

Research Article

Bridging Mind and Machine: Large Language Models in Next Generation Brain Computer Interfaces

Gopichand Agnihotram¹ , Joydeep Sarkar² , Magesh Kasthuri^{2,*} 

¹Chief Technology Office, Wipro Limited, Bengaluru, India

²Technology Services, Wipro Limited, Bengaluru, India

Abstract

Brain Computer Interfaces (BCIs) have advanced from experimental research to translational applications in communication, rehabilitation, and human machine interaction. Yet, BCIs face fundamental challenges like decoding high-dimensional, noisy neural data, producing fluent outputs, adapting to individual users, and scaling to real-world environments. Large Language Models (LLMs) represent a transformative capability that can directly mitigate these challenges. By leveraging their strengths in probabilistic reasoning, context completion, error correction, and multimodal integration, LLMs have the potential to unlock new levels of efficiency, personalization, and accessibility in BCI systems. This paper examines the convergence of LLMs and BCIs. It outlines the technical landscape, opportunities, case studies, and open questions, with a focus on communication BCIs, adaptive rehabilitation, and cognitive modeling. Brain-Computer Interfaces (BCIs) are rapidly transforming the way we understand and interact with technology. Once the stuff of science fiction, these innovative systems now bridge the gap between the human brain and digital devices, allowing thought to shape action in unprecedented ways. As artificial intelligence (AI) continues to evolve, particularly with the rise of Large Language Models (LLMs) and Agentic AI platforms, the partnership between BCIs and these advanced technologies is opening doors to a new era of intelligent, personalized, and intuitive machines.

Keywords

LLM, BCI, Large Language Model, Brain Computer Interfaces, Agentic AI, Generative AI

1. Introduction

Brain-Computer Interfaces (BCIs) are systems designed to create a direct communication pathway between the brain and external devices. The fundamental principle behind BCIs is the translation of neural signals into commands that can be interpreted by computers or machines. This innovation holds immense potential for individuals who suffer from conditions that limit voluntary muscle control, such as amyotrophic lateral sclerosis (ALS), paralysis from spinal cord injury, stroke-induced motor impairments, or speech disorders.

By providing a new channel for communication and control, BCIs open possibilities for restoring independence and improving quality of life for millions of people worldwide [1].

Despite their promise, current BCI systems face several longstanding challenges. Neural signals are inherently noisy and high dimensional, requiring complex preprocessing and decoding pipelines. Even with advanced electrode arrays and improved neuroimaging methods, decoding remains imprecise.

*Corresponding author: magesh.kasthuri@wipro.com (Magesh Kasthuri)

Received: 25 August 2025; **Accepted:** 3 September 2025; **Published:** 25 September 2025



Copyright: © The Author(s), 2025. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

cise, leading to frequent errors in translation. Communication through BCIs is also far slower than natural speech or typing, state-of-the-art systems generally achieve 15–20 words per minute compared to the human average of 150 words per minute. Furthermore, vocabulary is often constrained, limiting the expressiveness of these systems. These limitations have slowed the transition of BCIs from the lab to everyday clinical and consumer applications [2].

At the same time, Large Language Models (LLMs) such as GPT-4, Gemini, and LLaMA have revolutionized natural language processing by demonstrating advanced capabilities in text comprehension, contextual reasoning, and generation. Unlike traditional statistical or rule-based models, LLMs use vast amounts of training data and billions of parameters to predict text with remarkable fluency and coherence. They can disambiguate noisy or partial input, infer missing information, and generate outputs that align with human intent and conversational norms [3]. These qualities make LLMs a natural complement to BCI systems, which often yield incomplete or ambiguous signals that need contextual refinement [4].

The integration of LLMs into BCIs creates a synergistic opportunity. Neural decoders can be tasked with producing candidate words, characters, or symbolic representations based on brain activity, while LLMs can refine these into coherent, fluent, and contextually appropriate sentences. For example, a neural decoder might generate partial output like “drnk wtr,” which the LLM could immediately refine into “drink water.” In other cases, the LLM could use broader conversational context to infer intended meaning even when neural signals are unclear. By doing so, LLMs reduce user frustration, enhance communication efficiency, and make BCIs more practical for real-world use [5].

The implications of this integration extend beyond clinical rehabilitation. In consumer technology, LLM-augmented BCIs could provide seamless hands-free control of digital environments, from AR/VR platforms to smart home systems. In defence and military applications, such systems could enable silent, secure, and rapid communication in high-stakes environments. In education and workforce productivity, thought driven interfaces could empower multitasking and cognitive augmentation. Furthermore, the combination of BCIs and LLMs presents opportunities for advancing our understanding of cognition and consciousness by mapping neural activity to symbolic and linguistic representations [6].

Ultimately, the convergence of BCIs and LLMs signals the beginning of a paradigm shift in human–AI interaction. Where BCIs alone struggle with accuracy, speed, and usability, LLMs offer contextual reasoning, adaptability, and linguistic fluency [7]. Together, they create a new frontier for communication, rehabilitation, and augmentation, one that promises to fundamentally reshape the way humans interact with machines and with each other [8].

