic neural mechanisms [93,94]. As a consequence of these constraints, neural activity naturally exhibits structure in how neurons covary [50,52], in the distribution of population activity [79,82,95] and in the time course of population activity [54,96]. Using the state space view, these constraints imply that population activity tends to lie in specific regions of state space and follow particular paths as the activity changes over time (Figure 3E). Eliciting activity patterns that adhere to these intrinsic neural constraints is important when the objective of manipulating neural activity is to characterize how different neural population activity patterns causally relate to specific behaviors [1]. Perturbations that do not adhere to such constraints would be considered unnatural, but they can still be highly informative, such as for system identification (Box 2).

External perturbation methods, such as electrical or optical stimulation, can elucidate causal relationships between changes in neural activity originating in the targeted population and subsequent changes in both neural activity and behavior. These methods typically elicit unnatural population activity patterns by broadly affecting neurons, overriding biophysical processes and network-level interactions in neural populations [97,98] and leading to activity that may fall outside the regions of state space that neural activity naturally occupies (Figure 3F). As a consequence, perturbing neural activity using external forces can lead to unnatural perceptual experiences [99–101]. Recent developments in perturbation techniques, such as holographic optogenetics [9], have enhanced our ability to selectively manipulate neural activity along dimensions engaged during an animal's behavior, allowing us to generate activity patterns that resemble those observed during behavior. This has enabled the characterization of the dynamics of population activity that underlie behaviorally relevant computations (e.g., [102,103]). However, such perturbations still usually violate the network-level constraints that govern how neural activity in the targeted population is coordinated with activity in other interconnected brain regions. During an animal's behavior, a change in activity in any given population is triggered by inputs from upstream brain regions. When the targeted population is driven directly by external perturbation

methods, the relationship between activity in that region and its upstream brain regions would be altered. If we want to perturb activity in the targeted population in a manner that resembles how its activity would be driven by other brain regions, we would need to manipulate activity in multiple brain regions simultaneously.

By contrast, BCIs perturb neural activity through volitional forces internal to the brain, and thus, recruit the inputs from other neurons and brain areas needed to modulate the targeted neural population. Because of this, the population activity patterns elicited using BCIs obey network-level constraints due to network connectivity within the targeted population, as well as constraints due to connectivity across brain regions to which the targeted population is connected (Box 2). Even when novel population activity patterns are induced during BCI learning, these novel patterns are still constrained by physiological limitations and are learned through intrinsic synaptic and network mechanisms [47,104]. Additionally, how BCIs are designed can influence the degree to which elicited activity resembles activity during other behavioral contexts. For example, biomimetic decoder designs aim to match the mapping between neural activity and overt movements as closely as possible [105].

Concluding remarks

The main thesis of our review is that BCIs function as a tool for the causal perturbation of neural activity. Seen in this light, we can consider BCIs among the suite of tools available to manipulate neural activity, and we can expand the range of questions that can be addressed about neural circuit function. As with any perturbation tool, BCIs have a unique mixture of advantages and limitations in how they manipulate and perturb neural activity. These need to be considered when selecting BCIs as a causal perturbation method, and in interpreting the outcomes of BCI-based perturbations (Box 2 and Table 1). It is also possible to combine volitional activity modulation using BCIs with external perturbation methods to further expand the range of questions that can be addressed [59,60,106–108].

As a result of its reliance on volition to modulate neural activity, BCIs may be well suited to investigate complex cognitive functions that rely on volitional processes that are distributed across many brain areas (see Outstanding questions). These include learning, attention, decision making, motivation, emotion, motor control, and more [46–48,63,109,110]. Where in the brain these processes originate is uncertain and, likely, there is no single point of origin. BCIs can use volition to recruit the broader networks in which the targeted population is embedded, and doing so may be necessary to influence cognitive functions that rely on distributed circuits across the brain.

BCIs can also be used to repeatedly manipulate brain-wide neural activity that supports complex cognitive processes. Doing so could have significant therapeutic implications since repeated modulation of activity can induce long-lasting modifications of neural pathways [111–113]. As a result, BCI-based perturbations may be a promising tool for the treatment of cognitive and neurological disorders by potentially enabling long-term reconfiguration of neural activity across the brain.

