

Transcending natural evolution in the age of artificial intelligence

Pengwen GUO^{1,2†}, Xinyuan JIA^{1,2†}, Lu-Qi TAO^{2*}, Yi YANG^{1,2*} & Tian-Ling REN^{1,2*}

¹*School of Integrated Circuits, Tsinghua University, Beijing 100084, China*

²*Beijing National Research Center for Information Science and Technology (BNRist), Tsinghua University, Beijing 100084, China*

Received 6 May 2026/Revised 13 June 2026/Accepted 7 July 2026/Published online 31 August 2026

Abstract Artificial intelligence (AI) is undergoing a fundamental transition from digital representational models to embodied, physically situated systems. This position paper characterises this shift as the crossing of an embodiment threshold, where AI's operational capacity becomes integrated with physical environments and neural substrates. We argue that this transition is defined by a profound evolutionary tempo mismatch between rapid algorithmic iteration and the millennial-scale inertia of biological evolution. By synthesising advances in flexible bioelectronics and high-bandwidth neural decoding, we examine the biophysical mechanisms that could enable human-machine symbiosis while exposing the thermodynamic and material constraints that accompany this trajectory. In particular, we contrast the brain's highly efficient metabolic architecture with the resource-intensive scaling of silicon-based computation, highlighting the risk that the infrastructure required for large-scale augmentation may produce durable socio-technical stratification. We therefore contend that the central challenge of this era lies in governing neuro-computational co-evolution rather than merely regulating standalone AI systems. In this context, governance is not merely a regulatory constraint, but a deliberative mechanism for shaping whether, how, and under what conditions deeper human-machine integration should proceed. To address this shift, we propose an adaptive governance framework centred on infrastructural legibility, operational safety, and neural privacy. This roadmap outlines a strategic architecture for navigating the transition towards secure and equitable forms of human-machine symbiosis.

Keywords human-machine symbiosis, embodiment threshold, evolutionary tempo mismatch, brain-computer interface, adaptive governance, flexible bioelectronics

Citation Guo P W, Jia X Y, Tao L-Q, et al. Transcending natural evolution in the age of artificial intelligence. *Sci China Inf Sci*, 2026, 69(9): 193401, <https://doi.org/10.1007/s11432-026-5041-7>

1 Introduction

The historical trajectory of human biological and cultural adaptation can be understood as a gradual externalisation of physiological and cognitive functions. The emergence of spoken and written language extended working memory and enabled cumulative knowledge transmission across generations, facilitating increasingly complex forms of social coordination [1]. The industrial revolution shifted metabolic labour into mechanical systems governed by thermodynamic efficiency, and the development of classical computing formalised procedural logic and mathematical calculation in programmable architectures. In these earlier phases, technology largely functioned as a discrete instrument, activated and supervised by human operators within clearly bounded contexts.

The rapid proliferation of AI introduces a distinct dynamic. Algorithmic systems are no longer confined to narrow computational tasks. Instead, they increasingly operate as distributed infrastructures that influence scientific research, industrial systems, and everyday decision-making [2,3]. The long-articulated vision of ubiquitous computing has re-emerged as intelligence services become deeply embedded in consumer devices, networked environments, and edge architectures [2].

Over the past decades, machine learning has evolved from symbolic programs and expert systems to large-scale deep learning architectures and foundation models [4,5]. These systems process natural language, visual data, and heterogeneous sensor streams with increasing generality [6,7]. Empirical scaling laws indicate that performance improves predictably with training compute, data volume, and parameter count [8,9]. Phenomena such as few-shot prompting and in-context learning demonstrate that generalised models can adapt to new tasks without task-specific retraining [10]. Furthermore, contemporary multimodal systems are continuously updated and redeployed across global infrastructures with relatively low marginal costs.

* Corresponding author (email: taoluqi@tsinghua.edu.cn, yiyang@tsinghua.edu.cn, RenTL@tsinghua.edu.cn)

† These authors contributed equally to this work.

Table 1 Key concepts and operational definitions.

Concept	Operational definition
Embodiment threshold	The transition point at which an AI system moves beyond offline symbolic inference and enters a closed-loop relationship with a physical environment, the human body, or neural tissue. Operationally, this threshold is defined by three criteria: task-relevant interaction bandwidth, persistent or continuously available integration, and sensory-motor coupling. The threshold is therefore graded rather than fixed: robotic systems, wearable bioelectronics, and closed-loop BCIs represent progressively deeper forms of embodiment.
Evolutionary tempo mismatch	The quantifiable disparity between the millennial timescales of human biological adaptation and the sub-decadal iteration cycles of modern algorithmic systems, measured through comparative rates of functional change.
Directed co-evolution	A deliberate, governance-mediated process by which human biological capacities and machine intelligence are jointly shaped, in contrast to passive co-evolution driven solely by market forces or natural selection.
Biological-technological stratification	A durable socio-economic divide between individuals or populations who have access to performance-enhancing human-machine integration, including cognitive prosthetics and predictive health monitoring, and those who remain constrained by un-augmented biology.

These developments reposition AI as a field with broad scientific and economic implications [11]. Modern systems integrate perception, statistical inference, and real-time control, enabling them to generate persuasive content and autonomously coordinate physical devices at scale. A further critical transition is now emerging as scalable intelligence moves from digital representational domains into physical environments. The integration of vision-language-action models, advanced robotics, and brain-computer interfaces (BCIs) suggests a transition from advisory systems towards situated agents capable of direct material interaction and neural modulation [12–14].

As the rate of algorithmic iteration differs markedly from biological evolutionary timescales, this transition raises significant governance questions. Prevailing regulatory approaches have largely focused on privacy, intellectual property, and content integrity [15, 16]. However, systems that autonomously perceive, plan, and execute actions in physical or neural domains introduce complex safety and accountability challenges [17]. Biological adaptation unfolds over millennial timescales, whereas socio-technical systems can transform within years, creating temporal asymmetries in risk exposure. As algorithmic capabilities expand, the disparities between machine processing speeds and human cognitive limits have the potential to reshape labour markets, knowledge production, and geopolitical competition [18, 19].

Here we move beyond speculative narratives of AI risk to examine the empirical mechanisms underlying the embodiment threshold. We first quantify the evolutionary tempo mismatch between biological cognition and algorithmic iteration, establishing the central structural tension of the modern technological era. We then analyse how advances in flexible bioelectronics and neural decoding enable deeper forms of human-machine symbiosis while simultaneously revealing thermodynamic and material constraints that may drive socio-technical stratification. By bridging insights from soft-matter bioengineering and computational governance, we propose an adaptive roadmap structured around infrastructural legibility, operational safety, and neural privacy. We aim to provide a strategic framework for navigating the transition from tool-based AI towards an emerging regime of integrated, biologically coupled intelligence. This trajectory, however, is not an automatic consequence of technological scaling; it depends on deliberate institutional choices, social acceptance, and governance frameworks that can either enable, redirect, or restrain the depth of human-machine integration. The key concepts used throughout this position paper are summarised in Table 1.

