SVOA Neurology

e-ISSN 2753-9180

ScienceVolks Open Access

Review Article

The Milliseconds We Do Not Measure: Feedback Latency, Spike-Timing-Dependent Plasticity, and the Mechanistic Claims of Closed-Loop Brain–Computer Interfaces After Stroke

Anirudda Deshpande[1*], Vikram Khardenavis[2], Supriya Khardenavis[3], Gayatri Godse[4]

[1]Consultant Stroke Neurologist, Department of Elderly Care & Stroke, Altnagelvin Hospital, Londonderry, United Kingdom.

[2]Fellow in Epilepsy, Department of Neurology, Kokilaben Dhirubhai Ambani Hospital, Mumbai, Maharashtra, India.

[3]Specialty Doctor, Department of Ophthalmology, Altnagelvin Hospital, Western HSC Trust, Londonderry, United Kingdom.

[4]Senior Resident, Department of Ophthalmology, MGM Hospital Kamothe, Navi Mumbai, India.

***Corresponding Author:** Dr. Anirudda Deshpande, Consultant Stroke Neurologist, Department of Elderly Care & Stroke, Altnagelvin Hospital, Londonderry, United Kingdom.

DOI: https://doi.org/10.58624/SVOANE.2026.07.037

Received: August 19, 2026

Published: September 20, 2026

Citation: Deshpande A, Khardenavis V, Khardenavis S, Godse G. The Milliseconds We Do Not Measure: Feedback Latency, Spike-Timing-Dependent Plasticity, and the Mechanistic Claims of Closed-Loop Brain–Computer Interfaces After Stroke. *SVOA Neurology* 2026, 7:5, 297-313. doi.org/10.58624/SVOANE.2026.07.037

Abstract

Almost every review of brain–computer interfaces in stroke rehabilitation contains a version of the same sentence: by delivering sensory feedback time-locked to detected motor intent, the BCI creates the conditions for Hebbian, spike-timing-dependent plasticity. The sentence is repeated so consistently that it has acquired the status of furniture. It is also, taken literally, difficult to reconcile with the systems being described. Spike-timing-dependent plasticity is defined over a window of tens of milliseconds and reverses sign within that window. The characteristic timescale of clinically deployed closed-loop BCI is one to three orders of magnitude longer, dominated not by processing delay but by the decision window the classifier requires. If cellular STDP were the operative mechanism, most of the field's interventions would be operating well outside the rule they invoke. This is not an argument that these interventions do not work. It is an argument that a mechanism specified at the synapse has been used as a stand-in for one operating at the timescale of behaviour, and that the substitution has consequences. It makes latency look like an implementation detail rather than a design parameter; it licenses the pooling of systems whose plasticity mechanisms differ by orders of magnitude; and it discourages pursuit of the millisecond-scale interventions where timing genuinely is the active ingredient. We decompose the latency budget of closed-loop BCI, set it against the empirical time constants of timing-dependent plasticity in animal preparations and in human paired associative stimulation, and examine three resolutions — that the mechanism was never cellular STDP; that neuromodulatory eligibility traces bridge the gap; and that some component is genuinely timing-critical and is being discarded. We propose a timescale taxonomy of closed-loop systems, derive three falsifiable predictions, and specify a minimum latency reporting set for the field. All three resolutions appear partly correct, which is less satisfying than any one alone, and we think more honest.

Keywords: *Brain–Computer Interface; Spike-Timing-Dependent Plasticity; Latency; Closed-Loop Systems; Stroke Rehabilitation; Functional Electrical Stimulation; Eligibility Trace; Neuromodulation; Paired Associative Stimulation; Hebbian Plasticity*

1. Introduction

There is a sentence that appears, with minor variation, in the introduction of most papers on brain–computer interfaces for motor rehabilitation. It runs approximately as follows: because the BCI detects the cortical correlate of intended movement and delivers a contingent sensory consequence in real time, it establishes the temporal pairing required for Hebbian, spike-timing-dependent plasticity, thereby strengthening the residual corticomotor pathway. One of us has written that sentence. The other has reviewed a great many papers containing it, and has published systems it was used to justify [40, 41]. It is not a controversial claim within the field; it is closer to a formality, of the kind that reviewers skim and editors never query. It has been part of the furniture since the field's founding statements of purpose [39].

We want to take it seriously, which turns out to be an uncomfortable thing to do.

Spike-timing-dependent plasticity is not a loose synonym for associative learning. It is a specific, quantitatively characterised rule, first described in cultured and slice preparations, in which the sign and magnitude of a synaptic weight change depend on the interval between presynaptic and postsynaptic firing [1, 2]. The potentiating limb of that rule operates over roughly twenty to forty milliseconds and decays approximately exponentially; reversing the order of firing within a comparable interval produces depression rather than potentiation [3, 4]. The rule is asymmetric, it is sharp, and its defining feature is precisely the property that makes it attractive as a mechanistic justification: it cares about timing at the millisecond scale.

Now consider what a closed-loop BCI actually does between the patient's intention and the arrival of the corresponding afferent volley at cortex. It amplifies and digitises an EEG signal. It accumulates that signal into an analysis window, because a single sample carries no usable information about sensorimotor rhythm modulation; the window is typically between half a second and two seconds long. It extracts features from the window, classifies them, frequently applies temporal smoothing or a dwell-time criterion to suppress false positives, issues a command, waits for a stimulator or an actuator to reach a therapeutically meaningful output, and then waits again while the resulting afferent volley travels from the periphery to the sensorimotor cortex. The total is not tens of milliseconds. In the systems that have generated the field's clinical evidence base, it is hundreds of milliseconds at best and often several seconds.

This is not a small discrepancy that careful engineering will close. It is a mismatch of one to three orders of magnitude, and it sits between the mechanism the field cites and the systems the field builds.

We do not claim the observation is unprecedented. A minority literature has questioned the Hebbian framing directly, arguing that restorative BCIs are better understood through reinforcement and operant conditioning than through coincidence detection [45, 46], and that the parameters governing that reinforcement — threshold, difficulty, reward schedule — are themselves plasticity variables [47]. Every group that builds these systems knows what its analysis window is, and everyone who reads the plasticity literature knows how wide the STDP window is. What has not, so far as we can establish, been attempted is a systematic comparison of the two quantities, an account of which terms in the latency budget are irreducible, and a derivation of what follows for classification, trial design and reporting. That is the work this paper attempts.

There are three ways to react to the mismatch. The first is to conclude that BCI rehabilitation does not work, which is not our position and is not supported by the trial evidence; meta-analyses report consistent, if heterogeneous, benefit over conventional therapy [5, 6]. The second is to conclude that the mismatch does not matter, which is comfortable but requires an argument the literature has not supplied. The third — ours — is that the field has been using a cellular mechanism as a stand-in for a systems-level one, and that examining the substitution yields a more useful account of what closed-loop rehabilitation is doing and how it should be designed.