Outstanding questions

What aspects of neural activity are and are not volitionally controllable?

Are there differences in the level of volitional control of neural activity in different brain areas or neuronal subpopulations?

How widespread are the changes in neural activity induced by volitional control beyond the targeted population?

How can we design sensory feedback to be most useful to animals for enacting the desired volitional changes in neural activity?

For which brain disorders are BCI-based perturbations most amenable as a therapeutic tool?

Acknowledgments

We would like to thank Caroline Runyan and Steve Chase for valuable comments on the manuscript. We are grateful to the members of the Yu, Smith, Chase, and Batista laboratories for valuable discussions. This work was supported by The Swartz Foundation (A.M.), National Science Foundation (NSF) Neural and Cognitive Systems (NCS) Division of Research on Learning (DRL) 2124066 and 2123911 (B.M.Y., M.A.S., and A.P.B.), National Institutes of Health (NIH) Collaborative Research in Computational Neuroscience (CRCNS) R01 NS105318 (B.M.Y. and A.P.B.), NIH R01 NS129584 (A.P.B. and B.M.Y.), NIH CRCNS R01 MH118929 (B.M.Y. and M.A.S.), and Simons Foundation NC-GB-CULM-00003241-05 (B.M.Y.).

Declaration of interests

A.P.B. and B.M.Y. are inventors of pending US Patent Application No. 17/343,050, and M.A.S. and B.M.Y. are inventors of pending Patent Application No. PCT/US2024/010801. Both are related to using brain–computer interfaces as described in this review. J.S.M., M.A.S., and B.M.Y. are inventors of pending US Patent Application No. 63/697 896 and 63/795 999, which describe a closed-loop framework for external perturbations.

Declaration of generative AI and AI-assisted technologies in the writing process

After the manuscript text was written entirely by the authors, the authors used Llama 3.3 in order to occasionally verify the grammatical construction of some individual sentences, and to simplify their structure to improve clarity. After using this tool, the authors reviewed and edited the content and take full responsibility for the content of the publication.

