Title
From Neural Signal to Embodied Action: A Cognitive Framework for BCI-Mediated Intelligence

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Author
ding, yueyang

Publication Date
2026

From Neural Signal to Embodied Action: A Cognitive Framework for BCI-Mediated Intelligence

Anonymous CogSci submission

Abstract:Brain–computer interfaces (BCIs) are typically framed as channels for neural signal transmission, yet their implications for embodiment and agency remain insufficiently theorized. We argue that BCI-mediated systems provide a productive testbed for embodied and extended cognition. Drawing on enactivist and 4E theories, we propose five interaction paradigms: (1)behavior-based embodiment through minimal sensorimotor coupling;

(2)shared agency via human–machine co-negotiation;

(3)structural coupling through closed-loop co-adaptive plasticity;

(4)perceptual extension via integration of non-biological sensors;

(5)intentional binding through temporally aligned multimodal feedback.

We contend that BCIs function not as passive conduits but as active scaffolds that reconfigure embodiment, transforming tools into quasi-bodily extensions[2] while preserving cognitive sovereignty. This framework positions BCI-mediated intelligence as a dynamic interplay between neural intention and embodied action, offering a principled approach to studying agency, selfhood, and human–machine integration within 4E cognition.[23]

Keywords embodied intelligence, enactivism, sense of agency, extended cognition, human-robot interaction, brain-computer interface, sensorimotor coupling

Introduction

Research Background

If cognition is fundamentally embodied—emerging from the dynamic coupling of brain, body, and environment—then brain–computer interfaces (BCIs) pose a prima facie challenge: they enable users to control external agents through neural activity alone, seemingly bypassing bodily action. From an enactivist perspective, this appears to contradict the claim that cognition arises through sensorimotor engagement with the world (Varela et al., 1991).

We argue, however, that BCI-mediated interaction does not negate embodiment but reconfigures it. Neural signals in BCIs function not as disembodied commands, but as entry points into novel sensorimotor loops. Through sustained interaction[6], external agents such as robotic arms or virtual avatars can be incorporated into the user's peripersonal space, shifting from tools to quasi-bodily extensions. In this sense, BCIs instantiate forms of mediated embodiment consistent with extended and enactive accounts of cognition.[27]

Core Claim

In this paper, we propose a cognitive framework that reconceptualizes common BCI–agent interaction paradigms not as engineering architectures, but as distinct forms of emergent embodied intelligence. We identify five paradigms—behavior-based embodiment, shared agency, structural coupling, perceptual extension, and intentional binding—each reflecting key principles of 4E cognition (embodied, embedded, enactive, extended). Rather than treating BCIs as passive conduits for mental commands, we position them as active mediators that scaffold the co-construction of agency between human and machine.

In the following sections, we unpack these five paradigms and analyze how they reshape core constructs in cognitive science, including embodiment, agency, and perception.

A Cognitive Framework for BCI-Mediated Embodiment

Overview of BCI–Agent Interaction Paradigms

To move beyond abstract theorizing, we propose a quantifiable cognitive framework that reinterprets common BCI-agent interaction paradigms as distinct forms of embodied intelligence. These paradigms are not merely engineering architectures but represent analytically distinct patterns of real-time coupling. We summarize these five modes in Table 1, followed by a detailed exposition of each paradigm and its corresponding quantitative metrics.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Table 1: Five Modes of Embodied Human-Agent Interaction

Mode	Description
1. Behavior-Based Embodiment	Intelligent action arises from sensorimotor loops without explicit internal representations; neural activity participates in reactive sensorimotor loops, and feedback shapes neural states—forming a continuous perception–action loop.
2. Shared Agency through Complementary Control	Control is distributed by cognitive strengths: user sets high-level goals (e.g., "go left"), agent handles low-level execution (e.g., obstacle avoidance), grounded in joint action theory.
3. Structural Coupling via Closed-Loop Plasticity	Sustained interaction leads to mutual adaptation (Varela, 1979); user and agent co-evolve, reshaping cognitive habits for fluent coordination.
4. Perceptual Extension and Tool Incorporation	BCI-mediated agents act as sensory extensions; tactile or visual feedback from robots/avatars integrates into the user's perceptual field.
5. Intentional Binding and Sense of Agency	Multimodal feedback stabilizes agency ("I caused that") when intentions precede, match, and uniquely explain actions (Wegner, 2003).