The consequences are not merely semantic. If STDP is the stated mechanism, then latency is a first-order design parameter and should appear in every methods section and every device specification. It does not. If instead the operative mechanism is operant reinforcement of a neural state, or associative plasticity gated by neuromodulatory signals, then latency matters differently and less acutely — but so do several things the field currently ignores, including reward structure, agency, and the pharmacological state of the patient. Either way, the current position, in which a millisecond-scale mechanism is invoked to justify a second-scale intervention and nothing follows from the invocation, is the one position that cannot be right.

Section 2 sets out what timing-dependent plasticity actually specifies, including the human evidence, which is more precise than the rehabilitation literature usually acknowledges. Section 3 decomposes where the milliseconds go and identifies which terms are irreducible. Sections 4 and 5 state the mismatch and examine three candidate resolutions. Sections 6 to 8 turn the analysis into something usable: a taxonomy by operative timescale, three falsifiable predictions, and a minimum reporting set. Section 9 takes the strongest counterarguments seriously.

We should say at the outset that we disagree with each other about how much of this matters. That disagreement survives into the text, and we have not smoothed it over.

2. What Spike-Timing-Dependent Plasticity Actually Specifies

2.1 The rule and its time constants

Hebb's original formulation was famously imprecise about time [7]. The cell that fires and thereby takes part in firing another was not assigned a window, and for four decades the absence did not much matter, because there was no method capable of resolving it. That changed with the demonstration, in neocortical and hippocampal preparations, that the order and interval of pre- and postsynaptic spiking determine the direction of the resulting weight change [1, 2].

The quantitative picture that emerged is consistent in its essentials across preparations. Presynaptic activity preceding postsynaptic activity by a few to a few tens of milliseconds produces timing-dependent long-term potentiation. The reverse order, over a comparable interval, produces timing-dependent depression. The transition between the two is abrupt, occurring within a window of a few milliseconds around coincidence, and the magnitude of the effect in both directions decays with a time constant on the order of twenty milliseconds [3, 4, 8]. Beyond roughly a hundred milliseconds, in most preparations, the pairing produces no reliable weight change at all.

Two qualifications are important and are frequently omitted when the rule is cited in rehabilitation contexts. First, the canonical rule is not universal: the shape, width, and even the polarity of the window vary by synapse type, dendritic location, developmental stage, and the presence of neuromodulators [9, 10]. Second, timing is necessary but not sufficient. Firing rate, cooperativity between synapses, and postsynaptic depolarisation all interact with timing to determine the outcome, and at higher firing rates the timing dependence weakens considerably [11, 12]. STDP is better understood as one face of a multidimensional plasticity rule than as a complete account of synaptic learning.

Neither qualification rescues the rehabilitation claim. Widening the window from twenty to a hundred milliseconds, or acknowledging that rate matters too, does not accommodate a system operating at one to three seconds. If anything, the qualifications sharpen the problem, because they establish that the field's chosen mechanism has been characterised carefully enough for the mismatch to be measured rather than argued about.

2.2 The human evidence, which is better than the field admits

It is sometimes suggested, in conversation if not in print, that STDP is a slice-preparation phenomenon of uncertain relevance to the intact human brain, and that invoking it in a clinical paper is a permissible simplification. This is not correct, and the reason it is not correct is the most useful part of the argument we are making.

Paired associative stimulation provides a direct human analogue. When a peripheral nerve stimulus is paired repeatedly with a transcranial magnetic stimulus delivered to the contralateral motor cortex, the resulting change in corticospinal excitability depends critically on the interstimulus interval [13]. At an interval of approximately twenty-five milliseconds — chosen so that the afferent volley arrives at cortex just before the magnetic pulse — the protocol produces a durable, LTP-like increase in motor evoked potential amplitude. At an interval of approximately ten milliseconds, the same protocol produces an LTD-like decrease [14]. The direction of the induced plasticity reverses over a difference of fifteen milliseconds in an awake human being.

The picture is reinforced by state-dependent stimulation. When transcranial magnetic pulses are triggered on the instantaneous phase of the ongoing sensorimotor mu rhythm, pulses delivered at the trough produce reliably different plasticity outcomes from pulses delivered at the peak, a distinction spanning a few tens of milliseconds [15]. In primates, spike-triggered stimulation delivered through an autonomous implanted device strengthens or weakens corticospinal and corticocortical connections according to the delay imposed between the triggering spike and the stimulus, with the effect following the expected timing rule over free behaviour and sleep [16, 17].

The human motor system, in other words, demonstrably obeys a timing rule of the expected width. The rehabilitation field cannot simultaneously claim STDP as its mechanistic warrant and treat the timing specification as an idealisation.

2.3 The one part of the rehabilitation literature that gets this right

There is a body of work that has taken the timing constraint seriously, and it deserves more attention than it receives. Associative BCI protocols detect the negative peak of the movement-related cortical potential — a marker of imminent movement execution or vivid imagination — and deliver peripheral electrical stimulation timed so that the resulting afferent volley reaches the sensorimotor cortex within the associative window [18, 19, 37]. The same principle has been extended to lower limb work using an ankle-foot orthosis [38]. These systems are not attempting to be fast in a general sense; they are attempting to be correct about a specific interval, which is a different engineering problem. The reported outcome, both in healthy participants and in chronic stroke, is an increase in corticospinal excitability that scales with the precision of the pairing, and clinical benefit that has been demonstrated in controlled work [20].

This literature is usually filed alongside conventional BCI-FES in reviews and meta-analyses. Filing it there is a mistake, and it is the same mistake this paper is about: two interventions that share an acronym, and share the word closed-loop, are being pooled despite operating on different mechanisms at different timescales. We return to this in Section 6.

3. Where The Milliseconds Go

Anyone who has sat with a patient through a BCI session knows the delay is perceptible. The patient concentrates, something happens on the screen, and a moment later the stimulator fires. Clinicians read that pause as a technical inconvenience, or as evidence that the system is working hard. Almost nobody reads it as a mechanistic variable, which is what it is.

Discussion of BCI latency in the clinical literature, where it occurs at all, tends to treat it as a single quantity attributable to computational speed — as though the relevant question were processor performance. It is not. The dominant term in the latency budget of a conventional motor imagery BCI is not computation; it is the analysis window, which is an information-theoretic requirement rather than an engineering shortcoming. Decomposing the budget is therefore essential, because the terms differ in whether they can be reduced at all.

Table 1. Decomposition of end-to-end latency in a conventional closed-loop BCI for upper limb rehabilitation. Ranges are indicative of currently deployed systems rather than theoretical limits, and are derived from parameters stated in the trials cited in the final column together with standard device specifications; where a range rests on device specification rather than a published measurement this is stated. The penultimate column indicates whether the term is reducible by engineering effort, by paradigm redesign, or not at all.