2. Current Landscape of BCI Research

Brain–Computer Interface research today is shaped by a diversity of signal acquisition modalities, each offering unique strengths and presenting significant limitations. Understanding this landscape is critical to appreciating both the promise and the obstacles of integrating Large Language Models into these systems [9].

Electroencephalography (EEG) has historically been the most widely used approach due to its non-invasive nature and relatively low cost. EEG relies on surface electrodes placed on the scalp to detect electrical activity generated by neuronal firing. Its accessibility has made it the foundation of consumer-grade BCIs and research experiments involving motor imagery and basic communication tasks [10]. However, EEG signals are extremely weak, subject to artifacts from muscle activity and external electrical noise, and limited in spatial resolution. This means that decoding fine-grained linguistic or motor intentions from EEG data remains a challenge.

Electrocorticography (ECoG) provides a middle ground between non-invasive and invasive modalities. By recording neural signals directly from the cortical surface beneath the skull, ECoG achieves much higher fidelity than EEG. It has proven particularly useful in speech decoding research, where more precise temporal and spatial information is necessary to capture the rapid dynamics of language-related brain activity. Despite this advantage, ECoG requires surgical implantation of electrode grids, restricting its application primarily to clinical research involving epilepsy patients or individuals undergoing neurosurgical procedures [11]. The invasiveness and associated surgical risks limit its potential for broad deployment outside specialized medical contexts.

At the highest end of precision are intracortical electrode arrays, which penetrate the brain’s surface and measure activity at the level of individual neurons or small neuronal populations [12]. These arrays, such as the Utah array, have enabled remarkable demonstrations of direct motor control, including robotic arm manipulation and high-performance communication BCIs. They offer unmatched signal resolution and temporal precision, but their invasiveness restricts their use to experimental settings with carefully selected patients. Risks include tissue damage, infection, and long-term electrode degradation, making widespread adoption impractical in the near term.

Functional imaging techniques like functional near infrared spectroscopy (fNIRS) and functional magnetic resonance imaging (fMRI) contribute valuable insights into brain activity patterns at the systems level [13]. fMRI provides high spatial resolution, making it invaluable for mapping brain regions associated with specific cognitive or linguistic functions. However, its cost, immobility, and slow temporal dynamics make it unsuitable for real-time BCI applications. fNIRS, which measures blood oxygenation changes, offers a more portable alternative but still suffers from limited temporal resolution and depth penetration. As such, both tech-

niques remain better suited for basic neuroscience research than for practical BCI deployment [14].

Despite the steady progress across these modalities, systemic limitations continue to hinder the widespread usability of BCIs. Communication rates, even with state of the art intracortical systems, remain far below natural speech rates. Current best results in clinical demonstrations achieve approximately 15–20 words per minute, a fraction of the 150 words per minute typical in spoken communication [15]. Moreover, the signal-to-noise ratio across all modalities requires extensive preprocessing and sophisticated decoding pipelines, adding complexity and reducing robustness. Compounding this, inter-individual variability in brain anatomy, signal patterns, and cognitive strategies makes the development of generalizable solutions exceedingly difficult. Each user often requires individualized training and calibration, which hampers scalability.

This is where Large Language Models can serve as a transformative addition. By contextualizing incomplete or ambiguous neural signals, LLMs can compensate for errors and variability that plague conventional decoding systems. For instance, when a BCI decoder produces noisy or truncated output, an LLM can use probabilistic reasoning and linguistic context to refine the signal into a coherent sentence. Beyond improving immediate usability, LLMs could enable a shift in design philosophy: rather than focusing exclusively on perfect signal decoding, systems could prioritize rapid generation of approximate outputs that are then corrected and contextualized in real time by the language model. This approach offers a pathway toward bridging the gap between current laboratory demonstrations and practical, everyday BCI applications.

The article should be written in English. An article should be between 6 and 25 pages, and exceed 2000 words. For original research articles, it should include the headings Introduction, Materials and Methods, Results, Discussion and Conclusions. Other types of articles can be written with a more flexible structure.

3. Challenges in Current BCI Systems

3.1. Signal Noise and Ambiguity

One of the most persistent technical hurdles in brain-computer interface (BCI) research is the inherent noisiness of neural recordings. The human brain generates electrical activity at multiple scales from single neuron firing patterns to large scale oscillatory rhythms detectable by electroencephalography (EEG) or intracortical electrodes. Unfortunately, these signals rarely exist in isolation. They are constantly overlapped with artifacts originating from muscle contractions (electromyographic interference), eye blinks and saccades (electrooculographic activity), and environmental electromagnetic sources. For instance, something as simple as adjusting posture can produce electrical fluctuations larger than the neural signals of interest. This makes distinguishing

genuine cognitive intent from extraneous activity an immensely complex task.