2 Quantifying the evolutionary tempo mismatch

The foundational driver for investigating human-machine symbiosis is rooted in a quantifiable temporal and physical asymmetry. There is a substantial mismatch between the generational pace of natural biological adaptation and the iterative velocity of algorithmic scaling. Evolutionary biology and palaeoanthropology provide empirical baselines for this biological inertia. Anthropological evidence regarding skeletal loading, cortical folding, and bipedal biomechanics highlights the gradual pace at which biological bodies adapt to rapid technological and environmental change [20–22]. Major anatomical and physiological adaptations require millennia, encompassing thousands of generations, to achieve phenotypic fixation within a human population [23, 24]. While cultural and technological evolution has vastly expanded human cognitive reach through the creation of epistemic tools, the fundamental processing unit of the human central nervous system remains constrained by its evolutionary substrate.

The structural tension defining this era is most evident when the longitudinal trajectories of biological and technological change are superimposed (Figure 1). While the phenotypic characteristics of the human central




Figure 1 (Color online) The evolutionary tempo mismatch between biological cognition and technological capability. The blue curve illustrates the stabilisation of human brain size (reaching a functional plateau of approximately 1350 cm^3) over the past 300000 years, highlighting the inherent inertia of biological evolution. In contrast, the orange-to-red trajectory depicts the exponential acceleration in technological complexity, from the first primitive stone tools to the theoretical emergence of artificial general intelligence (AGI). While biological adaptation unfolds over millennial timescales dictated by metabolic and genomic constraints, the four successive industrial revolutions illustrate that socio-technical systems evolve on an increasingly compressed timeline. This divergence helps motivate the concept of an ‘embodiment threshold’, where the widening disparity between biological bandwidth and scalable computation may increase interest in more integrated forms of human–machine interaction, without making such integration inevitable.

nervous system, including cranial volume and synaptic transmission speeds, have reached a functional plateau, the complexity of the epistemic tools we deploy continues to scale exponentially. As illustrated, the time required for major anatomical shifts in *Homo sapiens* spans hundreds of thousands of years, yet the transition from the first industrial revolution to the era of ubiquitous computing and AGI has occurred within a mere three centuries. This divergence confirms that the fundamental processing unit of human biology remains constrained by its evolutionary substrate, creating a widening performance gap that cannot be bridged through natural selection alone.

Quantitatively, the human brain is an organ of remarkable architectural efficiency but functional physical strictures. It operates on a constrained metabolic energy budget of approximately 20 watts, largely dictated by the continuous ATP consumption required by sodium-potassium pumps to maintain neuronal resting potentials [25]. While the brain contains roughly 86 billion neurons interconnected by upwards of 100 trillion synapses, the fundamental speed of electrochemical neural transmission across synaptic clefts is inherently finite. These signals propagate at velocities between 1 and 100 metres per second, dictated heavily by the degree of axonal myelination. Furthermore, due to the absolute refractory periods required for ion channel repolarisation, the maximum firing rates of individual neurons rarely exceed 100 to 200 Hz. These physical properties, combined with the fixed architectural volume and thermal dissipation limits of the human cranium, impose upper bounds on individual cognitive bandwidth, working memory capacity, and parallel processing speed. Although the biological brain exhibits remarkable neuroplasticity, defined as the ability to dynamically rewire synaptic connections based on experiential learning, this plasticity is functionally bounded. It relies on slow and localised updating rules such as Hebbian spike-timing-dependent plasticity and must operate under the necessity of maintaining metabolic homeostasis [26–29]. On the evolutionary timescale, a single century of AI development is brief; yet, within this window, machines have acquired data-processing and generative abilities that human biology cannot match through genetic means alone.

By contrast, the structural mechanisms driving algorithmic capabilities are largely decoupled from biological homeostasis, though they face distinct macro-thermodynamic limits. The historical trajectory of machine learning provides a blueprint of accelerating computational capacity [3]. Successive architectural breakthroughs, from the

early development of convolutional neural networks that pioneered spatial hierarchies [30,31] to long short-term memory architectures for sequential data [32], established a highly effective toolkit for complex representation learning [4]. Crucially, the mathematical operationalisation of the backpropagation algorithm provided the optimisation principle that made these deep architectures computationally practical [33].

Today, scalable training frameworks, automatic differentiation paradigms, and distributed optimisation protocols allow researchers to coordinate tens of thousands of specialised processors [34,35]. Unlike biological neural networks, which rely on localised updating rules [36,37], deep learning leverages synchronised gradient descent operating at clock speeds in the gigahertz range and communicating via optical interconnects. Cloud platforms abstract this underlying hardware complexity to offer massive foundation models as a service [38]. These ubiquitous digital infrastructures now autonomously manage complex cognitive tasks including semantic memory storage, logical multi-step planning, and nuanced linguistic generation [39,40]. This transition effectively shifts judgement-intensive functions from individual humans to statistical systems.

Parallel to these algorithmic breakthroughs, hardware acceleration further widens the evolutionary tempo mismatch. Sustained scaling in semiconductor manufacturing has continuously enabled dense compute architectures tailored for highly parallel machine learning workloads [41,42]. Furthermore, the semiconductor industry is actively exploring computational paradigms beyond traditional von Neumann architectures. Emerging device concepts include memristive systems that seek to emulate biological synapses by physically unifying non-volatile memory storage and arithmetic processing in situ [43,44]. Two-dimensional materials, such as graphene and various transition metal dichalcogenides, also offer experimentally validated routes towards device integration architectures at the sub-silicon scale [45,46]. These alternative materials could bypass the impending thermodynamic and quantum tunnelling bottlenecks of current CMOS technology. In hyperscale data centres, intricate energy management dynamically integrates advanced liquid cooling systems and predictive workload scheduling to maximise computational output per joule.

Taken together, these hardware and software advances do not merely expand raw computational capacity. They underpin the tempo mismatch between a rapidly reconfigurable machine intelligence and a more inertial human biology. Frontier models driven by this hardware-software synergy can be updated, mathematically refined, and redeployed at a global scale [47]. This infrastructure supports iteration cycles in which complex neural architectures and safety mitigations can be reconfigured in weeks or months, rather than the generations required for biological adaptation. Viewed through this lens, human enhancement emerges as a systemic consequence of macro-level infrastructure design. This structural shift raises governance questions about who controls the externally mediated capacities that will influence our evolutionary future [48].