Stage	Typical contribution	Dominant determinant	Reducible?	Basis
Amplification, filtering, digitisation	1–20 ms	Hardware and anti-aliasing filter design	Yes, largely solved	Amplifier specification
Analysis window accumulation	500–2000 ms	Spectral resolution required to detect ERD/ERS reliably	Only by changing the decoded feature	[21]; windows stated in [22–27]

Table 1. *Continued...*

Stage	Typical contribution	Dominant determinant	Reducible?	Basis
Feature extraction and classification	5–50 ms	Algorithm complexity; usually trivial on modern hardware	Yes, already minor	[41, 43]
Decision smoothing / dwell-time criterion	0–1000 ms	False-positive tolerance chosen by the designer	Yes, at a cost in specificity	Design choice; where stated, [22–27]
Command transmission and actuator onset	10–100 ms	Communication protocol; motor or stimulator rise time	Partly	Device specification; [42]
Stimulus-to-contraction delay (FES)	30–80 ms	Neuromuscular activation dynamics (electromechanical delay)	No	■ authors to insert a primary source for electromechanical delay under FES
Peripheral afferent transit to cortex	20–40 ms	Conduction velocity and pathway length	No	Somatosensory afferent arrival time, as exploited in PAS [13]
Typical end-to-end total	**600–3000 ms**			

Several features of this decomposition deserve comment.

The analysis window dominates, and it is not a defect. Event-related desynchronisation in the alpha and beta bands is a spectral phenomenon; estimating band power with sufficient signal-to-noise to support a reliable single-trial decision requires accumulating data over a period that is inversely related to the frequency resolution needed [21]. A system that shortens its window buys latency at the direct cost of classification accuracy, which in a stroke population with degraded and variable signals is the currency it can least afford to spend. Connectivity and network features face the same constraint in a different form, since network estimation is if anything more data-hungry [43]; hybrid schemes combining cortical with residual muscle activity shorten the effective window, but only in patients who have residual muscle activity to combine [41]. This is the central tension of the field, and it is rarely stated as one: latency and accuracy trade against each other, and every design point on that curve implies a different plasticity mechanism.

Decision smoothing is the term most susceptible to unexamined inflation. Requiring that a classifier output persist for a dwell period before a command is issued is an entirely reasonable way to suppress spurious activations, and clinical systems often apply it because a stimulator firing at random is worse than one firing late. But it is a design choice, it is frequently unreported, and it can silently double the end-to-end latency of an otherwise well-characterised system.

The last two rows are the interesting ones. Even granting a hypothetical system with zero computational latency and no analysis window — a physical impossibility, but a useful limit case — an FES-based BCI cannot deliver cortical afferent input in less than roughly fifty to a hundred milliseconds after the command, because muscle takes time to contract and afferents take time to travel. There is a biological floor beneath the engineering. Any account of BCI-FES that depends on pairing within a twenty-millisecond window is therefore not merely difficult to implement; it is unachievable in principle for that class of device. Robotic and exoskeletal systems face a comparable floor, generally a worse one, because the mass of the limb and the actuator must both be accelerated before the patient feels anything; adaptive impedance control improves the therapeutic quality of that movement but does not make it arrive sooner [42]. Feedback delivered through a purely sensory channel — auditory or visual — escapes the mechanical floor and is correspondingly faster, though it also delivers a weaker and less specific afferent signal to sensorimotor cortex [44].

This is worth stating plainly because it changes what kind of claim the field can make. The mismatch is not a temporary artefact of immature technology that will resolve as processors improve. For peripherally actuated closed-loop systems, a meaningful part of it is permanent.

4. The Mismatch, Stated Plainly

Set the two preceding sections beside each other and the arithmetic is not subtle.

The potentiating window of timing-dependent plasticity is approximately twenty to forty milliseconds wide, with the sign of the effect reversing over an interval of a few milliseconds around coincidence. The end-to-end latency of a conventional closed-loop BCI is six hundred to three thousand milliseconds. We state this deliberately as a comparison of characteristic timescales rather than as a ratio of two commensurable intervals, for a reason developed in Section 9: the presynaptic term in a motor imagery BCI is a distributed cortical state rather than a discrete event, and a ratio would imply a precision of correspondence that the mapping does not support. The comparison of orders of magnitude is robust to that objection; a quoted multiple would not be.

Figure 1 places the relevant intervals on a single logarithmic axis. We have found that most colleagues, on seeing it, respond first that the axis must be wrong.

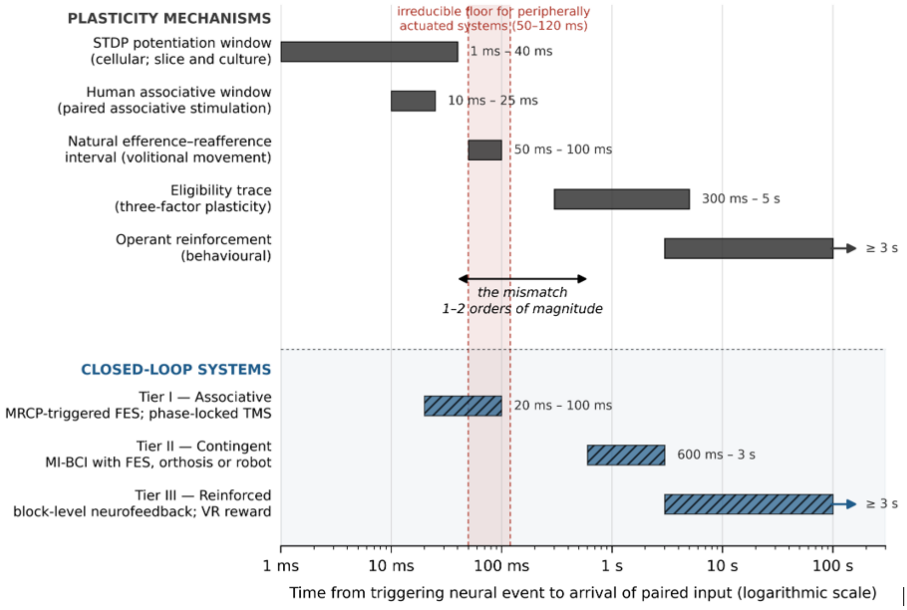

Figure 1. *Timescales of candidate plasticity mechanisms (upper band) and of closed-loop rehabilitation systems (lower band), on a common logarithmic axis. The potentiating window of spike-timing-dependent plasticity and the human associative window established by paired associative stimulation are both confined to tens of milliseconds. Conventional closed-loop BCI systems operate one to two orders of magnitude later, and are accommodated only by the eligibility-trace mechanism of three-factor plasticity. Tier definitions follow Table 2.*

There is a second problem, which we think is the more interesting of the two, and which we state with more care than we did in earlier drafts.