Together, these modes capture a progression from minimal sensorimotor coupling to stabilized forms of extended agency.

Specifically, we identify five modes of BCI-mediated embodiment, each interpreted through the lens of **body-perception-environment (BPE) dynamics** and quantified by three key metrics: *geometric constraint* (C_{geo}), *perceptual uncertainty* (U_{percept}), and *environmental dynamicity* (D_{env}):

- **Geometric Constraint (C_{geo}):** Measures the degree to which the agent's action space is physically bounded.
- **Perceptual Uncertainty (U_{percept}):** Captures the fidelity of sensory feedback relative to ground truth.
- **Environmental Dynamicity (D_{env}):** Reflects the rate and magnitude of environmental change encountered during interaction.

- **(1) Behavior-Based Embodiment**—Action emerges from reactive sensorimotor loops without explicit internal representation, where the agent's behavior is fundamentally constrained by C_{geo} (i.e., geometric limitations imposed by the physical body or workspace). In this mode, the "body" dominates perception and environmental interaction, as sensorimotor feedback directly modulates neural states without symbolic abstraction.
- **(2) Shared Agency through Complementary Control**—Human and machine divide cognitive labor across hierarchical control levels (e.g., user sets high-level goals, agent handles low-level execution), mitigating U_{percept} by leveraging the machine's precision in perceptual processing while preserving the user's intuitive sense of agency. This mode exemplifies **distributed cognition** across the BPE triad.
- **(3) Structural Coupling via Closed-Loop Plasticity**—Sustained interaction drives mutual adaptation between user and agent, where D_{env} (dynamic environmental changes) and U_{percept} jointly shape neural plasticity. Over time, the user's body schema co-evolves with the agent's control policies, creating a closed loop of co-adaptation that blurs the boundary between biological and artificial agency.
- **(4) Perceptual Extension and Tool Incorporation**—Artificial sensors become integrated into the user's perceptual field, expanding the "perception" dimension of BPE. This reduces U_{percept} by enriching sensory input and redefines the "body" to include the tool as an extension of the user's sensorimotor apparatus.
- **(5) Intentional Binding and the Sense of Agency**—Multimodal feedback stabilizes the experience of authorship over action, unifying C_{geo}, U_{percept}, and D_{env} to construct a coherent sense of agency. When feedback aligns with intention, U_{percept} diminishes, and the "body" (user + agent) is experienced as a single, intentional entity interacting with the environment.

In the following subsections, we unpack each mode as a window into how **embodiment is dynamically reconstructed** in technologically mediated cognition—revealing how BCI systems reconfigure the BPE triad through the interplay of C_{geo}, U_{percept}, and D_{env}.
amsmath

Quantitative Metrics: The CDU Framework

To formalize the interaction between the agent and its environment, we define three core quantitative metrics: Geometric Constraint (C), Perceptual Uncertainty (U), and Environmental Dynamicity (D).

$$C_{\text{geo}} = 1 - \frac{\text{DoF}_{\text{remaining}}}{\text{DoF}_{\text{initial}}} \tag{1}$$

$$U_{\text{percept}} = \frac{1}{1 + e^{k(x_0 - T_{rp})}} \tag{2}$$

$$D_{env} = \frac{1}{N}\sum_{i=1}^{N} \frac{\|V_i\|}{\|V\|_{\max}} \tag{3}$$

$$D_{\text{env}}^{\text{coupled}} = \frac{1}{N}\sum_{i=1}^{N} \left(\frac{\|v_i\|}{\|v\|_{\max}} \cdot \left[1 + \lambda \frac{a_i \cdot (-p_i)}{\|v\|_{\max}\|p_i\|} \right] \right) \tag{4}$$

Interpretation C_{geo} quantifies the reduction of motor freedom due to physical constraints (e.g., joint limits). U_{percept} models the inverse relationship between sensory noise and perceptual accuracy, modulated by the deviation from the true representation parameter T_{rp}. D_{env} captures the average intensity of environmental motion relative to system limits. The coupled formulation $D_{\text{env}}^{\text{coupled}}$ incorporates acceleration (a_i) to model how *changes in velocity* modulate perceived dynamicity, with λ acting as a gain controller.