3 Crossing the embodiment threshold towards situated symbiosis

Historically, AI has developed predominantly in disembodied forms. Large-scale foundation models and statistical optimisation algorithms operate over symbolic representations, including vast text corpora, static images, programming code, and structured tabular data, without an autonomous physical presence in the material world. Their outputs influence human cognition, corporate communication, and political decision-making, yet their operational locus remains largely confined to the digital realm of server infrastructures. In this configuration, intelligence is treated primarily as a mathematical optimisation problem over pre-collected representations rather than as situated action in the physical world.

This historical separation between algorithmic cognition and physical embodiment is an artefact of technological limitations rather than a conceptually necessary boundary. Foundational perspectives in robotics and theories of embodied cognition, drawing from ecological psychology and cybernetics, have long argued that robust intelligence emerges from continuous perception-action loops embedded in material environments. Early pioneering work on behaviour-based robotics challenged top-down representational models by demonstrating that adaptive behaviour can arise directly from real-time interaction with physical constraints [49]. The recent scaling of multimodal foundation models introduces structural pressure that blurs this separation [50,51]. As large models acquire general representations, high-fidelity multimodal perception, and tool-use behaviours through self-supervised pre-training, their competencies extend beyond linguistic prediction [52]. They increasingly move towards dynamic task coordination and complex decision sequencing. Once these models can integrate high-dimensional spatial inputs, interpret natural-language goals, and decompose physical tasks into executable motor commands, the strict boundary between digital inference and situated action increasingly dissolves [53–55]. Embodiment can therefore be understood as one possible extension of scalable intelligence into physical intervention, contingent on engineering feasibility, safety constraints, and social governance [56,57].




Figure 2 (Color online) Crossing the embodiment threshold towards human-machine integration. The evolutionary trajectory of artificial intelligence transitions across distinct operational domains. In the digital domain, disembodied intelligence functions primarily as a mathematical optimisation process over pre-collected symbolic representations, lacking autonomous physical presence. Driven by the scaling of multimodal foundation models, AI may cross an embodiment threshold when scalable intelligence becomes coupled to direct physical intervention. Within the physical environment, embodied intelligence operates through situated action, where robust capabilities emerge from continuous perception-action loops and dynamic task coordination. Finally, to transition into the neural domain, advances in flexible bioelectronics are actively engineered to overcome the profound biomechanical mismatch between rigid silicon and soft biological tissue. Such material innovation may enable deeper forms of human-machine integration, facilitating high-bandwidth, bidirectional information flows that produce a symbiotic redistribution of sensory perception and kinetic control.

The required interaction bandwidth is application-specific: wearable physiological monitoring, robotic control, and intracortical neural decoding operate at different temporal scales [54, 58, 59]. What distinguishes embodied AI is therefore not a fixed update frequency, but the presence of a task-relevant closed loop through which the system continuously senses, interprets, and influences physical, bodily, or neural states [49, 53].

One possible, but ethically and technically contested, frontier of this embodiment threshold is the integration of AI more directly with the human biological form [60, 61]. This symbiotic integration redistributes sensory perception, kinetic control, and physical endurance between humans and machines [19]. It produces forms of physical and cognitive augmentation alongside novel systemic dependencies, rather than acting as a simple task substitution (Figure 2).

The material gateway for crossing the embodiment threshold is formed by a suite of bioelectronic technologies that resolve the biomechanical mismatch between rigid computation and soft biological tissue, thereby enabling AI systems to acquire continuous bodily data streams through bioelectronic interfaces. To achieve this state of symbiosis, the scientific field must overcome profound biomechanical, immunological, and engineering challenges. Parallel advances in stretchable, flexible, and textile electronics are expanding the material substrate for embodied sensing and actuation on and within the human body. A historical barrier to human-machine integration has been the biomechanical mismatch between rigid electronics and soft biological tissue. Rigid silicon electronics and soft biological tissues differ by several orders of magnitude in mechanical stiffness. This mechanical impedance mismatch can cause micro-trauma, localised inflammation, and signal degradation during natural physiological movement.

However, the micro-scale engineering of advanced materials is actively crossing this threshold. Surface wrinkling techniques and related mechanical strategies improve the conformal contact and robustness of wearable sensors [62, 63]. These innovations allow computational elements to flex and interface seamlessly with the epidermis. Published roadmaps for flexible sensors highlight both the opportunities and the remaining manufacturing bottlenecks for large-scale clinical deployment [64]. Recent technology roadmaps for bioinspired computing hardware further indicate that the convergence of flexible bioelectronics and neuromorphic computation is approaching the material and engineering maturity required for stable human-machine interfaces, and key device-level challenges, including functional stability, scalable fabrication, and bio-interface compatibility, are being systematically addressed in ways that directly align with the requirements for crossing the embodiment threshold [65]. Furthermore, the mechanical buckling of nanoscale ribbons offers an additional structural route to achieving multi-axial stretchability in inorganic semiconductors [66]. Moving beyond surface contact sensors, pioneering work on multimaterial

fibres, self-healing conformal structures, and sophisticated fabric computing illustrates a paradigm shift [67–69]. It shows how cognitive and physiological augmentation can be engineered as a system property. Smart textile displays and highly sensitive triboelectric sensing arrays further expand the interface between human bodies, dynamic environments, and algorithmic computation [70, 71].

The miniaturisation of sensing elements offers routes to selectively filtered and multimodal perception. Intelligent eardrums equipped with edge computing processors can isolate acoustic signatures in high-noise environments [22]. Simultaneously, monolithically integrated in-textile wristband systems demonstrate how continuous biosensing can be embedded invisibly in garments [72]. These smart textiles thus convert the body surface into a distributed sensor array that streams data directly into AI models for real-time health inference, exemplifying how computation moves from detached servers onto the living body. This work provides a relevant example of device-level integration, where imaging, compression, and wireless transmission are combined within an in-sensor computing architecture [73]. These systems monitor glucose, cortisol, and complex protein markers, dynamically feeding streaming metabolic data into predictive health models [74, 75]. In the realm of communicative and motor restoration, wearable voice interfaces showcase the potential of embodied sensing. The development of artificial throats and ultra-thin graphene transducers illustrates how these interfaces can capture subtle myoelectric and mechanical throat vibrations [76–78]. By capturing neuromuscular signals and translating them into speech, such interfaces demonstrate how AI can be physiologically coupled to human expressive intent, shifting from external tool to an extension of the body’s own communicative apparatus. Such systems can restore speech modalities without requiring invasive surgical procedures. Moving deeper beneath the epidermal layer, bioengineering and microrobotic therapeutics are extending computational intervention directly into cellular and molecular scales. Artificial microtubule systems designed for computationally directed collective microcargo transport further illustrate how engineered micromechanics can extend therapeutic intervention [79–81].