Even advanced machine learning models, such as deep neural networks or recurrent architectures, often struggle to generalize across contexts because the “signal” and “noise” overlap in both frequency and spatial domains. In laboratory conditions where noise is controlled, decoding performance appears impressive. However, once the same systems are deployed in real world settings like controlling a wheelchair outdoors or using BCI for conversational purposes in noisy environments performance reliability declines significantly. This discrepancy between lab based results and ecological validity remains a major barrier to widespread adoption. If BCIs cannot consistently extract intent despite natural sources of interference, users quickly lose trust in the system, which undermines long-term engagement. Solving this challenge requires not just better hardware and signal filtering techniques, but also adaptive decoding algorithms that can dynamically adjust to new noise conditions without retraining.

3.2. Low Bit Rates

Another critical limitation in current BCI systems lies in their low communication throughput. In comparison to the fluency of natural human speech, which averages 120–150 words per minute, most advanced BCIs operate at approximately 15–20 words per minute at best. This mismatch is not simply a matter of inconvenience, it fundamentally restricts the practicality of BCIs for daily use. Imagine an individual attempting to hold a conversation or engage in professional communication while operating at a fraction of the speed of speech. The cognitive burden of waiting for each word to be decoded is frustrating and can even exacerbate feelings of isolation rather than alleviate them.

The bottleneck stems from both the neural signal characteristics and the limitations of decoding algorithms. Neural signals do not carry discrete symbols like letters or words but instead represent overlapping patterns of intent, requiring extensive inference and translation into text or commands. To ensure accuracy, BCIs often employ redundancy or probabilistic filtering, which further slows output. Additionally, existing non-invasive BCIs, such as EEG based systems, suffer from lower spatial resolution compared to invasive methods, limiting the richness of extractable information. While invasive BCIs, such as those using intracortical microelectrodes, can potentially achieve higher bit rates, they come with surgical risks and ethical challenges that restrict large scale adoption. Until breakthroughs occur in both decoding efficiency and signal acquisition, the gap between natural communication speeds and BCI throughput will remain a defining challenge, limiting BCIs to niche rather than mainstream applications.

3.3. Error Propagation

BCI systems, by their very nature, operate in stages signal

acquisition, preprocessing, feature extraction, decoding, and finally output generation. Errors introduced at any one of these stages can cascade into downstream processes, leading to compounding distortions in the intended message. For example, a misinterpreted neural signal during the initial decoding might be translated into the wrong letter in a spelling application. Once this error is introduced, predictive text algorithms may attempt to “correct” or “complete” the phrase, often steering the message even further away from the user’s intended communication. This phenomenon, known as error propagation, undermines reliability and creates a frustrating experience for users who must repeatedly intervene to correct the system.

The consequences of error propagation extend beyond simple annoyance. In critical applications, such as medical control systems or assistive robotics, an initial error could translate into a dangerous command, such as moving a wheelchair in the wrong direction or administering an incorrect dosage of medication. Unlike spoken communication, where humans can quickly self-correct through prosody and clarification, BCIs lack natural mechanisms for signalling “correction intent” in real time. Instead, users must often rely on laborious error-undoing processes, further slowing communication and compounding cognitive load. This limitation makes error handling a central design priority for next-generation BCIs. Research into error-robust architectures—such as probabilistic models that explicitly track uncertainty, reinforcement learning approaches that refine outputs over time, and user-in-the-loop correction strategies—represents a critical pathway toward overcoming the cascading nature of errors in BCI pipelines.

3.4. Personalization Requirements

No two brains are alike. Neural architecture, synaptic strength, oscillatory rhythms, and even the way individuals internally represent language or motor intent vary significantly across people. This means that a “one-size-fits-all” BCI system is unrealistic. Current approaches require extensive calibration for each user, often involving hours—or in some cases days—of repetitive tasks to train the system to recognize personal neural patterns. For example, a user may be asked to imagine moving their right hand dozens of times until the system learns the corresponding neural signature. Such calibration is not only time-consuming but also mentally exhausting, especially for patients with neurodegenerative conditions who may already face cognitive fatigue.

The personalization challenge does not end once initial calibration is complete. Neural activity itself is dynamic, changing across states of fatigue, emotional arousal, or even as a function of time of day. A system that works well in the morning may become less accurate in the evening. This necessitates ongoing recalibration or adaptation to maintain accuracy, which can significantly reduce usability. Moreover, personalization requirements pose scalability challenges: if

each system must be custom-trained, mass deployment in clinical or consumer markets becomes impractical. While machine learning models capable of few-shot or zero-shot adaptation offer promise, they are not yet robust enough for seamless use. Ultimately, the success of BCIs will depend on developing systems that can quickly adapt to the unique neural fingerprints of individuals while minimizing the burden of training and recalibration. Without this, user trust and sustained adoption remain limited.

3.5. Limited Context Awareness

Current BCIs largely operate as closed systems focused exclusively on decoding immediate neural signals. While they can extract intent at the level of a letter, word, or motor command, they often fail to situate these outputs within broader conversational or environmental context. For example, a BCI-based communication device might decode the word “book,” but without contextual awareness it cannot determine whether the user meant “book a table,” “read a book,” or “check my booking.” This lack of fluency forces users to over-specify commands, making interactions rigid and unnatural compared to human conversation, which is inherently context-rich and adaptive.