Because the timing rule is signed, arriving late is not necessarily a diluted version of arriving on time. In the anti-causal region — postsynaptic activity preceding presynaptic input — the rule predicts depression rather than weak potentiation. Whether a conventional BCI can enter that region at all depends entirely on an assignment the field has never made explicit: which population is to be treated as presynaptic and which as postsynaptic. If the afferent volley is the presynaptic term and the sustained cortical activity accompanying motor imagery is the postsynaptic term — the assignment implied by the PAS literature, where the afferent input arrives *before* the cortical event — then a system delivering afferent input hundreds of milliseconds *after* the onset of cortical activity is, formally, in the anti-causal configuration. If instead the cortical command is treated as presynaptic and the afferent return as postsynaptic, the ordering is causal but far outside the window.

We do not assert that conventional BCI therefore induces depression. We assert something weaker and, we think, more damaging: the field invokes a signed rule without specifying the sign convention, and the two available conventions give opposite predictions for the same device. A mechanism cited at this level of looseness cannot be doing the justificatory work it is asked to do. Whether the systems-level equivalent of anti-causal pairing carries any depressive consequence is not established, and so far as we can determine has never been asked. Prediction 2 in Section 7 is designed to ask it.

The third observation concerns the word closed-loop, which is doing a substantial amount of unexamined work in this literature. In control engineering, a closed loop is one in which the output is fed back and compared against a reference within the settling time of the system being controlled. Applied to synaptic plasticity, the relevant settling time is the plasticity window, and by that criterion most rehabilitation BCIs are not closed-loop systems with respect to the mechanism they claim to exploit. They are contingent systems: the feedback is reliably caused by the patient's neural state, which matters enormously for learning, but it does not arrive within the window that the cited rule specifies. Contingency and closed-loop control are different properties, and conflating them has allowed the field to import the mechanistic authority of the latter while delivering only the former.

We want to be careful here, because it would be easy to read the preceding paragraphs as a demolition, and that is not what they are. Contingency is a powerful teacher. Operant conditioning of cortical activity was demonstrated more than half a century ago, has been replicated across species and recording modalities since, and has never required millisecond precision to work [28, 29]. The distinction between a BCI that substitutes for lost function and one that restores it turns on exactly this — whether the system is trained to obey the patient or the patient is trained by the system [45, 46]. The problem is not that contingent systems fail to produce plasticity. The problem is that the field describes them using a mechanism they do not implement, and then draws no design conclusions from either the mechanism or its absence.

5. Three Resolutions, Only One of Which is Comfortable

5.1 Resolution one: it was never cellular STDP

The most straightforward response is that the literature has been speaking loosely, that nobody seriously believed individual synapses were being potentiated according to a twenty-millisecond rule, and that the intended claim was always about associative plasticity at the systems level — coincident activation of cortical motor representations and their afferent counterparts, over the hundreds of milliseconds that natural sensorimotor loops occupy.

This is probably true of what most authors mean. It is also insufficient, for three reasons.

First, systems-level associative plasticity is not a mechanism in the sense that STDP is a mechanism. It is a description of a phenomenon in search of one. Saying that coincident cortical and afferent activation strengthens their association restates the observation; it does not specify the rule, the window, the sign, or the conditions of failure. The reason authors reach for STDP is precisely that it supplies these things. Borrowing its specificity while disclaiming its constraints is not a defensible position.

Second, if the intended claim is about hundreds of milliseconds rather than tens, then the field owes an account of what the correct window is, and it does not have one. The interval between a voluntary motor command and its natural reafference is on the order of fifty to a hundred milliseconds. Systems operating at two seconds are not approximating that relationship either.

Third, and most practically: if timing does not matter at the tens-of-milliseconds scale, the field should stop citing evidence that depends on it. The associative BCI literature described in Section 2.3, the paired associative stimulation literature, and the primate spike-triggered stimulation work are all routinely cited as mechanistic support for conventional BCI-FES. They cannot function as support for a claim that explicitly discards their central variable.

5.2 Resolution two: eligibility traces, and why this is the good answer

The second resolution is the one we find genuinely persuasive, and it comes from computational and cellular neuroscience rather than from rehabilitation.

The problem the brain faces during reinforcement learning is formally identical to the problem we have been describing: a synapse must be modified on the basis of an outcome that arrives long after the activity that caused it. The solution, well characterised theoretically and increasingly well supported experimentally, is the eligibility trace — a transient, activity-induced synaptic tag that does not itself alter synaptic weight but renders the synapse susceptible to modification if a neuromodulatory signal arrives within a permissive interval [30, 31]. Coincident pre- and postsynaptic activity sets the tag; dopamine, acetylcholine or noradrenaline arriving subsequently converts it into a durable weight change.

The critical parameter is the lifetime of the trace, and it is far longer than the STDP window. Direct measurement in cortical synapses places the eligibility period for both potentiation and depression in the range of several seconds, with distinct traces for the two directions [31]. Work on dendritic spine enlargement gives a comparable figure, with dopamine effective when delivered within roughly one to two seconds of the triggering activity and ineffective beyond it [32]. Theoretical treatments converge on the same order of magnitude and make explicit that eligibility traces are the mechanism by which plasticity is reconciled with behavioural timescales [33].

This resolves the mismatch, and it does so in a way that is far more interesting than simply excusing it. It means that a BCI delivering feedback five hundred milliseconds after intent is not operating outside a plasticity mechanism; it is operating inside a different one, with different requirements. Specifically, three-factor plasticity requires a third factor. The neuromodulatory signal is not optional. If it is absent, weak, or mistimed, the eligibility trace decays and no weight change occurs regardless of how well the first two factors were paired.

The field adopted the language of Hebbian coincidence because it wanted a mechanism that explained why contingent feedback should work. Three-factor plasticity explains it better, tolerates the latencies these systems actually achieve, and specifies what else must be present. The design implications are immediate and underexploited. Reward structure ceases to be a motivational nicety and becomes a mechanistic requirement, which supplies a principled rather than merely intuitive rationale for gamified BCI environments [34], and a reason to treat adaptive difficulty as a plasticity intervention rather than an engagement feature [47]. Vagus nerve stimulation paired with rehabilitation moves from an empirically supported adjunct to a mechanistically predicted one, since it supplies precisely the cholinergic and noradrenergic third factor the model requires [35, 36]. Patient arousal, engagement, and dopaminergic tone become variables to measure rather than nuisances to control for. And the pharmacology of the stroke patient — the antipsychotics, the sedating antiemetics, the alpha-2 agonists that appear on drug charts without anyone considering their effect on plasticity — becomes a legitimate object of study.

There is also a constraint, and it is a sharp one. The eligibility trace has a lifetime of seconds, not tens of seconds. Systems operating at three to five seconds of end-to-end latency, or those in which reinforcement is delivered at the end of a trial block rather than contingently, may fall outside this window too. Eligibility traces widen the permissible interval by two orders of magnitude. They do not abolish it.

5.3 Resolution three: some of it is timing-critical, and we are discarding it

The third resolution is the one we cannot dismiss and would prefer to be wrong about.