Human-machine integration reaches a highly complex stage and faces significant obstacles in the domain of BCIs [12, 82]. Advanced interfaces strongly hint at near-term futures in which high-bandwidth information flows bidirectionally between living neural tissue and global digital networks [83]. Historically, rigid invasive neural interfaces have often suffered from foreign-body responses and progressive signal degradation, as mechanical mismatch can trigger inflammatory encapsulation around the electrode-tissue interface. Today, flexible and conformal interface designs aim to mitigate this immune response to improve long-term safety [82, 84]. When designed to decode cortical signals in real time, these interfaces represent the deepest level of coupling considered in this position paper, where algorithmic systems interact with neural signals associated with intention, communication, or motor control. Innovations include ultra-thin polymer threads inserted by precision surgical robots designed to avoid disrupting delicate cortical vasculature [85]. Similarly, endovascular stents can be deployed through the jugular vein directly into the motor cortex to bypass the risks associated with open cranial surgery [86]. When designed safely and ethically, such interfaces hold the potential to restore lost sensorimotor function and make abstract computational patterns graspable to the human mind. However, this intimate integration depends on reliable and continuous on-body power delivery. This dependency necessitates the simultaneous development of body-coupled power transmission, implantable energy harvesting, high-performance fibre batteries, and ultrasound-driven harvesters designed for powering deep-tissue implants [87–90].

4 Thermodynamic limits and the risk of socio-technical stratification

The theoretical emergence of directed co-evolution through embodied intelligence and neuro-computational prosthetics does not imply an escape from the laws of physics or the complexities of neurobiology. The translation of raw biological neural signals into actionable digital commands requires immense algorithmic sophistication. Biological neural signals, whether measured via local field potentials, electrocorticography, or single-unit action potential spikes, are inherently noisy, chaotic, and non-stationary. Recent landmark clinical studies have successfully utilised deep learning architectures, specifically recurrent neural networks and highly optimised transformer models, to decode these noisy signals into accurate and real-time synthesised text and speech for patients with severe paralysis [58, 59]. In these clinical trials, decoding rates have exceeded 60 to 70 words per minute, approaching the speed of natural human conversation. Here, AI acts as a dynamic translation layer. It utilises its pre-trained linguistic priors to map biological intent into precise digital action.

However, these technical advantages are linked to systemic dependencies. They translate into cognitive and physical capabilities that are structurally difficult for un-augmented humans to match. A single frontier model can process and synthesise more text in an hour than a person can read in a lifetime, track millions of interrelated variables across complex thermodynamic simulations, and maintain fine-grained logs of its own systemic behaviour. Coupled




Figure 3 (Color online) Thermodynamic limits and material constraints on AI capability scaling. Estimated training-related carbon emissions (tons CO₂ equivalent) for representative AI models are plotted against model parameter counts on logarithmic axes. The carbon-emission estimates are compiled from the 2025 AI Index Report, which combines developer-reported emissions with calculator-based estimates using available information on training hardware, training duration, cloud provider, and training region. Because these values are not normalised to a single hardware platform, data-centre efficiency, electricity mix, or experimental-run protocol, the fitted trend line should be interpreted as an indicative association rather than a deterministic scaling law. Horizontal reference lines indicate emissions associated with representative human activities, including a single transcontinental passenger flight, average annual per-capita emissions, and lifetime passenger-car use, providing interpretable benchmarks for the environmental scale of large-model training.

with advanced robotics, such systems perceive and manipulate the physical world with sub-millimetre precision. They maintain around-the-clock performance without experiencing biological fatigue. Domain-specific variants directly assist in scientific discovery, industrial optimisation, and personalised recommendations in healthcare. As artificial systems autonomously assist in designing new metamaterials, pharmaceutical drugs, and next-generation semiconductor devices, they feed back into subsequent generations of hardware and software [91]. This dynamic creates an accelerating loop in which algorithms actively improve the physical substrates on which they run.

Yet, the mathematical and computational dynamics that enable this accelerating loop also highlight physical planetary resource constraints and systemic vulnerabilities [92, 93]. To contextualise the material footprint of this scaling trajectory, Figure 3 visualises the estimated carbon emissions associated with training representative foundation models across several generations of deep learning systems.

When plotted against model parameter counts on logarithmic axes, the data indicate an approximate association between model scale and reported training-related emissions, while also reflecting substantial variation introduced by training protocol, hardware efficiency, and electricity supply. Training state-of-the-art frontier models demands substantial volumes of computation, gigawatt-scale electrical energy generation, and extensive specialised data centre capacity [94]. These hard material and energy constraints translate technical scale directly into socio-political asymmetries of agency. They can concentrate the capacity to shape or redirect mediated forms of human augmentation largely within the domains of actors who possess the capital to control compute clusters, energy grids, and physical infrastructure. Progress at the frontier of machine intelligence is precariously dependent on a geopolitically contested global supply chain. This complex chain encompasses the fabrication of advanced extreme ultraviolet lithography machines, the engineering of semiconductor substrates, the resource-intensive extraction of rare-earth elements, and the construction of high-capacity electricity infrastructure. Inference at a global scale consumes significant power and quantities of cooling water. These material realities raise immediate concerns regarding the ecological sustainability of ubiquitous computing, its carbon emissions, and its competition with other societal sectors for finite planetary resources [95].

Model performance also remains sensitive under algorithmic distribution shifts, adversarial prompt injection, and long-horizon causal planning [96,97]. This fragility reminds us that statistical pattern recognition over vast datasets is distinct from robust causal human understanding [98,99]. Many systems still struggle with systematic reasoning, counterfactual causal inference, and value-sensitive judgement [100,101]. The longstanding vision of ubiquitous computing is increasingly intertwined with these resource constraints. This dynamic echoes the historical argument in technological economics that scalable computation and heavy physical capital often outweigh hand-designed knowledge and localised human labour [102].

Furthermore, large models learn primarily from historically constrained human corpora and are mathematically optimised via reinforcement learning for majority preferences and standardised safety filters [13]. Consequently, they are argued to drift towards forms of cognitive convergence. They systemically reproduce familiar arguments, homogenised logical paths, and dominant cultural styles rather than exploring heterodox regions of the idea space. When such cognitively converged systems mediate search engines, legal drafting, academic peer review, and eventually the cognitive prosthetics integrated into human brains, they can narrow the diversity of ideas that human researchers and institutions routinely explore. At scale, the automation of these judgement-intensive functions has implications for how epistemic authority, legal responsibility, and democratic decision-making power are distributed between biological humans and opaque statistical optimisation systems [103].