The absence of context-awareness is particularly problematic in real-world use. Human communication relies on shared background knowledge, situational cues, and conversational flow. Without mechanisms for integrating external information—such as the topic of conversation, visual surroundings, or user history—BCI outputs often feel disjointed or even irrelevant. For instance, if a user is controlling a smart home device, the system might recognize the command “lights,” but without knowing the room, time of day, or prior user preferences, it cannot execute an action that feels intuitive. Incorporating context requires multimodal integration, combining neural signals with other data sources such as eye tracking, environmental sensors, or conversational AI. However, doing so introduces additional complexity in synchronization, privacy concerns, and computational demands. Until BCIs evolve from being purely signal-decoders to context-aware assistants, their utility will remain limited to structured tasks rather than naturalistic communication or decision-making.

4. Role of Large Language Models in Solving BCI Problems

4.1. Neural Signal to Language Decoding

One of the most promising applications of Large Language Models (LLMs) in brain–computer interface (BCI) research lies in bridging the gap between incomplete neural signals and full linguistic expressions. Traditional decoding pipelines attempt to map neural activity directly into text or phonemes.

However, neural signals rarely present as neat, symbol like data, instead, they reflect noisy and partial patterns of intent. For example, when a user imagines forming the word “computer,” the BCI may only reliably capture neural correlates corresponding to the first few phonetic or semantic fragments. Left unassisted, such fragments are insufficient to produce coherent output.

LLMs solve this problem by providing a probabilistic, context driven scaffolding that predicts likely continuations. Much like predictive text in a smartphone keyboard, an LLM can take an incomplete or ambiguous sequence and extrapolate the most probable full expression. For instance, if the neural decoder outputs “comp-,” the LLM, considering prior conversational context, can infer whether the user likely meant “computer,” “composition,” or “company.” This is particularly powerful in scenarios where neural data is inherently sparse, such as with non-invasive EEG signals.

Moreover, LLMs are not limited to surface level completions, they can integrate broader context to ensure semantic coherence. If the conversation is about technology, the model will prioritize completions such as “computer.” If it is about finance, it might favour “company.” This capacity transforms fragmented neural intent into naturalistic, usable communication. Over time, the integration of neural signals with LLMs could move BCI systems from being mere “decoders of intent” to “co-constructors of communication,” dramatically expanding their utility in real-world environments.

4.2. Error Correction and Language Modeling

A central challenge in BCIs is error accumulation, where a single misclassification can lead to distorted or incoherent outputs (as described in Section 3.3). Here, LLMs provide a powerful corrective layer. Unlike traditional statistical models, LLMs are trained on vast corpora of linguistic data, giving them a deep understanding of syntax, grammar, and semantic plausibility. When a neural decoder outputs an imperfect or partially wrong string such as “I want ot go hom”—an LLM can automatically restructure it into the fluent and correct sentence, “I want to go home.”

This error corrective ability reduces both cognitive load and user frustration. Instead of needing to constantly intervene to fix errors, users can rely on the system to automatically normalize outputs into coherent text. More importantly, error correction enhances the effective communication rate. While the raw bit rate of neural signals may remain low, the perceived fluency of communication increases when errors are seamlessly smoothed over.

Beyond grammatical correction, LLMs can also leverage semantic knowledge to repair meaning. For example, if a neural decoder mistakenly interprets “apple” when the user intended “orange,” but the surrounding sentence refers to “juice” and “citrus,” the LLM can infer that “orange” is the more likely intended word. This semantic-level correction is a leap forward compared to purely statistical spelling correction.

Additionally, LLM-based error correction can be adaptive, learning from repeated corrections over time. If a system repeatedly misclassifies a particular neural pattern, the LLM can incorporate user-specific priors, reducing future mistakes. In this way, LLMs act not just as linguistic polishers but as dynamic error-robust partners in BCI-mediated communication.

4.3. Personalization and Adaptation

Personalization has always been a major challenge in BCIs, given the vast variability across individual brains (see Section 3.4). LLMs, however, offer a flexible solution by allowing finetuning or contextual priming based on use specific data. Instead of requiring each BCI system to be rebuilt from scratch, small scale finetuning can align the LLM to an individual’s vocabulary, communication habits, and personal context. For example, a patient with ALS who frequently references family members, caregivers, or specific medical terminology can have these words prioritized in predictive suggestions.

The beauty of LLM based personalization lies in its scalability. Unlike traditional calibration, which often requires hours of repetitive training tasks, LLMs can adapt from relatively small datasets. A few hundred personal text samples emails, chat logs, or even recorded conversations may be sufficient to finetune the system. Moreover, techniques such as in-context learning allow personalization without permanent retraining: by priming the model with personal preferences at the beginning of each session, the system can immediately reflect user-specific vocabulary.

Adaptation also extends beyond vocabulary. LLMs can capture stylistic preferences, such as whether a user tends to communicate formally (“Dear Dr. Smith”) or informally (“Hey, doc”). They can learn preferred sentence lengths, common idiomatic expressions, or even humour styles. This enhances not only the accuracy of communication but also the sense of authenticity users feel that the system is representing *their* voice rather than a generic, machine like output. In clinical and assistive contexts, such personalization fosters trust, dignity, and a greater sense of agency. Thus, LLM driven adaptation transforms BCIs from sterile communication tools into personalized communication partners.