The associative BCI protocols described in Section 2.3 achieve what conventional systems do not: they place the afferent volley inside the classical plasticity window, and they do so by decoding a feature — the movement-related cortical potential — that is available before the movement rather than during it. The predictive nature of the decoded signal is what buys the time. The reported effects on corticospinal excitability are consistent with the timing rule, scale with the precision of the pairing, and have translated into clinical benefit in controlled trials [18, 19, 20]. Similarly, phase-locked stimulation exploits an oscillatory feature that is intrinsically predictable a few tens of milliseconds ahead [15]. What unites these approaches is that they decode something about the immediate future rather than the immediate present, which is the only way to spend less time than the plasticity window allows. Magnetoencephalography, with its superior temporal and spatial resolution, is an obvious and underused platform for developing such predictive decoders in stroke [48].

If cellular-scale timing genuinely contributes, then conventional systems are leaving a mechanism on the table — not because of any physical impossibility, but because the field has treated latency as an incidental property rather than a target. The uncomfortable implication is that a decade of BCI-FES trials may have been measuring the operant and associative components of an intervention while omitting the timing-dependent component entirely, and reporting the result as the efficacy of the whole.

5.4 What follows if all three are partly true

We think all three are partly true, and we recognise this is the least satisfying conclusion available.

Conventional BCI-FES most likely produces its benefit through operant reinforcement of a cortical state, gated by neuromodulatory eligibility, with a systems-level associative component operating over hundreds of milliseconds. Cellular-scale timing-dependent plasticity contributes little or nothing to these systems. This does not make them ineffective; it makes them mislabelled, and it means their optimisation should be pursued along axes the field currently neglects — reward, arousal, neuromodulatory state, pharmacological context — rather than the axis it names but does not measure.

A distinct class of intervention, exemplified by MRCP-triggered and phase-locked stimulation, does exploit cellular-scale timing, does require millisecond precision, and should be developed and evaluated separately. Pooling the two in meta-analyses produces an average of two different treatments and an estimate that describes neither.

6. A Timescale Taxonomy of Closed-Loop Rehabilitation Systems

If the operative plasticity mechanism depends on latency, then latency should organise how we classify these systems. It does not currently organise anything. We propose the following taxonomy, offered as a working instrument rather than a settled classification, and note that it cuts across the conventional MI/MA/SMR paradigm distinction rather than replacing it.

Table 2. Proposed classification of closed-loop rehabilitation systems by end-to-end latency and the plasticity mechanism plausibly recruited at that timescale.

Tier	End-to-end latency	Plausible dominant mechanism	Representative implementations	Design priority
Tier I: Associative	< 100 ms	Timing-dependent plasticity; classical associative window	MRCP-triggered peripheral stimulation; phase-locked TMS or FES; spike-triggered stimulation	Precision of the interval; predictive decoding of pre-movement features
Tier II: Contingent	100 ms – 3 s	Three-factor plasticity; eligibility traces gated by neuromodulation	Conventional MI-BCI with FES, orthosis or exoskeleton	Neuromodulatory state; reward structure; reliability of contingency
Tier III: Reinforced	3 s – trial level	Operant conditioning; declarative and strategic learning	Neurofeedback with block-level scoring; VR reward at task completion	Motivation, engagement, practice volume
Tier 0: Non-contingent	Not applicable	Afferent stimulation without cortical pairing	Fixed-schedule FES; open-loop robotics	Not a closed-loop intervention; should not be pooled with the above

The taxonomy makes several things visible. Tier 0 systems are not BCIs in any meaningful sense, yet they appear in comparisons and occasionally in meta-analyses under the BCI-FES heading. This is more consequential than a labelling error: it dilutes effect estimates with an intervention that lacks the mechanism under investigation.

Tier I and Tier II are different treatments. They require different equipment, different expertise, and different outcome logic. A trial demonstrating Tier I efficacy provides no evidence for Tier II and vice versa.

The taxonomy also reframes non-response. A patient who fails to achieve control on a Tier II system has failed at a task requiring sustained self-regulation over seconds under neuromodulatory conditions nobody measured; the same patient on a Tier I system faces a task requiring only that a pre-movement potential be detectable. These are not the same failure, and treating them as one has helped keep so-called BCI illiteracy a technical mystery rather than a tractable clinical question [49, 50].

Most importantly, the taxonomy is falsifiable in a way that paradigm-based classification is not. It predicts that within-tier heterogeneity in outcome should be smaller than between-tier heterogeneity, and that specific interventions should shift outcome in tier-specific ways — neuromodulatory priming should help Tier II disproportionately, while millisecond jitter should degrade Tier I disproportionately and leave Tier III untouched.

7. Three Falsifiable Predictions

Following the practice we adopted in earlier work, we state our position in a form that permits it to be shown wrong. Each prediction is testable with existing technology; the sample sizes given are illustrative, derived from the assumptions stated, and are intended to establish feasibility rather than to substitute for a protocol-stage power calculation.

Prediction 1 — the latency gradient. Within Tier I systems, deliberately increasing the interval between detected motor intent and afferent arrival from approximately 25 ms to approximately 200 ms will substantially attenuate the induced change in corticospinal excitability, with the relationship following an approximately exponential decay. Within Tier II systems, the same manipulation applied across the range 300 ms to 1500 ms will produce no detectable difference in outcome. If instead a smooth latency–response relationship is observed across the whole range, our tier boundary is wrong and the mechanism is continuous rather than categorical. *Design:* within-participant, counterbalanced, MEP amplitude as the primary outcome. On the effect sizes reported for interval-dependent PAS [13, 14] — differences of roughly one standard deviation between optimal and non-optimal intervals — a within-participant design with four latency levels detects the gradient at 80% power with approximately 20 participants.

Prediction 2 — the sign reversal. In a Tier I associative protocol, deliberately delivering afferent stimulation in the anti-causal configuration will produce a measurable reduction in corticospinal excitability relative to sham, rather than merely a smaller increase. This is the strongest test of whether the classical rule genuinely governs these systems. *Staging and ethics:* deliberately inducing an LTD-like change in a stroke population is not a first-in-human step, and we do not propose it as one. The prediction should be tested first in healthy participants, where the equivalent manipulation is already established and reversible within a session [14]; extension to a stroke population is justified only if the healthy-participant effect is confirmed, and then only with a single session, an inhibitory-direction outcome measured neurophysiologically rather than functionally, and pre-specified stopping rules. A null result in healthy participants would substantially weaken Resolution three and should end the line of enquiry.