References

1. Jazayeri, M. and Afraz, A. (2017) Navigating the neural space in search of the neural code. *Neuron* 93, 1003–1014
2. Cohen, M.R. and Newsome, W.T. (2004) What electrical microstimulation has revealed about the neural basis of cognition. *Curr. Opin. Neurobiol.* 14, 169–177
3. Histed, M.H. *et al.* (2013) Insights into cortical mechanisms of behavior from microstimulation experiments. *Prog. Neurobiol.* 103, 115–130
4. Liu, A. *et al.* (2018) Immediate neurophysiological effects of transcranial electrical stimulation. *Nat. Commun.* 9, 5092
5. Yavari, F. *et al.* (2018) Basic and functional effects of transcranial electrical stimulation (tES)—an introduction. *Neurosci. Biobehav. Rev.* 85, 81–92
6. Hallett, M. (2007) Transcranial magnetic stimulation: a primer. *Neuron* 55, 187–199
7. Lefaucheur, J.-P. (2019) Chapter 37 - Transcranial magnetic stimulation. In *Handbook of Clinical Neurology* (160) (Levin, K.H. and Chauvel, P., eds), pp. 559–580, Elsevier
8. Fenno, L. *et al.* (2011) The development and application of optogenetics. *Annu. Rev. Neurosci.* 34, 389–412
9. Adesnik, H. and Abdeladim, L. (2021) Probing neural codes with two-photon holographic optogenetics. *Nat. Neurosci.* 24, 1356–1366
10. Bolz, J. *et al.* (1989) Pharmacological analysis of cortical circuitry. *Trends Neurosci.* 12, 292–296
11. Martin, J.H. and Ghez, C. (1999) Pharmacological inactivation in the analysis of the central control of movement. *J. Neurosci. Methods* 86, 145–159
12. Roth, B.L. (2016) DREADDs for neuroscientists. *Neuron* 89, 683–694
13. Smith, K.S. *et al.* (2016) DREADDs: use and application in behavioral neuroscience. *Behav. Neurosci.* 130, 137–155
14. Rabut, C. *et al.* (2020) Ultrasound technologies for imaging and modulating neural activity. *Neuron* 108, 93–110
15. Murphy, K.R. *et al.* (2024) Optimized ultrasound neuromodulation for non-invasive control of behavior and physiology. *Neuron* 112, 3252–3266.e5
16. Szczepanski, S.M. and Knight, R.T. (2014) Insights into human behavior from lesions to the prefrontal cortex. *Neuron* 83, 1002–1018
17. Vaidya, A.R. *et al.* (2019) Lesion studies in contemporary neuroscience. *Trends Cogn. Sci.* 23, 653–671
18. Brooks, V.B. (1983) Study of brain function by local, reversible cooling. In *Reviews of Physiology, Biochemistry and Pharmacology, Volume 95* (Vol. 95), pp. 1–109, Springer
19. Banerjee, A. *et al.* (2021) Using focal cooling to link neural dynamics and behavior. *Neuron* 109, 2508–2518
20. Shenoy, K.V. and Yu, B.M. (2021) Brain-machine interfaces. In *Principles of Neural Science* (6th edn) (Kandel, E.R. *et al.*, eds), McGraw Hill
21. Serruya, M.D. *et al.* (2002) Instant neural control of a movement signal. *Nature* 416, 141–142
22. Taylor, D.M. *et al.* (2002) Direct cortical control of 3D neuroprosthetic devices. *Science* 296, 1829–1832
23. Carmena, J.M. *et al.* (2003) Learning to control a brain–machine interface for reaching and grasping by primates. *PLoS Biol.* 1, e42
24. Musallam, S. *et al.* (2004) Cognitive control signals for neural prosthetics. *Science* 305, 258–262
25. Santhanam, G. *et al.* (2006) A high-performance brain–computer interface. *Nature* 442, 195–198
26. Hochberg, L.R. *et al.* (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442, 164–171
27. Collinger, J.L. *et al.* (2013) High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* 381, 557–564
28. Bouton, C.E. *et al.* (2016) Restoring cortical control of functional movement in a human with quadriplegia. *Nature* 533, 247–250
29. Ajiboye, A.B. *et al.* (2017) Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration. *Lancet Lond. Engl.* 389, 1821–1830
30. Pandarinath, C. *et al.* (2017) High performance communication by people with paralysis using an intracortical brain-computer interface. *eLife* 6, e18554
31. Flesher, S.N. *et al.* (2021) A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science* 372, 831–836
32. Metzger, S.L. *et al.* (2023) A high-performance neuroprosthesis for speech decoding and avatar control. *Nature* 620, 1037–1046
33. Willett, F.R. *et al.* (2023) A high-performance speech neuroprosthesis. *Nature* 620, 1031–1036
34. Card, N.S. *et al.* (2024) An accurate and rapidly calibrating speech neuroprosthesis. *N. Engl. J. Med.* 391, 609–618
35. Kamiya, J. (1971) Biofeedback training in voluntary control of EEG alpha rhythms. *Calif. Med.* 115, 44
36. Sterman, M.B. and Egner, T. (2006) Foundation and practice of neurofeedback for the treatment of epilepsy. *Appl. Psychophysiol. Biofeedback* 31, 21–35
37. Sitaram, R. *et al.* (2017) Closed-loop brain training: the science of neurofeedback. *Nat. Rev. Neurosci.* 18, 86–100
38. Sitaram, R. *et al.* (2024) Mechanisms of brain self-regulation: psychological factors, mechanistic models and neural substrates. *Philos. Trans. R. Soc. B Biol. Sci.* 379, 20230093
39. Watanabe, T. *et al.* (2017) Advances in fMRI real-time neurofeedback. *Trends Cogn. Sci.* 21, 997–1010
40. Haugg, A. *et al.* (2021) Predictors of real-time fMRI neurofeedback performance and improvement – a machine learning mega-analysis. *NeuroImage* 237, 118207
41. Loriette, C. *et al.* (2021) Neurofeedback for cognitive enhancement and intervention and brain plasticity. *Rev. Neurol. (Paris)* 177, 1133–1144
42. Nicolelis, M.A.L. (2003) Brain-machine interfaces to restore motor function and probe neural circuits. *Nat. Rev. Neurosci.* 4, 417–422
43. Orsborn, A. and Carmena, J.M. (2013) Creating new functional circuits for action via brain-machine interfaces. *Front. Comput. Neurosci.* 7, 157
44. Wander, J.D. and Rao, R.P.N. (2014) Brain-computer interfaces: a powerful tool for scientific inquiry. *Curr. Opin. Neurobiol.* 25, 70–75
45. Moxon, K.A. and Foffani, G. (2015) Brain-machine interfaces beyond neuroprosthetics. *Neuron* 86, 55–67
46. Golub, M.D. *et al.* (2016) Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Curr. Opin. Neurobiol.* 37, 53–58