System State Function The overall system state S (survival or performance efficacy) emerges from the interaction of these metrics:

$$S(C, U, D_{env}) = S_{\max} \cdot \left(\frac{C}{C - \kappa_t \cdot U} \right) \cdot \left(\frac{1}{1 - \lambda \cdot D_m} \right) \quad (5)$$

Eq. 5 integrates geometric constraint, perceptual uncertainty, and environmental dynamicity. The parameter κ_t represents the temporal sensitivity of the constraint, while D_m is the maximum tolerable dynamicity. This formulation predicts a critical phase transition: as U approaches C/κ_t, the denominator vanishes, causing system collapse ($S \to 0$).

Parameter Sensitivity: Empirical Validation

Together, these results ground the CDU framework in empirical dynamics:

• Figure 1 shows how environmental dynamics (D) are modulated by λ, revealing the agent's capacity to anticipatechange through coupled constraint-uncertainty feedback.

• Figure 2 reveals how perceptual uncertainty (U) is **nonlinearly modulated** by k_{danger}, exposing the cognitive cost of precisionin threat representation**and the trade-off inherent in high-gain regulation**.

Both figures demonstrate that CDU is not a static taxonomy, but a dynamic control system — where parameters like λ and k danger act as embodied regulators that continuously reconfigure the agent's state (S) in response to environmental and perceptual constraints.

Experiment 1: Acceleration Weight (λ) Variation We compared three experimental groups to examine how acceleration weighting affects behavioral outcomes (see Fig. 1):

- Group A ($\lambda = 0$): Velocity-only response; baseline behavior with delayed threat detection.
- Group B ($\lambda = 1$): Acceleration-dominated response; hyperreactive behavior causing instability.
- Group C ($\lambda = 0.5$): Balanced integration; anticipatory and robust navigation.

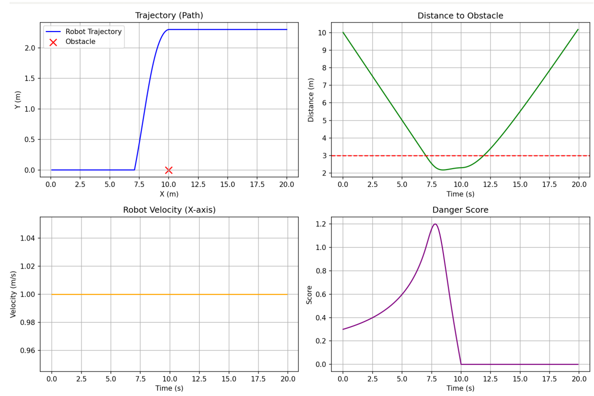

Figure 1: Acceleration weighting (λ) modulates environmental dynamics to govern navigational resilience. Trajectories under varying λ values reveal how the coupling term in $D_{\text{env}}^{\text{coupled}}$ (Eq. 4) reshapes behavioral trade-offs between velocity stability and threat anticipation. At $\lambda = 0$, agents exhibit delayed responses, reflecting low dynamic sensitivity. At $\lambda = 1$, trajectories become erratic, indicating overcompensation—a hallmark of high dynamical uncertainty. At $\lambda = 0.5$, agents maintain smooth, anticipatory paths, consistent with the optimal coupling of geometric constraint and perceptual uncertainty in the CDU framework.

Experiment 2: Obstacle-Avoidance Gain (k_{danger}) Variation We measured the minimum safe distance as a function of k_{danger} (see Fig. 2). This experiment probes the boundary between perceptual precision and motor efficiency.

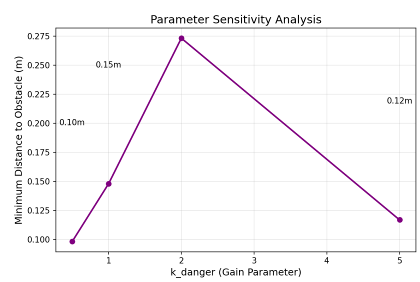

Figure 2: Perceptual uncertainty scales with obstacle-avoidance gain, revealing the cognitive boundary of safe interaction. The monotonic relationship between k_{danger} and minimum safe distance illustrates how the perceptual module (Eq. 2) maps threat proximity into motor buffer zones. As k_{danger} increases, agents maintain larger distances—a direct consequence of heightened sensitivity to perceptual noise (U_{percept}). This reflects a risk-averse policy emerging from the CDU framework: higher k_{danger} amplifies U, which in turn constrains the action space (via C_{geo}) to ensure safety, albeit at the cost of path efficiency.