These intersecting material and epistemic trends produce a structural asymmetry with socio-economic consequences, threatening to precipitate a societal challenge that we term biological-technological stratification. The widespread adoption of these technologies in cognitive and physical automation is poised to reshape the global labour market. While algorithms have the potential to increase efficiency across industries, they present displacement risks for workers in the manufacturing, clerical, and logistics sectors. High-skilled professions will still experience profound transformations. This is especially true in industries like healthcare, finance, and law, where AI improves, mediates, or replaces human decision-making processes. If access to advanced forms of human-machine integration, including high-bandwidth cognitive neural prosthetics and predictive health monitoring, is constrained by cost, geopolitics, or fragile supply chains, humanity faces the risk of durable cognitive and economic divergence. This could manifest as a structural divide between an augmented population with enhanced cognitive bandwidth and physical endurance, and an un-augmented population constrained strictly by traditional biological limits. Governments and global industries must therefore invest in proactive reskilling initiatives and robust public digital infrastructure. They must ensure that displaced workers can transition to new roles within emerging hybrid fields while guaranteeing that the economic and biological benefits of technology adoption are shared equitably. At this critical point, human sovereignty is no longer only a matter of passive cognitive oversight, but of maintaining democratic oversight over physical agency, morphological risk allocation, and the material resources required for our evolutionary adaptation [104].

5 An adaptive governance roadmap for human-machine integration

The structural asymmetries detailed above demand an operationalised reimagining of global technological oversight [105]. Algorithms driving cognitive augmentation can be dynamically redesigned, scaled, and redeployed on timescales of mere months. In contrast, human biological evolution and legal institutions adapt on far longer generational timescales, remaining constrained by material, environmental, and geopolitical limits. Legal and ethical frameworks face challenges to keep pace with urgent questions about biometric privacy, cognitive data governance, intellectual property rights, distributed physical liability, and democratic responsibility when AI is intimately braided into consequential human decisions. Because model weights, cross-border data flows, and the atmospheric carbon emissions generated by training clusters inherently do not respect national boundaries, establishing a robust global governance framework becomes the central institutional realisation of the directed co-evolution paradigm [104,106,107]. While biological adaptation relies on the passive filter of natural selection, our transition into human-machine symbiosis must be actively mediated by adaptive legal and ethical architectures.

National strategies and fragmented sectoral regulations, while politically important, will remain insufficient if highly capable systems can be trained in permissive jurisdictions with weak oversight and subsequently exported worldwide via cloud infrastructure or physically embodied hardware. As these technologies expand globally, they encounter friction arising from cross-cultural and regulatory disparities. Sovereign countries possess varying legal frameworks and historical norms regarding privacy, technological ethics, and acceptable use. For instance, European regulatory frameworks mandate stringent and risk-tiered data privacy controls, prioritising fundamental human rights, algorithmic explainability, and citizen safety over frictionless corporate innovation [16]. Conversely, the United States adopts a relatively more lenient and market-driven stance. It relies on sector-specific guidelines,


Governance dimensions	Phase I: Infrastructural legibility	Phase II: Embodied safety & autonomy	Phase III Transnational equity
 Compute & ecological accountability	• Mandatory compute & energy disclosure • Training data provenance	• Public-interest evaluation platforms • Immutable incident logging	• Transnational oversight bodies (e.g., CERN/IAEA models) • Localised edge computing networks
 Operational safety & auditing	• Verifiable interface visibility • Rigorous capability red-teaming	• Harmonised engineering standards • Algorithmic drift assessment	• Durable cross-border infrastructure • Public-interest oversight of training clusters
 Neural privacy & data sovereignty	• Legal neuroprivacy classification • Strict medical data protections	• Anti-commodification of brain signals • Protection of biological intent	• Democratised access to computational models • Preventative health monitoring as human right
 Hybrid user rights & liability	• Standardised disclosure at release	• Functional right to disconnect • Physical and software fail-safes	• Modernised liability frameworks • Untangling biological vs. algorithmic intent

Figure 4 (Color online) An adaptive governance roadmap for human–machine integration. Rather than relying on static treaties, this continuous framework triggers regulatory intervention based on measurable technological indicators and the depth of symbiotic integration. The matrix maps four critical governance dimensions across three sequential phases. Phase I (infrastructural legibility) establishes empirical visibility into macro-level compute substrates, mandating the standardised disclosure of gigawatt energy utilisation and training data provenance, alongside the strict medical classification of raw neural signals to establish baseline neuroprivacy. As artificial systems cross the embodiment threshold to act autonomously in physical or neural domains, Phase II (embodied safety & autonomy) enforces harmonised engineering standards, evaluates algorithmic drift, prevents the commodification of biological intent, and guarantees a functional right to disconnect via robust software and physical fail-safes. Finally, to mitigate the systemic risk of biological-technological stratification, Phase III (transnational equity) necessitates the modernisation of liability frameworks and the establishment of transnational oversight bodies, drawing on models such as CERN or the IAEA, to democratise access to computational infrastructure and treat predictive health monitoring as a fundamental human right.

litigation-based precedents, and voluntary corporate safety commitments, maintaining a focus on fostering rapid technological innovation [108,109]. Meanwhile, AI development, data aggregation, and generative algorithms in China are legally tethered to national strategic interests, social stability, and data sovereignty, creating a highly centralised regulatory environment [110]. These regulatory differences inevitably lead to jurisdictional conflicts when ubiquitous and embodied systems are deployed across global borders.

Creating a globally harmonised and adaptive governance framework therefore becomes increasingly critical, particularly concerning cognitive data privacy, algorithmic transparency, and ethical standards for human–machine integration [111–113]. To translate this macroscopic socio-technical framework into actionable policy, governance must not rely on static treaties that quickly become technologically obsolete, nor should it consist of fragmented regulatory mandates. Instead, we propose a continuous and narrative-driven roadmap of conditional oversight, where successive phases of regulatory intervention are triggered by measurable technological indicators and the depth of human–machine integration (Figure 4).