47. Orsborn, A.L. and Pesaran, B. (2017) Parsing learning in networks using brain-machine interfaces. *Curr. Opin. Neurobiol.* 46, 76–83
48. Andersen, R.A. *et al.* (2022) Exploring cognition with brain-machine interfaces. *Annu. Rev. Psychol.* 73, 131–158
49. 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
50. Hwang, E.J. *et al.* (2013) Volitional control of neural activity relies on the natural motor repertoire. *Curr. Biol. CB* 23, 353–361
51. Law, A.J. *et al.* (2014) Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. *J. Neurophysiol.* 112, 1528–1548
52. Sadtler, P.T. *et al.* (2014) Neural constraints on learning. *Nature* 512, 423–426
53. Neely, R.M. *et al.* (2018) Volitional modulation of primary visual cortex activity requires the basal ganglia. *Neuron* 97, 1356–1368.e4
54. Oby, E.R. *et al.* (2025) Dynamical constraints on neural population activity. *Nat. Neurosci.* 28, 383–393
55. Ganguly, K. *et al.* (2011) Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* 14, 662–667
56. Oby, E.R. *et al.* (2019) New neural activity patterns emerge with long-term learning. *Proc. Natl. Acad. Sci. U. S. A.* 116, 15210–15215
57. Zhou, X. *et al.* (2019) Distinct types of neural reorganization during long-term learning. *J. Neurophysiol.* 121, 1329–1341
58. Gulati, T. *et al.* (2017) Neural reactivations during sleep determine network credit assignment. *Nat. Neurosci.* 20, 1277–1284
59. Prsa, M. *et al.* (2017) Rapid integration of artificial sensory feedback during operant conditioning of motor cortex neurons. *Neuron* 93, 929–939.e6
60. Athalye, V.R. *et al.* (2018) Evidence for a neural law of effect. *Science* 359, 1024–1029
61. Koralek, A.C. *et al.* (2012) Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483, 331–335
62. Wander, J.D. *et al.* (2013) Distributed cortical adaptation during learning of a brain–computer interface task. *Proc. Natl. Acad. Sci.* 110, 10818–10823
63. Bassett, D.S. and Khambhati, A.N. (2017) A network engineering perspective on probing and perturbing cognition with neurofeedback. *Ann. N. Y. Acad. Sci.* 1396, 126–143
64. Schafer, R.J. and Moore, T. (2011) Selective attention from voluntary control of neurons in prefrontal cortex. *Science* 332, 1568–1571
65. Vyas, S. *et al.* (2018) Neural population dynamics underlying motor learning transfer. *Neuron* 97, 1177–1186.e3
66. Losey, D.M. *et al.* (2024) Learning leaves a memory trace in motor cortex. *Curr. Biol.* 34, 1519–1531.e4
67. Motiwala, A. *et al.* (2024) Probing movement preparation dynamics using BCI-based causal perturbations Computation and Systems Neuroscience (COSYNE) Abstracts, Lisbon, Portugal T-29
68. Cunningham, J.P. and Yu, B.M. (2014) Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17, 1500–1509
69. Gallego, J.A. *et al.* (2017) Neural manifolds for the control of movement. *Neuron* 94, 978–984
70. Kao, T.-C. and Hennequin, G. (2019) Neuroscience out of control: control-theoretic perspectives on neural circuit dynamics. *Curr. Opin. Neurobiol.* 58, 122–129
71. Gurnani, H. and Cayco Gajic, N.A. (2023) Signatures of task learning in neural representations. *Curr. Opin. Neurobiol.* 83, 102759
72. Brocker, D.T. and Grill, W.M. (2013) Principles of electrical stimulation of neural tissue. *Handb. Clin. Neurol.* 116, 3–18
73. Cole, E.R. *et al.* (2024) Irregular optogenetic stimulation waveforms can induce naturalistic patterns of hippocampal spectral activity. *J. Neural Eng.* 21, 036039
74. Kuncel, A.M. and Grill, W.M. (2004) Selection of stimulus parameters for deep brain stimulation. *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.* 115, 2431–2441
75. Tafazoli, S. *et al.* (2020) Learning to control the brain through adaptive closed-loop patterned stimulation. *J. Neural Eng.* 17, 056007