Prediction 3 — neuromodulatory rescue. In Tier II systems, pairing with a neuromodulatory intervention that supplies the third factor — vagus nerve stimulation being the best-evidenced candidate — will produce a disproportionately larger benefit than the same intervention paired with a Tier I system, because Tier I already achieves plasticity through a pathway that does not depend on eligibility-trace conversion. If VNS benefit is equivalent across tiers, the three-factor account is not doing the explanatory work we attribute to it. *Feasibility:* this is the most demanding of the three, because it requires a tier-by-intervention interaction rather than a main effect. Taking the between-group difference reported in VNS-REHAB [35] as the Tier II effect and assuming the interaction is of similar magnitude, detection at 80% power requires on the order of four times the sample of the original trial — beyond a single centre, and realistically a candidate for a multi-site or registry-embedded design. We state it nonetheless, because a prediction that is expensive to test is not thereby unfalsifiable, and because a smaller mechanistic surrogate — cortical excitability rather than function — could provide an early signal.

8. What Should Change in Reporting

8.1 A minimum latency specification

The recommendation that follows from all of the above is unglamorous and overdue. Latency should be reported as a primary system specification in every clinical BCI publication, with the same routineness as sample size. Table 3 sets out the minimum set we propose. It is deliberately short, and every item is measurable with equipment already present in any laboratory running these systems.

Table 3. Proposed minimum latency reporting set for closed-loop neurorehabilitation studies. Items are stated so that absence can be reported explicitly rather than by omission; where a parameter does not apply to the system under study, that should be recorded rather than left blank.

#	Item	Definition	Why it matters
1	Analysis window length and overlap	Duration of data accumulated per decision, and the step between successive windows	The dominant and least reducible term in the budget; determines the tier
2	Decoded feature	Band power, connectivity, MRCP, oscillatory phase, hybrid EEG–EMG	Predictive features permit Tier I latencies; state features do not
3	Dwell-time or smoothing criterion	Duration and rule by which a classifier output must persist before a command issues	Frequently unreported; can double end-to-end latency
4	Command-to-actuator-onset delay	Measured interval from command issue to mechanical or electrical output onset	Device-specific; not inferable from the paradigm description
5	Stimulation parameters sufficient to estimate contraction onset	Pulse width, frequency, intensity relative to motor threshold	Allows the electromechanical delay term to be reconstructed by readers
6	Measured end-to-end latency	Neural feature onset to peripheral afferent arrival at cortex, measured rather than inferred	The single figure that determines which plasticity mechanism is available
7	Trial-to-trial variability of item 6	Standard deviation or interquartile range across trials	A mean of 400 ms with an SD of 300 ms delivers consistent pairing at no timescale
8	Method used to measure items 4 and 6	Oscilloscope, synchronised triggers, high-speed video, or inference from specification	Distinguishes measured from assumed values
9	Tier assignment	Tier 0, I, II or III per Table 2, with the latency figure that justifies it	Enables like-with-like pooling in synthesis
10	Reward and neuromodulatory context	Reward schedule; concurrent VNS or equivalent; plasticity-relevant medication	The third factor, without which the Tier II mechanism does not operate

None of this is difficult to measure. That it is almost never reported reflects the fact that nobody has been asked for it, which is a problem that editors and reviewers are unusually well placed to solve. We would encourage journals publishing in this area to adopt items 1, 3, 6 and 9 as a minimum condition of methodological completeness.

8.2 A pre-specified protocol for the companion synthesis

Our claim that these parameters are seldom reported rests, at present, on our own reading of the literature rather than on a documented count, and we are explicit that this is the principal empirical limitation of the present paper. We therefore state here the protocol for the synthesis that would settle it, which we are undertaking separately and which we recommend to any group attempting a comparable exercise.

Search MEDLINE, Embase, IEEE Xplore and Scopus for controlled trials of closed-loop BCI in stroke from 2005 to present. For each included report, extract verbatim any statement corresponding to items 1 to 10 of Table 3, recording absence explicitly rather than by omission, and consult supplementary materials and prior methodological papers from the same group before recording a value as unreported. Where a figure can be inferred but is not stated, record it as inferred and give the basis. Report the proportion of trials in which each item is recoverable, and assign each trial to a tier where the data permit. Reporting should follow PRISMA-ScR [55]. The trials at references [19] and [22–27] constitute a natural pilot set, spanning FES, orthotic, robotic and visual-feedback implementations and both Tier I and Tier II designs.

We note that the outcome of this exercise is not neutral with respect to our argument. If end-to-end latency proves to be recoverable in a majority of trials, the reporting recommendation in Section 8.1 loses much of its force, though the taxonomy and the mechanistic argument do not.

9. Limitations and the Counterargument we Find Hardest

Several objections deserve to be stated in their strongest form rather than their most convenient one.

The most serious is this: the entire argument assumes that the relevant temporal relationship is between the BCI's detection event and the arrival of afferent input at cortex, and that this maps onto the pre- and postsynaptic activity of the STDP rule. It may not. Cortical neurons during motor imagery are not firing a single well-defined spike whose timing can be meaningfully compared with an afferent volley; they are engaged in sustained, distributed, rhythmic activity lasting seconds. If the relevant presynaptic activity is a state rather than an event, then the notion of an interval between it and the afferent input may be ill-posed, and both our critique and the claim it criticises would rest on the same category error.

We think this objection has real force, and we have taken two steps in response. The first is to state the mismatch as a comparison of characteristic timescales rather than as a numerical ratio, since a ratio presupposes exactly the correspondence in question (Section 4). The second is to observe that the objection does not cut symmetrically. Our recommendations — report latency, classify by timescale, design for the third factor — survive the objection intact, because they do not depend on the interval being well defined. The claim we are criticising does depend on it: a rule defined over an interval cannot be invoked as mechanistic warrant if the interval is ill-posed. The objection, if correct, removes the field's justification and leaves our prescription standing.

A second objection is that the magnitude of the mismatch may not matter if the intervention is repeated thousands of times. Weak, poorly timed pairings delivered with sufficient volume might accumulate into meaningful change. This is plausible, though it sits awkwardly with the observation that the classical rule is signed rather than merely graded, and it makes a prediction the field could test: outcome should scale with repetition count more steeply in Tier II than in Tier I systems.

A third objection comes from the metaplasticity literature and is more subtle than it first appears. The state of the cortex at the moment of pairing determines what any given pairing does, and post-stroke cortex is not in a stable state: excitability, inhibitory tone and the threshold for potentiation all shift over the weeks following infarction [51, 52]. A protocol delivering identical timing to the same patient at two weeks and at six months may produce opposite effects for reasons that have nothing to do with latency. This does not undermine our argument, but it does mean that latency is one of at least two temporal variables in play, the other being where in the recovery trajectory the intervention lands [53, 54]. Any serious test of our predictions must control for both.

A fourth is that we have concentrated on EEG-based systems with peripheral actuators, and the analysis does not transfer straightforwardly to intracortical BCIs, where both the recorded signal and the delivered stimulation can operate at the relevant timescale. That is correct, and it is arguably the strongest argument for pursuing invasive approaches in severe stroke, though the risk-benefit calculation there is a separate discussion.