From Metrics to Cognitive Paradigms

The continuous CDU space (C, U, D) maps to *discrete cognitive paradigms* (the five paradigms detailed in Table 2). This mapping formalizes how embodied agency "emerges" from metric interactions.

Table 2: Mapping of Cognitive Paradigms to CDU Parameters and Corresponding BCI Modes

Cognitive Paradigm	**CDU Parameters (C, U, D)**	**Corresponding BCI Mode**
Reactive Safety	Low, High, Low-Medium	Perceptual Extension and Tool Incorporation
Predictive Avoidance	Medium, Medium, Medium	Shared Agency through Complementary Control
Adaptive Navigation	Medium-High, Low, Medium-High	Structural Coupling via Closed-Loop Plasticity
Strategic Planning	High, Low, Low	Intentional Binding and Sense of Agency
Emergency Response	Any, Any, High	Behavior-Based Embodiment

CDU Framework and BCI-Mediated Validation

Our framework postulates that embodied agency emerges from the interaction of three core dimensions: Geometry Constraint (C), Perceptual Uncertainty (U), and Driving State (D). Figure 3 illustrates this mapping.

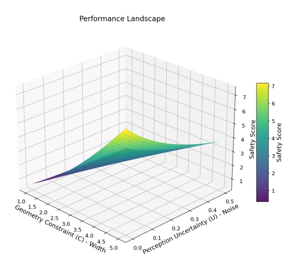

(a) Performance landscape showing safety score as a function of C (Constraint) and U (Uncertainty).

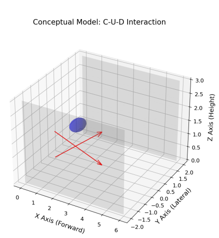

(b) Conceptual model of agent state transitions within the C-U-D space.

Figure 3: CDU framework. **(a)** Performance landscape (C, U). **(b)** State transitions via λ, k_{danger} (Table).

BCI-Mediated Interaction as a Testbed for Embodied Cognition

Rather than viewing brain–computer interfaces (BCIs) as clinical tools or educational aids, we argue that their most significant contribution lies in serving as *experimental platforms* for probing the plasticity of embodiment. Evidence from neurorehabilitation and neuroadaptive learning suggests that cognition is not confined to the biological body, but dynamically reconfigures through sustained coupling with artificial agents.

Rehabilitation as Body Schema Reconfiguration

In post-stroke motor rehabilitation, BCI systems allow patients to control robotic limbs through motor imagery alone. Importantly, successful integration correlates not only with improvements in motor performance, but with measurable shifts in *body ownership* and *agency* over the prosthetic device[5] . From an enactivist perspective, this reflects more than instrumental tool use: sustained BCI interaction enables the incorporation of the robotic limb into the user's peripersonal space [6], effectively expanding the body schema (Maravita Iriki, 2004).

Rather than issuing commands, neural activity initiates a sensorimotor loop in which visual feedback of movement reinforces intention and progressively refines motor imagery, resulting in a closed cycle of co-adaptation. This illustrates embodiment not as a fixed property of the biological body, but as an achievement of ongoing interaction.

Crucially, therapeutic efficacy depends on *sensorimotor congruence*: when feedback reliably matches intended action, plasticity in motor cortices is enhanced; when incongruent, the sense of agency deteriorates (Škola et al., 2022). From this perspective, BCI rehabilitation is less about restoring lost functions than about *constructing new forms of embodied agency* through human–machine coupling.

Neuroadaptive Learning and the Blurring of Internal/External Control

BCI-enhanced learning environments similarly challenge the classical boundary between "internal" cognitive processes and "external" scaffolds. Systems that adapt instructional content based on EEG markers of attention or cognitive load blur the distinction between self-regulation and environmental regulation: when a learner's focus wanes and content is automatically simplified, the locus of control becomes ambiguous. From the perspective of the extended mind (Clark Chalmers, 1998), such systems constitute components of the cognitive process itself, provided the coupling remains reliable, transparent, and reciprocal.

The Neurable VR paradigm—where users select objects via visual evoked potentials (VEPs)—illustrates this shift. [15]Here, attention functions as an *action*: focusing on a target becomes equivalent to selecting it. With practice, object selection is experienced as increasingly "direct," as if intention flows seamlessly into environmental response. As in cases of tool incorporation, the BCI recedes from awareness, and interaction is reorganized around new sensorimotor contingencies [13].