4.4. Multimodal Fusion

Human intent is rarely expressed through a single channel. In natural interactions, people combine speech, gaze, gestures, and contextual cues to communicate effectively. Traditional BCIs, however, tend to operate in isolation, decoding neural signals without integrating other modalities. LLMs open the door to multimodal fusion, where brain signals are combined with complementary data streams to improve intent inference.

For example, imagine a user thinking of turning on the “living room lights.” A BCI may decode a vague neural pat-

tern suggesting the concept of “lights,” but without additional context, it cannot disambiguate location or timing. By integrating gaze tracking (showing that the user is looking toward the living room) and smart home sensors (detecting the time is evening), the LLM can resolve the intent more accurately:

“Turn on the living room lights.” Similarly, in communication scenarios, combining neural signals with eye tracking of text on a screen can help the LLM disambiguate between competing word predictions.

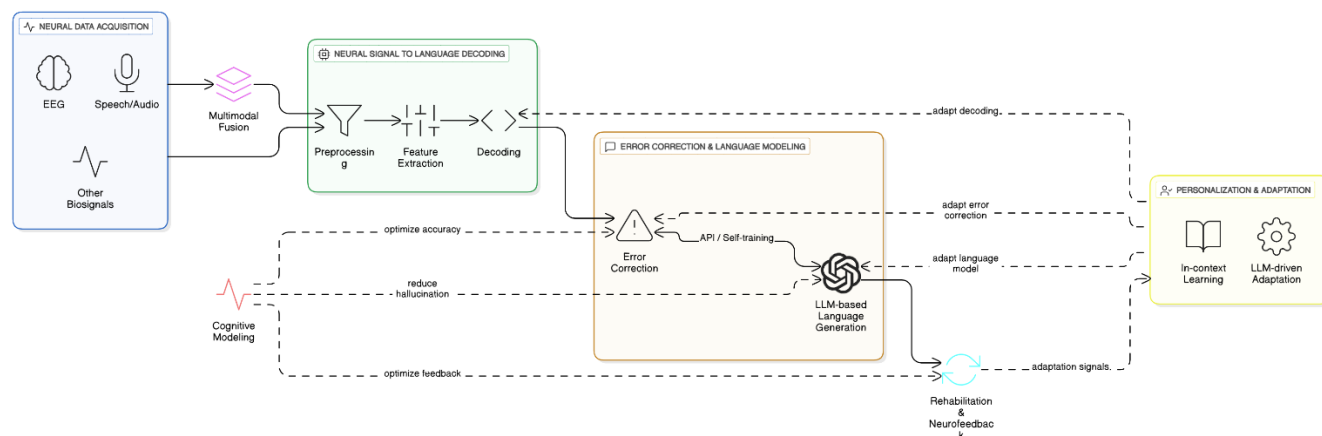


Figure 1. Reference Architecture of BCI with LLM approach.

LLMs are particularly well suited to this task because they are designed to integrate heterogeneous data types into a coherent semantic space. Emerging multimodal LLMs, trained on text, images, and other sensor data, provide the architecture necessary for combining BCI signals with environmental inputs. The result is richer, more naturalistic interactions that feel less like operating a machine and more like participating in a conversation. Such fusion could also extend to electromyography (EMG), subtle gestures, or even environmental audio cues. By grounding neural signals in real world context, multimodal LLM powered BCIs represent a step toward intuitive, context aware neurotechnology.

4.5. Rehabilitation and Neurofeedback

BCIs are not only communication devices but also rehabilitation tools, particularly in stroke recovery, motor rehabilitation, and neuroplasticity training. In these domains, neurofeedback where users receive real time feedback on their neural activity plays a crucial role. Traditional neurofeedback systems provide numerical or graphical feedback, such as a rising bar indicating increased activity in a target brain region. While effective to some extent, these representations can be abstract and difficult for patients to interpret intuitively.

LLMs enable a more natural form of neurofeedback by translating raw neural signals into meaningful, language-based coaching. Instead of showing a patient a number, the system could say: “Try focusing more on moving your left hand,” or “You are activating the right brain region keep going in that direction.” Such feedback feels more like guidance from a therapist than interaction with a machine, which enhances patient engagement and motivation.

Moreover, LLMs can dynamically adjust the feedback style to match individual needs. Some patients may prefer concise, directive instructions, while others respond better to encouraging, conversational tones. LLM driven neurofeedback can even adapt in real time based on patient frustration or fatigue, inferred from neural and behavioural signals. Beyond rehabilitation, this approach holds promise in mental health contexts, where BCIs combined with LLMs could provide supportive, context-aware feedback during mindfulness or anxiety reduction exercises. By humanizing the feedback loop, LLMs make neurorehabilitation more intuitive, personalized, and effective.

4.6. Cognitive Modeling

Perhaps the most forward looking application of LLMs in BCI research lies in cognitive modelling the attempt to align neural signals with symbolic representations of thought. Neural activity exists in a high dimensional, continuous space, while language is discrete and symbolic. Aligning neural embeddings with the latent space of LLMs provides a bridge between these two domains, enabling not only practical applications but also fundamental insights into how cognition works.