Fifth, as set out in Section 8.2, the systematic extraction of latency reporting from the trial literature has not yet been completed, and until it is, the claim that the field does not report latency rests on our reading of it. We regard the completion of that synthesis as the immediate next step and have specified its protocol so that others may perform it independently, or contradict us.

We should also record where the two of us disagree. One of us (GP) considers the eligibility-trace resolution close to decisive and the timing-critical component correspondingly minor, on the grounds that the engineering constraints described in Section 3 make Tier I implementations impractical outside research settings. One of us (AD) considers this too quick, and suspects that a decade of pooling has hidden a genuine effect that would be visible if Tier I systems were evaluated separately. The text reflects both positions. We did not resolve it, and readers should treat Section 5.4 as a shared position on an unshared question.

10. Conclusion

The claim that closed-loop brain–computer interfaces work by inducing spike-timing-dependent plasticity has been repeated often enough to stop being examined. Examining it produces an uncomfortable result: the mechanism operates over tens of milliseconds, the systems operate over hundreds or thousands, and for peripherally actuated devices a substantial part of the gap is not closable by any engineering effort.

We do not conclude that the field is wrong about efficacy. We conclude that it is wrong about mechanism, and that the error is expensive. It has made latency invisible as a design parameter; it has permitted the pooling of interventions whose mechanisms differ by orders of magnitude in timescale; it has obscured the fact that the most plausible actual mechanism — neuromodulator-gated plasticity operating on eligibility traces — has design requirements the field is not meeting, concerning reward, arousal and pharmacological state; and it has allowed a small body of genuinely timing-critical work to be filed alongside interventions that discard the variable it was built around.

The remedy is not complicated. Measure and report latency, using something like the ten items in Table 3. Classify systems by the timescale they actually operate on. Stop pooling interventions that differ by two orders of magnitude in the property their shared mechanistic justification depends on. And, where the mechanism turns out to be three-factor rather than two-factor, design for the third factor deliberately rather than hoping the patient supplies it.

We began with a sentence that everyone writes and nobody checks. We are not certain that checking it will change what the field does. We are reasonably certain that not checking it has already cost more than anyone has counted.

Declarations

Author Contributions

Recorded using the CRediT taxonomy. AD: Conceptualisation; writing — original draft; writing — review and editing; visualisation. VK, SK, GG: Conceptualisation; methodology; supervision; writing — review and editing; resources. All authors read and approved the final version of the manuscript.

Acknowledgements

We thank our parents & teachers for their blessings and best wishes.

Funding

Authors: None of the authors received any funding for this paper.

Competing Interests

Authors: No competing interests for any of the authors.

Data Availability Statement

This manuscript is a narrative and methodological review. No new primary datasets were generated. The reporting set presented as Table 3 and the protocol in Section 8.2 are provided to support independent replication of the planned synthesis.

Ethics Statement

This manuscript is a review article and does not report primary research involving human participants, animals, or biological specimens. No ethics approval was required. The research proposed in Section 7 would require ethical approval in each jurisdiction where it is undertaken; the staging and safeguards we consider necessary for Prediction 2 are set out there.

Abbreviations

BCI, brain–computer interface; EEG, electroencephalography; EMG, electromyography; ERD, event-related desynchronisation; ERS, event-related synchronisation; FES, functional electrical stimulation; LTD, long-term depression; LTP, long-term potentiation; MEP, motor evoked potential; MI, motor imagery; MRCP, movement-related cortical potential; PAS, paired associative stimulation; PRISMA-ScR, Preferred Reporting Items for Systematic Reviews and Meta-Analyses extension for Scoping Reviews; STDP, spike-timing-dependent plasticity; TMS, transcranial magnetic stimulation; VNS, vagus nerve stimulation; VR, virtual reality.

References

1. Markram H, Lübke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science. 1997;275(5297):213–215.

2. Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci. 1998;18(24):10464–10472.

3. Feldman DE. The spike-timing dependence of plasticity. Neuron. 2012;75(4):556–571.

4. Caporale N, Dan Y. Spike timing-dependent plasticity: a Hebbian learning rule. Annu Rev Neurosci. 2008;31:25–46.

5. Cervera MA, Soekadar SR, Ushiba J, et al. Brain-computer interfaces for post-stroke motor rehabilitation: a meta-analysis. Ann Clin Transl Neurol. 2018;5(5):651–663.

6. Bai Z, Fong KNK, Zhang JJ, Chan J, Ting KH. Immediate and long-term effects of BCI-based rehabilitation of the upper extremity after stroke: a systematic review and meta-analysis. J Neuroeng Rehabil. 2020;17(1):57.

7. Hebb DO. The Organization of Behavior: A Neuropsychological Theory. New York: Wiley; 1949.

8. Dan Y, Poo MM. Spike timing-dependent plasticity: from synapse to perception. Physiol Rev. 2006;86(3):1033–1048.

9. Shulz DE, Jacob V. Spike-timing-dependent plasticity in the intact brain: counteracting spurious spike coincidences. Front Synaptic Neurosci. 2010;2:137.

10. Pawlak V, Wickens JR, Kirkwood A, Kerr JND. Timing is not everything: neuromodulation opens the STDP gate. Front Synaptic Neurosci. 2010;2:146.

11. Sjöström PJ, Turrigiano GG, Nelson SB. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron. 2001;32(6):1149–1164.

12. Lisman J, Spruston N. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. Nat Neurosci. 2005;8(7):839–841.

13. Stefan K, Kunesch E, Cohen LG, Benecke R, Classen J. Induction of plasticity in the human motor

cortex by paired associative stimulation. Brain. 2000;123(3):572–584.

14. Wolters A, Sandbrink F, Schlottmann A, et al. A temporally asymmetric Hebbian rule governing plasticity in the human motor cortex. J Neurophysiol. 2003;89(5):2339–2345.

15. Zrenner C, Desideri D, Belardinelli P, Ziemann U. Real-time EEG-defined excitability states determine efficacy of TMS-induced plasticity in human motor cortex. Brain Stimul. 2018;11(2):374–389.

16. Jackson A, Mavoori J, Fetz EE. Long-term motor cortex plasticity induced by an electronic neural implant. Nature. 2006;444(7115):56–60.

17. Nishimura Y, Perlmutter SI, Eaton RW, Fetz EE. Spike-timing-dependent plasticity in primate corticospinal connections induced during free behavior. Neuron. 2013;80(5):1301–1309.

18. Mrachacz-Kersting N, Kristensen SR, Niazi IK, Farina D. Precise temporal association between cortical potentials evoked by motor imagination and afference induces cortical plasticity. J Physiol. 2012;590(7):1669–1682.

19. Mrachacz-Kersting N, Jiang N, Stevenson AJT, et al. Efficient neuroplasticity induction in chronic stroke patients by an associative brain-computer interface. J Neurophysiol. 2016;115(3):1410–1421.