In the initial phase of infrastructural legibility, governance must respond directly to the rapid scaling of compute paradigms. This foundational phase must focus on establishing empirical and verifiable visibility into both the macro-level compute substrate and the micro-level biological interface [114,115]. This requires a regulatory shift away from corporate secrecy towards the standardised and mandatory disclosure of compute allocation, gigawatt energy utilisation, and precise training data provenance at the point of frontier model release [116,117]. Governance at this stage demands the implementation of immutable incident logging [118], rigorous red-teaming protocols

probing for capabilities, and the establishment of independent public-interest evaluation platforms [119]. Crucially, as biological integration deepens through wearable and non-invasive sensors, governance must strictly define and legally protect the concept of neuroprivacy. The raw physiological and electrical signals generated by a human brain interfacing with a neural prosthetic, including continuous EEG recordings or spike train telemetry, must be legally classified under stringent medical data protections. This classification must prevent their extraction, commodification, or monetisation as standard consumer telemetry data by technology vendors, ensuring that the biological intent of the user cannot be harvested for behavioural manipulation.

As artificial systems decisively cross the embodiment threshold by acting autonomously in the physical world or integrating deeply with the human body via smart biomaterials, governance must seamlessly transition to a stricter second phase dedicated exclusively to the safety, autonomy, and legal rights of the hybrid user [120, 121]. This phase explicitly connects higher-risk capability tiers to harmonised international engineering standards. Regulatory health bodies must establish rigorous clinical protocols for assessing both biocompatibility and algorithmic drift in neural interfaces [122]. They must require empirical proof of long-term signal stability and immune-response mitigation over chronic timelines before permitting commercialisation. Furthermore, socio-technical design and legal frameworks must guarantee a functional right to disconnect. This principle requires engineering robust physical and software fail-safe mechanisms that allow individuals to securely sever an algorithmic connection without suffering physiological withdrawal, psychological trauma, or critical functional loss. Liability frameworks will also require profound and nuanced modernisation. Courts will need new jurisprudential tools to untangle the conscious intent of a biological operator from the automated predictive actions of an integrated corporate algorithmic model when a hybrid entity causes physical or economic harm. Errors in this paradigm are no longer confined to misinformation; they manifest directly as physical harm, system failure, or infrastructural disruption.

A distinct regulatory challenge arises when brain-computer interfaces employ algorithms that update their parameters online based on incoming neural data. Traditional audit paradigms assume a frozen decoder that can be tested against a hold-out validation set, but in an online-learning BCI, decoder adaptation may interact with the user's own neural plasticity, producing a co-adaptive system rather than a static medical device [123]. Operational safety in this setting therefore cannot rely only on pre-deployment certification; it must also govern how model updates are proposed, validated, deployed, monitored, and reversed. In practice, the adaptive decoder should be separated from a fixed safety layer that constrains the allowable output space, such as the range of motor commands, stimulation parameters, or device-level actions. Proposed model updates should enter a sandboxed validation stage before they can affect physical effectors or neural stimulation. Each update should be versioned and logged, with records of the model state, neural-data distribution, output commands, user override events, and safety-relevant anomalies. Predefined drift and performance thresholds should trigger rollback to the last validated decoder or to a conservative fallback mode. This workflow extends the concept of independent AI auditing [114, 115] and formal safety envelopes for autonomous systems [124], shifting operational safety from one-time model evaluation towards continuous oversight of the update pipeline itself.

Addressing the ultimate open challenges in transnational infrastructure and evolutionary equity constitutes the most advanced phase of this governance roadmap. Recognising the plausible risk of biological-technological stratification, international discussions must actively explore and deploy mechanisms to democratise access to advanced computational models. This necessitates exploring policies aimed at providing baseline access to computational infrastructure, ensuring that cognitive augmentation and predictive preventative health monitoring are treated as fundamental human rights rather than luxury commodities. Achieving this requires the establishment of durable cross-border infrastructure for global registries and the rigorous public-interest oversight of massive training clusters. Such architectures could potentially draw organisational and structural lessons from transnational bodies like the International Atomic Energy Agency or the European Organization for Nuclear Research, explicitly adapted for the governance of planetary-scale computation [125]. Additionally, actively exploring and funding decentralised architectures and localised edge computing provides a critical technical pathway to mitigate excessive societal reliance on monopolistic corporate cloud servers. These localised models must be capable of running highly efficient and quantised algorithms on personal hardware under strict thermal envelopes. Global ethics and health guidance must emphasise inclusive participation, ensuring that the trajectory of directed co-evolution is not dictated solely by the geopolitical interests of the entities controlling the advanced semiconductor supply chain.

Operationalising the roadmap through existing institutions, the four governance dimensions shown in Figure 4 can be grounded in existing institutional structures across the three implementation phases of infrastructural legibility, embodied safety and autonomy, and transnational equity. First, compute and ecological accountability can be linked to international standards bodies and public evaluation institutions, including ISO/IEC JTC 1 and IEEE, which are already developing frameworks for AI management and auditing, such as ISO/IEC 42001:2023 [121]; however, extending these frameworks to mandatory compute-reporting templates, energy-use disclosure, and cross-

jurisdictional enforcement remains an open challenge [114, 115]. Second, operational safety and auditing align with the mandate of medical-device and safety regulators, including the U.S. FDA and China's NMPA, whose current protocols for implantable neural devices are well established for relatively static products but are not yet designed for the algorithmic drift characteristic of continuously learning systems. Third, neural privacy and data sovereignty can be overseen by data-protection and health-data authorities operating under frameworks such as the GDPR [15] and China's Personal Information Protection Law [110], provided that raw neural signals are explicitly classified as sensitive biometric or medical data. Fourth, hybrid user rights and liability require coordination among courts, product-liability regimes, insurance systems, and legislative bodies, because future disputes may need to distinguish biological intent, algorithmic recommendation, and device-level execution in tightly coupled human-machine systems. At the transnational level, treaty-based coordination inspired by CERN or the IAEA may further support equitable access and public-interest oversight of large-scale computational infrastructure [125], although this would depend on geopolitical will and funding parity. By anchoring these governance dimensions to specific actors, we emphasise that the roadmap is not an abstract ideal but a set of institutionally grounded and negotiable mandates.

6 Conclusion and outlook

Transcending the traditional constraints of natural evolutionary inertia in the age of rapidly scaling machine intelligence does not necessarily signify the obsolescence of biological judgement. Nor does it advocate for a fatalistic surrender to deterministic algorithms. Rather, it demands a scientifically grounded and ethically robust recalibration of the entire paradigm of human-machine symbiosis. The accelerating transition from disembodied representational AI to embodied and situated intelligence marks a measurable and pivotal transition. This trajectory extends to intimate and symbiotic cognitive prosthetics. At this threshold, AI transitions from a digital advisory tool to an operative system within the physical, physiological, and social fabric of human existence.