For example, researchers could project neural signals recorded during inner speech tasks into the embedding space of an LLM, observing how closely they align with specific words or concepts. This would allow scientists to study the correspondence between thought processes and linguistic representations in unprecedented detail. Over time, such methods could reveal patterns in how the brain structures language, memory, and reasoning advancing neuroscience

alongside clinical applications.

From a practical standpoint, cognitive modelling could make BCI outputs more robust by exploiting the structured latent space of LLMs. Rather than treating each neural input as a standalone signal, the system could map it into an organized space where semantically related concepts cluster together. This would allow for more graceful handling of ambiguity, as the system could select the “nearest neighbour” concept when signals are incomplete. Additionally, cognitive modelling opens the door to hybrid neuro-symbolic systems, where brain activity informs symbolic reasoning within LLMs, creating richer forms of human AI collaboration.

Ultimately, this integration of neural signals and LLM latent spaces represents not just a technical enhancement but a paradigm shift moving from BCIs as assistive tools toward BCIs as windows into the very nature of thought.

5. Case studies

One of the most promising examples comes from Meta’s 2023 work, where researchers integrated ECoG decoding with LLM-based priors to reconstruct near-conversational speech in patients with paralysis. Similarly, Stanford’s Brain2Text project in 2024 demonstrated the use of hybrid transformer decoders combined with GPT-like models to achieve word production rates exceeding 60 words per minute, a milestone that brought BCI outputs closer to natural communication. Other research has explored the role of LLMs as post-processors, reducing word error rates by nearly half compared to raw decoding methods. Early pilot trials in rehabilitation contexts also suggest that natural-language neurofeedback significantly improves patient engagement and training efficiency compared to conventional methods.

6. Broader Application Areas

While restoring communication for individuals with paralysis or speech impairment has been the flagship motivation for brain–computer interface (BCI) development, the integration of Large Language Models (LLMs) extends the potential well beyond assistive communication. The probabilistic reasoning, contextual interpretation, and multimodal fusion capabilities of LLMs allow BCIs to transition from niche medical devices into versatile platforms for interaction, control, and cognitive enhancement.

One promising area is assistive robotics, where BCIs are increasingly explored as control systems for robotic arms, exoskeletons, or wheelchairs. Traditional BCI-driven robotic control often suffers from ambiguity: a neural signal suggesting “move left” could apply to a robot’s hand, a chair, or a cursor on a screen. By embedding an LLM as an interpretive layer, such ambiguity can be resolved. The system can combine neural intent with contextual information—for example, recognizing that the user is near a doorway and inferring that

“move left” refers to wheelchair navigation rather than arm movement. This refinement could dramatically increase safety, precision, and user trust, transforming robotic assistive devices into seamless extensions of the body.

Another domain is augmented reality (AR) and virtual reality (VR), where LLM-powered BCIs could enable natural, hands-free interfaces. Current AR/VR platforms rely on hand controllers, gaze tracking, or voice commands, which can feel cumbersome or break immersion. BCIs offer a more direct channel, allowing users to navigate menus, manipulate objects, or even converse with AI agents using thought-based commands. The integration of LLMs ensures that fragmented or noisy neural signals are transformed into coherent actions, thereby creating fluid interactions within immersive environments. Such applications are particularly attractive for gaming, training simulations, and remote collaboration, where efficiency and realism are critical.

Beyond physical and virtual interaction lies the domain of cognitive augmentation. BCIs combined with LLMs could support enhanced learning, rapid memory retrieval, and complex problem solving. For instance, a student might use a BCI to silently query background information during a lecture, with the LLM interpreting neural intent and supplying structured explanations. Similarly, professionals could benefit from rapid recall of technical details or contextual knowledge, effectively creating a “thought-to-knowledge” interface. In the long run, this could blur the boundaries between natural cognition and artificial intelligence, opening new possibilities for education, research, and decision-making.

Finally, workplace productivity tools represent a near-term, practical frontier. Imagine a future in which office software, coding environments, or knowledge management platforms accept thought-driven commands mediated by LLM powered BCIs. Instead of navigating through menus, users could simply think, “summarize this report” or “open the last draft I was editing,” with the LLM parsing intent and generating the appropriate action. Such systems could reduce reliance on manual input devices and accelerate workflows, particularly for individuals engaged in knowledge-heavy tasks. Importantly, they would also improve accessibility for users with motor impairments, making productivity tools more inclusive.

Taken together, these application areas suggest that LLM enhanced BCIs are not limited to solving communication deficits but are poised to become foundational technologies for human–machine symbiosis. From restoring mobility and autonomy to augmenting cognition and creativity, the scope of BCI applications expands significantly once LLMs are incorporated as intelligent interpreters of neural intent.