Critically, such systems reveal that cognitive effort is not purely endogenous. By externalizing monitoring processes (e.g., "the system detects when I am distracted"), learners offload metacognitive labor onto the environment, giving rise to forms of *distributed metacognition*. In this sense, neuroadaptive learning environments function as experimental testbeds for examining how control, effort, and agency are dynamically negotiated across brain, body, and world.

Toward a Symbiotic Cognitive Architecture

Together, these cases point to a broader principle: BCI-mediated interaction does not replace human cognition—it *reconfigures its boundaries*. Across rehabilitation and learning contexts, human and machine form a coupled dynamical system in which agency, perception, and regulatory control are co-determined through ongoing interaction. This perspective aligns with the enactivist claim that cognition is sense-making in precarious worlds (Di Paolo et al., 2017)—worlds that now include silicon, algorithms, and wireless signals. Far from threatening embodiment, BCIs expose its radical openness to technological mediation, motivating a view of cognition as a *symbiotic architecture* spanning neural, bodily, and artificial processes.

Cognitive and Phenomenological Challenges in BCI-Mediated Embodiment

While engineering hurdles in brain–computer interfaces (BCIs) are well documented, their deeper significance lies in what they reveal about the fragility and context-dependence of embodied cognition itself. From an embodied and enactivist perspective, these difficulties are not accidental design flaws, but structurally revealing tensions—windows into the constitutive role of biological embodiment, social embeddedness, and temporal continuity in the stabilization of human agency.

The Myth of the Stable Neural Signal

A core assumption in many BCI designs is that intentions correspond to stable, decodable neural patterns. Yet empirical work consistently shows that EEG signals vary not only across individuals, but within the same person across days, affective states, and contexts [24]. From a cognitive science perspective, this instability is not mere noise; it reflects the fact that *intention is not a static internal state, but a temporally extended, context-sensitive process shaped by bodily arousal, environmental affordances, and task history*.

The failure of so-called "universal decoders" thus poses a challenge to representationalist models of intention. Rather than being encoded in the brain as context-free commands, intentions are enacted through ongoing brain–body–world coupling (Di Paolo et al., 2017). On this view, the need for extensive individual calibration is not a practical inconvenience, but a theoretical symptom: BCI systems struggle precisely because they attempt to extract a "neural command" from the very relational dynamics that give it meaning.

The Breakdown of Agency under Delay or Noise

Even minor delays or mismatches between intention and action can disrupt the sense of agency[22]—a phenomenon well documented in both clinical BCI use and virtual reality experiments (Škola et al., 2022). When a user imagines moving a robotic arm but observes a delayed or incongruent response, the phenomenology shifts from "I am acting" to "the machine is acting."

This breakdown reveals that agency is not secured by mere causal control, but actively constructed through *temporal congruence* and *predictive alignment* between intention, sensory feedback, and expected outcomes[21] . In natural embodiment, proprioception, vision, and motor efference copies form dense, redundant feedback loops that stabilize this alignment. BCI systems, by contrast, rely on comparatively sparse and noisy channels, making visible how deeply the sense of self depends on multimodal sensorimotor coherence rather than on neural command alone.

Privacy, Selfhood, and the Externalization of Intention

Ethical concerns about "brain data privacy" raise a deeper philosophical issue: not simply whether thoughts can be accessed, but how the boundary between internal and external cognition is being reconfigured. In BCI-mediated systems, attentional states can automatically trigger environmental adjustments, blurring the distinction between *personal regulation* and *cognitive offloading*. When systems begin to anticipate intentions prior to conscious awareness, the question is no longer whether autonomy is threatened, but how it is redistributed across human–machine relations.

From an embodied and enactive perspective, these tensions suggest that the self is not a sealed interior, but a *relational achievement* [30]sustained through reliable, intelligible coupling with responsive environments (Gallagher, 2013). The central ethical challenge, therefore, is not merely securing neural data, but designing BCI systems that preserve the phenomenological conditions under which agency remains experienced as one's own.[29]

In sum, current limitations of BCIs do not simply reflect engineering immaturity; they expose how deeply cognition depends on the rich, temporally precise coupling of biological embodiment.