The central challenge of the coming decades is to proactively engineer the socio-technical systems that mediate this directed co-evolution. We must ensure that these systems support collective human agency, maintain epistemic and cognitive diversity, and bolster societal resilience against systemic shocks. It is imperative to critically navigate the material, energetic, and thermodynamic constraints of capability scaling. Failure to do so risks entrenching infrastructural dependencies on corporate oligopolies, or worse, exacerbating existing global wealth inequalities into durable socio-technical stratification. Addressing this profoundly complex landscape requires a continuous and internationally coordinated roadmap. This framework must seamlessly align infrastructural legibility (e.g., empirical evaluation of frontier models) with rigorous neural privacy standards, operational safety benchmarks, and democratised access to the computational substrates of adaptation.

Whether human-machine convergence can responsibly extend biological capacities will not be determined solely by compute scaling laws or by the biological limitations of organic neural networks. Rather, our collective institutional capacity to deliberate, design, and enforce scientific frameworks has now become the primary engine of our evolutionary future. These frameworks must preserve the fundamental human right to reflect, to dissent, and to maintain democratic oversight over the transformative technologies we purposefully deploy into the physical world.

Acknowledgements This work was supported by National Key R&D Program of China (Grant No. 2022YFB3204100) and National Natural Science Foundation of China (Grant Nos. 92580209, U25A6021, 62405157).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- 1 Boyd R. The evolution of human uniqueness. *Span J Psychol*, 2017, 19: 97
- 2 Weiser M. The computer for the 21st century. *Sci Am*, 1991, 265: 94–104
- 3 Jordan M I, Mitchell T M. Machine learning: trends, perspectives, and prospects. *Science*, 2015, 349: 255–260
- 4 Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. 2012. ArXiv:1206.5538
- 5 Schmidhuber J. Deep learning in neural networks: an overview. *Neural Netws*, 2015, 61: 85–117
- 6 Zhao W X, Zhou K, Li J, et al. A survey of large language models. 2023. ArXiv:2303.18223
- 7 Minaee S, Mikolov T, Nikzad N, et al. Large language models: a survey. 2024. ArXiv:2402.06196
- 8 Kaplan J, McCandlish S, Henighan T, et al. Scaling laws for neural language models. 2020. ArXiv:2001.08361
- 9 Hoffmann J, Borgeaud S, Mensch A, et al. Training compute-optimal large language models. 2022. ArXiv:2203.15556
- 10 Brown T B, Mann B, Ryder N, et al. Language models are few-shot learners. 2020. ArXiv:2005.14165
- 11 Jarrahi M H. Artificial intelligence and the future of work: human-AI symbiosis in organizational decision making. *Bus Horiz*, 2018, 61: 577–586