76. Yang, Y. *et al.* (2021) Modelling and prediction of the dynamic responses of large-scale brain networks during direct electrical stimulation. *Nat. Biomed. Eng.* 5, 324–345
77. Nejatbakhsh, A. *et al.* (2023) Predicting the effect of microstimulation on macaque prefrontal activity based on spontaneous circuit dynamics. *Phys. Rev. Res.* 5, 043211
78. Minai, Y. *et al.* (2024) MiSO: optimizing brain stimulation to create neural activity states. *Adv. Neural Inf. Proces. Syst.* 37, 24126–24149
79. Hennig, J.A. *et al.* (2018) Constraints on neural redundancy. *eLife* 7, e36774
80. Clancy, K.B. *et al.* (2014) Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* 17, 807–809
81. Mitani, A. *et al.* (2018) Brain-computer interface with inhibitory neurons reveals subtype-specific strategies. *Curr. Biol.* 28, 77–83.e4
82. Golub, M.D. *et al.* (2018) Learning by neural reassociation. *Nat. Neurosci.* 21, 607–616
83. Dadarlat, M.C. *et al.* (2024) Activity-dependent recruitment of inhibition and excitation in the awake mammalian cortex during electrical stimulation. *Neuron* 112, 821–834.e4
84. Laferriere, S. *et al.* (2020) Hierarchical Bayesian optimization of spatiotemporal neurostimulations for targeted motor outputs. *IEEE Trans. Neural Syst. Rehabil. Eng.* 28, 1452–1460
85. Saalmann, Y.B. *et al.* (2023) Microscale multicircuit brain stimulation: achieving real-time brain state control for novel applications. *Curr. Res. Neurobiol.* 4, 100071
86. Shang, C.-F. *et al.* (2024) Real-time analysis of large-scale neuronal imaging enables closed-loop investigation of neural dynamics. *Nat. Neurosci.* 27, 1014–1018
87. Minai, Y. *et al.* (2025) OMiSO: adaptive optimization of state-dependent brain stimulation to shape neural population states. *arXiv*, Published online July 10, 2025. https://doi.org/10.48550/arXiv.2507.07858
88. Cowley, B.R. *et al.* (2020) Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron* 108, 551–567.e8
89. Hennig, J.A. *et al.* (2021) Learning is shaped by abrupt changes in neural engagement. *Nat. Neurosci.* 24, 727–736
90. Kaufman, M.T. *et al.* (2014) Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440–448
91. Crevecoeur, F. *et al.* (2016) Dynamic multisensory integration: somatosensory speed trumps visual accuracy during feedback control. *J. Neurosci.* 36, 8598–8611
92. Dadarlat, M.C. *et al.* (2015) A learning-based approach to artificial sensory feedback leads to optimal integration. *Nat. Neurosci.* 18, 138–144
93. Saxena, S. and Cunningham, J.P. (2019) Towards the neural population doctrine. *Curr. Opin. Neurobiol.* 55, 103–111
94. Ebitz, R.B. and Hayden, B.Y. (2021) The population doctrine in cognitive neuroscience. *Neuron* 109, 3055–3068
95. Athalye, V.R. *et al.* (2017) Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* 93, 955–970.e5
96. Athalye, V.R. *et al.* (2023) Invariant neural dynamics drive commands to control different movements. *Curr. Biol.* 33, 2962–2976.e15
97. Griffin, D.M. *et al.* (2011) Hijacking cortical motor output with repetitive microstimulation. *J. Neurosci. Off. J. Soc. Neurosci.* 31, 13088–13096
98. O'Shea, D.J. *et al.* (2022) Direct neural perturbations reveal a dynamical mechanism for robust computation. *bioRxiv*, Published online December 16, 2022. https://doi.org/10.1101/2022.12.16.520768
99. Fernandez, E. *et al.* (2021) Visual percepts evoked with an intracortical 96-channel microelectrode array inserted in human occipital cortex. *J. Clin. Invest.* 23, 131
100. Godlove, J. *et al.* (2014) Comparing temporal aspects of visual, tactile, and microstimulation feedback for motor control. *J. Neural Eng.* 11, 046025