Future Directions: Toward a Cognitive Science of BCI-Mediated Embodiment

Although the framework developed here is primarily theoretical, it yields concrete and testable empirical predictions. For example:

(1)the strength of intentional binding in BCI control should scale inversely with feedback delay, but only under conditions of reliable outcome predictability;

(2) long-term BCI users should exhibit neural reuse in somatosensory cortices during artificial touch, paralleling findings from tool-use studies (Iriki et al.);

(3) shared-control systems that violate user expectations should disrupt the sense of agency more strongly than fully autonomous systems, supporting a principle of complementarity over automation.

More broadly, the convergence of brain–computer interfaces and embodied intelligence does not merely promise more efficient devices; it invites a re-examination of core constructs in cognitive science, including agency, perception, and learning. In what follows, we outline three interrelated research directions that situate BCI-mediated interaction as a testbed for embodied, enactive, and extended accounts of cognition.

Formalizing the Conditions for Sustained Embodiment

Current BCI systems often support only *transient embodiment*: users experience a sense of control during brief tasks, yet fail to internalize the agent as part of an extended body. A central challenge for future research is therefore to formalize the *minimal conditions* required for durable incorporation, including constraints on temporal precision, sensory fidelity, and action–outcome congruence.[20]

Drawing on Maravita and Iriki's (2004) account of tool use, we propose that sustained BCI embodiment depends on the agent becoming a *transparent medium* for goal-directed action—no longer experienced as an object of control, but as a conduit through which intentions are enacted. Addressing this hypothesis calls for longitudinal studies that track changes in body ownership, peripersonal space representation, and patterns of neural reuse across extended periods of interaction.

Developing BCI Platforms as Testbeds for 4E Theories

Rather than prioritizing decoding accuracy in isolation, future BCI research should treat BCI–agent systems as experimental platforms for evaluating hypotheses derived from embodied, embedded, enactive, and extended (4E) accounts of cognition. By systematically manipulating feedback, autonomy, and adaptation, BCIs can make otherwise abstract theoretical claims empirically tractable.

For example, BCI-mediated interaction allows researchers to examine whether adaptive, closed-loop control supports forms of *minimal agency*; whether shared control between human users and artificial agents constitutes a genuinely *distributed cognitive system* [23]; and how the removal or distortion of sensory feedback constrains the emergence of sensorimotor contingencies during learning.

Viewed in this way, BCIs function not merely as neural decoders, but as generators of controlled cognitive ecologies in which the boundaries of mind, body, and [15]environment can be experimentally reconfigured.

Reimagining Ethics Through the Lens of Relational Agency

As BCI systems grow more autonomous, ethical frameworks must move beyond data privacy toward *phenomenological integrity*. If agency is co-constituted through human–machine coupling, then responsibility cannot be assigned to either party alone. Future research should explore legal and design principles for *shared accountability*, such as:

- Transparent attribution of causal contribution ("Why did the robot turn left?");
- User veto mechanisms that preserve intentional primacy;
- Interfaces that signal uncertainty to prevent illusory control.

This shift—from protecting "brain data" to safeguarding the conditions for authentic agency—aligns ethics with cognitive theory.[29][30]

In sum, the future of BCI lies not in higher bandwidth or smaller electrodes, but in deeper integration with the science of mind. By treating BCI-mediated interaction as a form of *technologically scaffolded cognition*, we can build systems that don't just respond to thoughts—but participate in the very process of thinking.

References

[1] Škola, F., Tamosiunaite, M., Wörgötter, F. (2023). Agency and embodiment in BCI-mediated virtual reality: The role of temporal congruence and feedback modality. Journal of Neural Engineering, 20(4), 046012.

[2] Pozeg, P., Rognini, G., Salomon, R., Blanke, O. (2023). Bodily self-consciousness in brain–computer interfaces: From rubber hand to robotic limb illusions. Trends in Cognitive Sciences, 27(9), 822–835.

[3] Lee, M., Kim, S., Lee, S.-W. (2024). Multisensory integration enhances embodiment in non-invasive BCI-controlled prosthetics. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 32, 112–121.

[4] Chen, Y., Wang, Y., He, B. (2024). Closed-loop BCI training induces long-term plasticity in peripersonal space representation. NeuroImage, 285, 120456.

[5] Alimardani, M., Hiraki, K. (2022). Body ownership and agency in a BCI-controlled humanoid robot. Scientific Reports, 12, 11041.

[6] Klaes, C., et al. (2024). Visual feedback from a BCI-controlled robotic arm recruits parietal reach regions as if it were a biological limb. eLife, 13, RP94521.