20. Mrachacz-Kersting N, Stevenson AJT, Jørgensen HRM, et al. Brain state-dependent stimulation boosts functional recovery following stroke. Ann Neurol. 2019;85(1):84–95.

21. Pfurtscheller G, Lopes da Silva FH. Event-related EEG/MEG synchronization and desynchronization: basic principles. Clin Neurophysiol. 1999;110(11):1842–1857.

22. Biasiucci A, Leeb R, Iturrate I, et al. Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke. Nat Commun. 2018;9(1):2421.

23. Ramos-Murguialday A, Broetz D, Rea M, et al. Brain-machine interface in chronic stroke rehabilitation: a controlled study. Ann Neurol. 2013;74(1):100–108.

24. Ang KK, Chua KSG, Phua KS, et al. A randomised controlled trial of EEG-based motor imagery brain-computer interface robotic rehabilitation for stroke. Clin EEG Neurosci. 2015;46(4):310–320.

25. Pichiorri F, Morone G, Petti M, et al. Brain-computer interface boosts motor imagery practice during stroke recovery. Ann Neurol. 2015;77(5):851–865.

26. Frolov AA, Mokienko O, Lyukmanov R, et al. Post-stroke rehabilitation training with a motor-imagery-based brain-computer interface (BCI)-controlled hand exoskeleton: a randomised controlled multicenter trial. Front Neurosci. 2017;11:400.

27. Cheng N, Phua KS, Lai HS, et al. Brain-computer interface-based soft robotic glove rehabilitation for stroke. IEEE Trans Biomed Eng. 2020;67(12):3339–3351.

28. Fetz EE. Operant conditioning of cortical unit activity. Science. 1969;163(3870):955–958.

29. Sitaram R, Ros T, Stoeckel L, et al. Closed-loop brain training: the science of neurofeedback. Nat Rev Neurosci. 2017;18(2):86–100.

30. Frémaux N, Gerstner W. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. Front Neural Circuits. 2016;9:85.

31. He K, Huertas M, Hong SZ, et al. Distinct eligibility traces for LTP and LTD in cortical synapses. Neuron. 2015;88(3):528–538.

32. Yagishita S, Hayashi-Takagi A, Ellis-Davies GCR, et al. A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science. 2014;345(6204):1616–1620.

33. Gerstner W, Lehmann M, Liakoni V, Corneil D, Brea J. Eligibility traces and plasticity on behavioral time

scales: experimental support of neoHebbian three-factor learning rules. Front Neural Circuits. 2018;12:53.

34. Ma Z-Z, Wu J-J, Cao Z, et al. Motor imagery-based brain–computer interface rehabilitation programmes enhance upper extremity performance and cortical activation in stroke patients. J Neuroeng Rehabil. 2024;21(1):91.

35. Dawson J, Liu CY, Francisco GE, et al. Vagus nerve stimulation paired with rehabilitation for upper limb motor function after ischaemic stroke (VNS-REHAB): a randomised, blinded, pivotal, device trial. Lancet. 2021;397(10284):1545–1553.

36. Hays SA, Rennaker RL, Kilgard MP. Targeting plasticity with vagus nerve stimulation to treat neurological disease. Prog Brain Res. 2013;207:275–299.

37. Niazi IK, Mrachacz-Kersting N, Jiang N, Dremstrup K, Farina D. Peripheral electrical stimulation triggered by self-paced detection of motor intention enhances motor evoked potentials. IEEE Trans Neural Syst Rehabil Eng. 2012;20(4):595–604.

38. Xu R, Jiang N, Mrachacz-Kersting N, et al. A closed-loop brain-computer interface triggering an active ankle-foot orthosis for inducing cortical neural plasticity. IEEE Trans Biomed Eng. 2014;61(7):2092–2101.

39. Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G, Vaughan TM. Brain-computer interfaces for communication and control. Clin Neurophysiol. 2002;113(6):767–791.

40. Prasad G, Herman P, Coyle D, McDonough S, Crosbie J. Applying a brain-computer interface to support motor imagery practice in people with stroke for upper limb recovery: a feasibility study. J Neuroeng Rehabil. 2010;7:60.

41. Chowdhury A, Raza H, Meena YK, Dutta A, Prasad G. An EEG-EMG correlation-based brain-computer interface for hand orthosis supported neuro-rehabilitation. J Neurosci Methods. 2019;312:1–11.

42. Chowdhury A, Nishad SS, Meena YK, Dutta A, Prasad G. Hand-exoskeleton assisted progressive neurorehabilitation using impedance adaptation based challenge level adjustment method. IEEE Trans Haptics. 2019;12(2):128–140.

43. Rathee D, Chowdhury A, Meena YK, Dutta A, McDonough S, Prasad G. Brain-machine interface driven post-stroke upper-limb functional recovery correlates with beta-band mediated cortical networks. IEEE Trans Neural Syst Rehabil Eng. 2019;27(5):1020–1031.

44. McCreadie KA, Coyle DH, Prasad G. Is sensorimotor BCI performance influenced differently by mono, stereo, or 3-D auditory feedback? IEEE Trans Neural Syst Rehabil Eng. 2014;22(3):431–440.

45. Grosse-Wentrup M, Mattia D, Oweiss K. Using brain-computer interfaces to induce neural plasticity and restore function. J Neural Eng. 2011;8(2):025004.

46. Gharabaghi A. What turns assistive into restorative brain-machine interfaces? Front Neurosci. 2016;10:456.

47. Bauer R, Gharabaghi A. Reinforcement learning for adaptive threshold control of restorative brain-computer interfaces: a Bayesian simulation. Front Neurosci. 2015;9:36.

48. Buch E, Weber C, Cohen LG, et al. Think to move: a neuromagnetic brain-computer interface (BCI) system for chronic stroke. Stroke. 2008;39(3):910–917.

49. Soekadar SR, Birbaumer N, Slutzky MW, Cohen LG. Brain–machine interfaces in neurorehabilitation of stroke. Neurobiol Dis. 2015;83:172–179.

50. Vidaurre C, Blankertz B. Towards a cure for BCI illiteracy. Brain Topogr. 2010;23(2):194–198.

51. Ziemann U, Siebner HR. Modifying motor learning through gating and homeostatic metaplasticity. Brain Stimul. 2008;1(1):60–66.

52. Turrigiano GG, Nelson SB. Homeostatic plasticity in the developing nervous system. Nat Rev Neurosci. 2004;5(2):97–107.

53. Zeiler SR, Krakauer JW. The interaction between training and plasticity in the post-stroke brain. Curr Opin Neurol. 2013;26(6):609–616.

54. Krakauer JW, Carmichael ST, Corbett D, Wittenberg GF. Getting neurorehabilitation right: what can be learned from animal models? Neurorehabil Neural Repair. 2020;34(1):13–22.

55. Tricco AC, Lillie E, Zarin W, et al. PRISMA extension for scoping reviews (PRISMA-ScR): checklist and explanation. Ann Intern Med. 2018;169(7):467–473.