101. Hughes, C.L. *et al.* (2021) Perception of microstimulation frequency in human somatosensory cortex. *eLife* 10, e65128
102. Daie, K. *et al.* (2021) Targeted photostimulation uncovers circuit motifs supporting short-term memory. *Nat. Neurosci.* 24, 259–265
103. Vinograd, A. *et al.* (2024) Causal evidence of a line attractor encoding an affective state. *Nature* 634, 910–918. https://doi.org/10.1038/s41586-024-07915-x
104. Hennig, J.A. *et al.* (2021) How learning unfolds in the brain: toward an optimization view. *Neuron* 109, 3720–3735
105. Shenoy, K.V. and Carmena, J.M. (2014) Combining decoder design and neural adaptation in brain-machine interfaces. *Neuron* 84, 665–680
106. O'Doherty, J.E. *et al.* (2011) Active tactile exploration using a brain–machine–brain interface. *Nature* 479, 228–231
107. Klaes, C. *et al.* (2014) A cognitive neuroprosthetic that uses cortical stimulation for somatosensory feedback. *J. Neural Eng.* 11, 056024
108. Eaton, R.W. *et al.* (2017) Operant conditioning of neural activity in freely behaving monkeys with intracranial reinforcement. *J. Neurophysiol.* 117, 1112–1125
109. Shanechi, M.M. (2019) Brain–machine interfaces from motor to mood. *Nat. Neurosci.* 22, 1554–1564
110. Gallego, J.A. *et al.* (2022) Going beyond primary motor cortex to improve brain-computer interfaces. *Trends Neurosci.* 45, 176–183
111. Jackson, A. and Zimmermann, J.B. (2012) Neural interfaces for the brain and spinal cord—restoring motor function. *Nat. Rev. Neurol.* 8, 690–699
112. Ganguly, K. and Poo, M. (2013) Activity-dependent neural plasticity from bench to bedside. *Neuron* 80, 729–741
113. Courtine, G. and Sofroniew, M.V. (2019) Spinal cord repair: advances in biology and technology. *Nat. Med.* 25, 898–908
114. Suminski, A.J. *et al.* (2010) Incorporating feedback from multiple sensory modalities enhances brain–machine interface control. *J. Neurosci.* 30, 16777–16787
115. Murphy, B.K. and Miller, K.D. (2009) Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron* 61, 635–648
116. Rajeswaran, P. *et al.* (2025) Assistive sensory-motor perturbations influence learned neural representations. *bioRxiv*, Published online April 2, 2025. https://doi.org/10.1101/2024.03.20.585972
117. Chang, J.C. *et al.* (2024) De novo motor learning creates structure in neural activity that shapes adaptation. *Nat. Commun.* 15, 4084
118. Soldado-Magraner, J. *et al.* (2025) Robustness of working memory to prefrontal cortex microstimulation. *J. Neurosci.* e2197242025 https://doi.org/10.1523/JNEUROSCI.2197-24.2025
119. Li, N. *et al.* (2016) Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459–464