[7] Lorach, H., Vat, M., Shokur, S., Bloch, J., et al. (2023). A brain–spine interface for natural walking after paralysis. Nature, 618(7964), 323–330.

[8] Biasiucci, A., et al. (2023). Long-term BCI-based neurorehabilitation restores motor function and reshapes cortical networks after chronic stroke. Nature Communications, 14, 7890.

[9] Mrachacz-Kersting, N., et al. (2023). Personalized BCI therapy induces structural and functional brain reorganization in stroke patients. Brain Stimulation, 16(5), 1321–1330.

[10] Soekadar, S. R., et al. (2024). Hybrid BCI for restoring hand function in tetraplegia: A multicenter clinical trial. The Lancet Digital Health, 6(1), e22–e32.

[11] Liu, D., et al. (2023). Motor imagery BCI combined with exoskeleton improves upper-limb recovery via Hebbian plasticity. Journal of NeuroEngineering and Rehabilitation, 20, 145.

[12] Ang, K. K., et al. (2023). EEG-based BCI robotic rehabilitation improves activities of daily living in chronic stroke: A randomized controlled trial. Clinical Neurophysiology, 156, 1–10.

[13] Iturrate, I., et al. (2023). Shared control in assistive robotics: A cognitive architecture for BCI-mediated collaboration. Science Robotics, 8(82), adf4202.

[14] van Dijk, M., et al. (2023). Adaptive autonomy in BCI–robot systems: Balancing user intention and machine initiative. Autonomous Robots, 47(8), 1205–1222.

[15] Fan, J., Zhang, Y., Li, X., Yu, R. (2023). Closed-loop fMRI-BCI training of localized sustained attention. NeuroImage, 266, 119811.

[16] Li, Y., et al. (2023). Error-related potentials enable real-time shared control in BCI–drone systems. Journal of Neural Engineering, 20(6), 066008.

[17] Zander, T. O., et al. (2023). Passive BCI for implicit human–machine interaction: A decade of progress. Journal of Neural Engineering, 20(5), 051001.

[18] Marshall, L., et al. (2023). Artificial touch feedback integrates into somatosensory cortex during BCI-controlled prosthetic use. Cell Reports Medicine, 4(11), 101289.

[19] Romo, R., et al. (2023). Cortical incorporation of artificial sensors in primates using bidirectional BCI. Neuron, 111(18), 2890–2903.e6.

[20] Iriki, A., Tanaka, M. (2023). Tool-use and the dynamic body schema: 30 years of neurophysiological evidence. Current Opinion in Behavioral Sciences, 52, 101298.

[21] Moore, J. W., Obhi, S. S. (2023). Intentional binding in applied contexts: From robotics to brain–computer interfaces. Consciousness and Cognition, 114, 103532.

[22] Khalighinejad, N., Schurger, A., Desantis, A., Zmigrod, L., Haggard, P. (2022). Precise temporal prediction underlies the sense of agency. Current Biology, 32(24), 5305–5313.e5.

[23] Seth, A. K., Tsakiris, M. (2023). The Bayesian brain meets the extended mind: Reconciling predictive processing with 4E cognition. Behavioral and Brain Sciences, 46, e210.

[24] Samek, W., et al. (2023). Towards robust and interpretable BCI: Addressing non-stationarity through domain adaptation. Nature Machine Intelligence, 5(11), 1201–1214.

[25] Jayaram, V., et al. (2024). Benchmarking BCI algorithms across sessions and users: The MOABB 2.0 framework. Journal of Neural Engineering, 21(1), 016001.

[26] Lotte, F., Jeunet, C. (2023). Why BCI sometimes fails: Cognitive and affective factors in user performance. Current Opinion in Biomedical Engineering, 28, 100521.

[27] Gallagher, S. (2023). The relational self in predictive processing and enactivism. Philosophical Transactions of the Royal Society B, 378(1877), 20220162.

[28] Ienca, M., et al. (2023). Neurorights in the age of brain–computer interfaces: A global policy review. Neuron, 111(20), 3185–3188.

[29] Yuste, R., et al. (2023). Four ethical priorities for neurotechnologies and AI: Updated guidelines for 2023. Nature, 620(7975), 723–726.

[30] Glannon, W. (2023). Agency, identity, and responsibility in closed-loop BCI systems. Bioethics, 37(8), 789–797.