7. Risks and Ethical Considerations

7.1. Over-Prediction

The predictive strength of LLMs, while beneficial for completing fragmented neural signals, carries the inherent

risk of “over-prediction” or hallucination. In the context of BCIs, this could mean the system inserts unintended words, phrases, or even entire sentences into a user’s communication stream. While such errors may seem minor in casual settings, they pose serious concerns in sensitive contexts such as medical communication, legal testimony, or professional correspondence, where every word carries weight. Unlike a smartphone autocorrect error, a hallucination in BCI mediated speech could misrepresent intent, damage trust, or even cause harm if commands are misinterpreted by assistive devices. Preventing over-prediction requires careful calibration of confidence thresholds, transparent uncertainty indicators, and user-in-the-loop correction methods to ensure the system amplifies rather than distorts the user’s voice.

7.2. Privacy Concerns

Privacy is a central ethical issue in the development of LLM-powered BCIs. By their nature, BCIs access neural activity that is often intimate and involuntary. When combined with LLMs capable of contextualizing these signals, there arises the risk of inadvertently extracting or revealing private thoughts that extend beyond what the user intends to share. For example, a fleeting association or subconscious signal could be amplified into a communicative output, unintentionally disclosing sensitive personal details. Furthermore, if neural data is stored or processed by external systems, concerns about surveillance, data misuse, and unauthorized access become magnified. Ethical BCI deployment must therefore prioritize strict data governance, encryption, consent frameworks, and mechanisms that give users granular control over what is shared and retained, ensuring mental privacy remains inviolable.

7.3. Bias Amplification

Another risk lies in the biases inherent to LLMs, which are often trained on massive datasets drawn from the internet. These datasets inevitably contain cultural, gender, racial, and ideological biases, which the models may replicate or even amplify. In a BCI context, such bias is particularly troubling because it could directly affect how a user’s intent is represented to others. For instance, predictive completions may favor certain language patterns, inadvertently altering tone or introducing stereotypes into communication. In assistive applications, biased outputs could further marginalize already vulnerable populations, undermining inclusivity and fairness. Addressing this challenge requires the incorporation of bias-detection mechanisms, continual auditing, and transparent reporting of how models are trained and deployed, ensuring that BCI outputs reflect user intent faithfully rather than embedding societal prejudices.

7.4. Safety and Alignment

Finally, ensuring the safety and alignment of

LLM-powered BCIs is critical, especially as these systems extend into high-stakes domains like healthcare, mobility, and workplace integration. Unlike traditional software, BCIs mediate a direct pathway between human cognition and machine action, meaning misalignment could result in unintended or even dangerous outcomes. For example, a misinterpreted neural signal amplified by an LLM could cause a robotic arm to perform an unsafe motion. To mitigate these risks, alignment mechanisms must ensure that user intent remains the guiding principle at every stage of processing. This includes maintaining human-in-the-loop oversight, implementing override mechanisms, and designing robust guardrails that prioritize safety and consent. Ethical alignment also demands transparency, allowing users to understand how their signals are being interpreted and corrected, thereby fostering trust in systems that directly mediate their thoughts.

8. Conclusion

The path forward for Large Language Model (LLM)-powered Brain-Computer Interfaces (BCIs) requires both technological innovation and ethical foresight. At the technical level, future research should emphasize the creation of hybrid architectures, where neural decoders, LLM post-processors, and multimodal integration layers function as a seamless pipeline. Such systems would not only decode raw neural signals but also refine them with linguistic fluency and contextual awareness, enabling communication that feels natural and reliable. A parallel priority lies in the design of personalization protocols that allow continuous adaptation to each individual’s neural and linguistic patterns. Rather than requiring lengthy calibration sessions, BCIs should support fluid fine-tuning, ensuring that the system evolves alongside its user and remains accurate under changing cognitive or environmental conditions.

Equally important is the rethinking of evaluation metrics. Traditional measures such as bit rate capture only the raw throughput of information, yet they fail to reflect communication quality or user experience. New metrics must emphasize fluency, fidelity to user intent, error robustness, and overall satisfaction. Similarly, progress will hinge on the development of multimodal BCI frameworks, where neural data is complemented by gaze, gesture, and environmental cues. By embedding brain signals into a richer context, these systems can achieve greater accuracy, intuitiveness, and functional diversity, moving beyond communication into domains such as assistive robotics, immersive environments, and cognitive augmentation.

At the societal level, governance frameworks must evolve in parallel with technical development. Issues of privacy, safety, bias, and alignment cannot remain afterthoughts. Neural data represents some of the most intimate information available, and any leakage or misinterpretation could have profound consequences. Transparent data handling, robust user consent, and safety-first alignment protocols will be

essential for ensuring that LLM-powered BCIs serve as trustworthy extensions of human cognition rather than sources of risk.

Looking ahead, the convergence of LLMs and BCIs represents more than incremental progress it marks a paradigm shift in human machine interaction. Over the next decade, clinical-grade systems for patients with paralysis or speech impairments are likely to evolve into consumer technologies that enable thought-driven interaction with everyday devices. As LLMs improve in multimodal reasoning, reduce bias, and achieve tighter alignment with human values, they will no longer serve as secondary processors but as copilots in communication, interpreting, structuring, and enhancing human thought. The implications are profound: healthcare will be transformed by more effective assistive devices, education could benefit from personalized cognitive augmentation, and workplace productivity may be reshaped by seamless, thought-driven tools.

In conclusion, Large Language Models provide a transformative pathway to address the longstanding challenges of BCIs. By combining probabilistic reasoning, contextual prediction, and multimodal integration, they bring us closer to fluent, adaptive, and accessible communication for all. Yet the promise of this convergence can only be realized if it is pursued responsibly, with careful attention to ethical design, personalization, and safety. The next generation of BCIs has the potential not only to restore lost abilities but to redefine the boundaries of human–AI symbiosis—reshaping how we communicate, create, and think in the decades to come.

Abbreviations

ALS	Amyotrophic Lateral Sclerosis
BCI	Brain Computer Interfaces
EEG	Electroencephalography
LLM	Large Language Models

Author Contributions

Gopichand Agnihotram: Writing – original draft, Conceptualization, Data curation

Joydeep Sarkar: Writing – original draft, Investigation, Methodology

Magesh Kasthuri: Visualization, Writing – review & editing

Conflicts of Interest


The authors declare no conflicts of interest.

References

- [1] Moses, D. A., Metzger, S. L., Liu, J. R., Anumanchipalli, G. K., Makin, J. G., Sun, P. F.,... & Chang, E. F. (2021). Neuroprosthesis for decoding speech in a paralyzed person with anarthria. *New England Journal of Medicine*, 385(3), 217-227.
- [2] Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature*, 593(7858), 249-254.
- [3] Ye, Z., Ai, Q., Liu, Y. *et al.* Generative language reconstruction from brain recordings. *Commun Biol* 8, 346 (2025). <https://doi.org/10.1038/s42003-025-07731-7>
- [4] Okpala, I., Golgoon, A., & Kannan, A. R. (2025). Agentic AI systems applied to tasks in financial services: modeling and model risk management crews. *arXiv preprint arXiv: 2502.05439*.
- [5] Wang Y, Tang Y, Wang Q, et al. Advances in brain computer interface for amyotrophic lateral sclerosis communication. *Brain-X*. 2025; 3: e70023. <https://doi.org/10.1002/brx2.70023>
- [6] Liu, Y., Ye, H., & Li, S. (2025). LLMs Help Alleviate the Cross-Subject Variability in Brain Signal and Language Alignment. *arXiv preprint arXiv: 2501.02621*.
- [7] Car à, A. (2025). Towards Predictive Communication: The Fusion of Large Language Models and Brain–Computer Interface. *Sensors*, 25(13), 3987.
- [8] Benster, T., Wilson, G., Elisha, R., Willett, F. R., & Druckmann, S. (2024). A cross-modal approach to silent speech with llm-enhanced recognition. *arXiv preprint arXiv: 2403.05583*.
- [9] Wang, P., Zheng, H., Dai, S., Wang, Y., Gu, X., Wu, Y., & Wang, X. (2024). A survey of spatio-temporal eeg data analysis: from models to applications. *arXiv preprint arXiv: 2410.08224*.
- [10] Li, H., Chen, Y., Wang, Y., Ni, W., & Zhang, H. (2025). Foundation Models for Cross-Domain EEG Analysis Application: A Survey. *arXiv preprint arXiv: 2508.15716*.
- [11] Inoue, M., Sato, M., Tomeoka, K., Nah, N., Hatakeyama, E., Arulkumaran, K.,... & Sasai, S. (2025). A Silent Speech Decoding System from EEG and EMG with Heterogenous Electrode Configurations. *arXiv preprint arXiv: 2506.13835*.
- [12] Dong, Y., Wang, S., Huang, Q., Berg, R. W., Li, G., & He, J. (2023). Neural decoding for intracortical brain–computer interfaces. *Cyborg and Bionic Systems*, 4, 0044.
- [13] Homer, M. L., Nurmikko, A. V., Donoghue, J. P., & Hochberg, L. R. (2013). Sensors and decoding for intracortical brain computer interfaces. *Annual review of biomedical engineering*, 15(1), 383-405.
- [14] Han, Y., Ziebell, P., Riccio, A., & Halder, S. (2022). Two sides of the same coin: adaptation of BCIs to internal states with user-centered design and electrophysiological features. *Brain-Computer Interfaces*, 9(2), 102-114.

- [15] Luo, S., Rabbani, Q., & Crone, N. E. (2022). Brain-computer interface: applications to speech decoding and synthesis to augment communication. *Neurotherapeutics*, 19(1), 263-273.

Biography



Gopichand Agnihotram is a Distinguished Member of Technical Staff (Senior Principal Member) and Director of AI/ML in Wipro Limited.



Joydeep Sarkar is Distinguished Member of Technical Staff (Principal Member) & Senior Architect in Wipro Limited.



Magesh Kasthuri is Distinguished Member of Technical Staff (Master) & Chief Architect in Wipro Limited.

Research Field

Gopichand Agnihotram: Artificial Intelligence, Generative AI, Agentic AI, Fake voice detection

Joydeep Sarkar: Artificial Intelligence, Generative AI, Agentic AI, Fake voice detection

Magesh Kasthuri: Artificial Intelligence, Generative AI, Agentic AI, AI for Data, Cloud technology