- 12 Lebedev M A, Nicoletti M A L. Brain-machine interfaces: past, present and future. *Trends Neurosci*, 2006, 29: 536–546
- 13 Biever C. ChatGPT broke the Turing test—the race is on for new ways to assess AI. *Nature*, 2023, 619: 686–689
- 14 An Z, Yu X, Wang C, et al. Embodied intelligence: recent advances and future perspectives. *Innov Inform*, 2025, 1: 100001–100008
- 15 Kuner C, Bygrave L A, Docksey C. The EU General Data Protection Regulation: A Commentary. Oxford: Oxford University Press, 2020
- 16 European Parliament, Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 (Artificial Intelligence Act). *OJ L* 2024/1689. 2024
- 17 Awad E, Dsouza S, Kim R, et al. The moral machine experiment. *Nature*, 2018, 563: 59–64
- 18 Bostrom N, Sandberg A. Cognitive enhancement: methods, ethics, regulatory challenges. *Sci Eng Ethics*, 2009, 15: 311–341
- 19 Raisamo R, Rakkolainen I, Majaranta P, et al. Human augmentation: past, present and future. *Int J Hum Comput Stud*, 2019, 131: 131–143
- 20 Chirchir H, Kivell T L, Ruff C B, et al. Recent origin of low trabecular bone density in modern humans. *Proc Natl Acad Sci USA*, 2015, 112: 366–371
- 21 Gruss L T, Schmitt D. The evolution of the human pelvis: changing adaptations to bipedalism, obstetrics and thermoregulation. *Phil Trans R Soc B*, 2015, 370: 20140063
- 22 Dayer G, Guy F, Mackaye H T, et al. Postcranial evidence of late Miocene hominin bipedalism in Chad. *Nature*, 2022, 609: 94–100
- 23 Aiello L C, Wheeler P. The expensive-tissue hypothesis: the brain and the digestive system in human and primate evolution. *Curr Anthropol*, 1995, 36: 199–221
- 24 Ruff C B, Trinkaus E, Holliday T W. Body mass and encephalization in Pleistocene Homo. *Nature*, 1997, 387: 173–176
- 25 Balasubramanian V. Brain power. *Proc Natl Acad Sci USA*, 2021, 118: 2107022118
- 26 Karbowski J. Metabolic constraints on synaptic learning and memory. *J Neurophysiol*, 2019, 122: 1473–1490
- 27 Attwell D, Laughlin S B. An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab*, 2001, 21: 1133–1145
- 28 Markram H, Lüscher B, Frotscher M, et al. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 1997, 275: 213–215
- 29 Bi G Q, Poo M M. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 1998, 18: 10464–10472
- 30 Rosenblatt F. The perception: a probabilistic model for information storage and organization in the brain. *Psychol Rev*, 1958, 65: 386–408
- 31 Lecun Y, Bottou L, Bengio Y, et al. Gradient-based learning applied to document recognition. *Proc IEEE*, 1998, 86: 2278–2324
- 32 Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*, 1997, 9: 1735–1780
- 33 Rumelhart D E, Hinton G E, Williams R J. Learning representations by back-propagating errors. *Nature*, 1986, 323: 533–536
- 34 Chen T, Li M, Li Y, et al. MXNet: a flexible and efficient machine learning library for heterogeneous distributed systems. 2015. [ArXiv:1512.01274](https://arxiv.org/abs/1512.01274)
- 35 Barham P, Chowdhery A, Dean J, et al. Pathways: asynchronous distributed dataflow for ML. 2022. [ArXiv:2203.12533](https://arxiv.org/abs/2203.12533)
- 36 Hopfield J J. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA*, 1982, 79: 2554–2558
- 37 Kohonen T. The self-organizing map. *Neurocomput*, 1998, 21: 1–6
- 38 Bommasani R, Hudson D A, Adeli E, et al. On the opportunities and risks of foundation models. 2021. [ArXiv:2108.07258](https://arxiv.org/abs/2108.07258)
- 39 Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. 2017. [ArXiv:1706.03762](https://arxiv.org/abs/1706.03762)
- 40 Devlin J, Chang M-W, Lee K, et al. BERT: pre-training of deep bidirectional transformers. 2019. [ArXiv:1810.04805](https://arxiv.org/abs/1810.04805)
- 41 Moore G E. Cramming more components onto integrated circuits, reprinted from *Electronics*, volume 38, number 8, April 19, 1965, pp.114 ff. *IEEE Solid-State Circuits Soc Newsl*, 2006, 11: 33–35
- 42 Sevilla J, Heim L, Ho A, et al. Compute trends across three eras of machine learning. 2022. [ArXiv:2202.05924](https://arxiv.org/abs/2202.05924)
- 43 Zidan M A, Strachan J P, Lu W D. The future of electronics based on memristive systems. *Nat Electron*, 2018, 1: 22–29
- 44 Lin Y, Zhang Q, Gao B, et al. Uncertainty quantification via a memristor Bayesian deep neural network for risk-sensitive reinforcement learning. *Nat Mach Intell*, 2023, 5: 714–723
- 45 Wang S, Liu X, Xu M, et al. Two-dimensional devices and integration towards the silicon lines. *Nat Mater*, 2022, 21: 1225–1239
- 46 Wang S, Liu X, Zhou P. The road for 2D semiconductors in the silicon age. *Adv Mater*, 2022, 34: 2106886
- 47 Rush K, Charles Z, Garrett Z, et al. DrJAX: scalable and differentiable MapReduce primitives in JAX. 2024. [ArXiv:2403.07128](https://arxiv.org/abs/2403.07128)
- 48 Erman E, Furendal M. The democratization of global AI governance and the role of tech companies. *Nat Mach Intell*, 2024, 6: 246–248
- 49 Brooks R A. Intelligence without representation. *Artif Intell*, 1991, 47: 139–159
- 50 Driess D, Xia F, Sajjadi M S M, et al. PaLM-E: an embodied multimodal language model. 2023. [ArXiv:2303.03378](https://arxiv.org/abs/2303.03378)
- 51 Zhang Y, Tian J, Xiong Q. A review of embodied intelligence systems: a three-layer framework integrating multimodal perception, world modeling, and structured strategies. *Front Robot AI*, 2025, 12: 1668910
- 52 Long X, Zhao Q, Zhang K, et al. A survey: learning embodied intelligence from physical simulators and world models. 2025. [ArXiv:2507.00917](https://arxiv.org/abs/2507.00917)
- 53 Levine S, Finn C, Darrell T, et al. End-to-end training of deep visuomotor policies. 2015. [ArXiv:1504.00702](https://arxiv.org/abs/1504.00702)
- 54 Billard A, Kragic D. Trends and challenges in robot manipulation. *Science*, 2019, 364: 8414
- 55 Brohan A, Brown N, Carbajal J, et al. RT-2: vision-language-action models transfer web knowledge to robotic control. 2023. [ArXiv:2307.15818](https://arxiv.org/abs/2307.15818)
- 56 Yifan C, Wei M, Wang X, et al. Embodied AI: a survey on the evolution from perceptive to behavioral intelligence. *SmartBot*, 2025, 1: 70003
- 57 Kang L, Zhou H, Song X, et al. Advances and innovations in the multi-agent robotic system challenge. 2026. [ArXiv:2601.18733](https://arxiv.org/abs/2601.18733)
- 58 Metzger S L, Littlejohn K T, Silva A B, et al. A high-performance neuroprosthesis for speech decoding and avatar control. *Nature*, 2023, 620: 1037–1046
- 59 Willett F R, Kunz E M, Fan C, et al. A high-performance speech neuroprosthesis. *Nature*, 2023, 620: 1031–1036
- 60 Zaitsev V, Gvirsman O, Ben Hanan U, et al. A locust-inspired miniature jumping robot. *Bioinspir Biomim*, 2015, 10: 066012
- 61 Lisondra M, Benhabib B, Nejat G. Embodied AI with foundation models for mobile service robots: a systematic review. 2025. [ArXiv:2505.20503](https://arxiv.org/abs/2505.20503)
- 62 Lacour S P, Chan D, Wagner S, et al. Mechanisms of reversible stretchability of thin metal films on elastomeric substrates. *Appl Phys Lett*, 2006, 88: 204103
- 63 Lee G, Zarei M, Wei Q, et al. Surface wrinkling for flexible and stretchable sensors. *Small*, 2022, 18: 2203491
- 64 Luo Y, Abidian M R, Ahn J H, et al. Technology roadmap for flexible sensors. *ACS Nano*, 2023, 17: 5211–5295
- 65 Wang S, Li Z, Pei M, et al. Technology roadmap of bioinspired computing hardware. *ACS Nano*, 2026, 20: 8102–8163
- 66 Sun Y, Choi W M, Jiang H, et al. Controlled buckling of semiconductor nanoribbons for stretchable electronics. *Nat Nanotech*, 2006, 1: 201–207
- 67 Abouraddy A F, Bayindir M, Benoit G, et al. Towards multimaterial multifunctional fibres that see, hear, sense and communicate. *Nat Mater*, 2007, 6: 336–347
- 68 Liu J, Jiang S, Xiong W, et al. Self-healing kirigami assembly strategy for conformal electronics. *Adv Funct Mater*, 2021, 32: 2109214
- 69 Chen M, Liu J, Li P, et al. Fabric computing: concepts, opportunities, and challenges. *Innovation*, 2022, 3: 100340
- 70 Choi H W, Shin D W, Yang J, et al. Smart textile lighting/display system with multifunctional fibre devices for large scale smart home and IoT applications. *Nat Commun*, 2022, 13: 814

- 71 Shang R, Chen H, Cai X, et al. Machine learning-enhanced triboelectric sensing application. *Adv Mater Technol*, 2024, 9: 2301316
- 72 Ma X, Wang P, Huang L, et al. A monolithically integrated in-textile wristband for wireless epidermal biosensing. *Sci Adv*, 2023, 9: 2763
- 73 Wang Y, Yang Y, Yang N, et al. In-sensor wireless computing for intelligent remote sensing. *Nat Sens*, 2026, 1: 425–435
- 74 Xu C, Song Y, Han M, et al. Portable and wearable self-powered systems based on emerging energy harvesting technology. *Microsyst Nanoeng*, 2021, 7: 25
- 75 Xu C, Solomon S A, Gao W. Artificial intelligence-powered electronic skin. *Nat Mach Intell*, 2023, 5: 1344–1355
- 76 Tao L Q, Tian H, Liu Y, et al. An intelligent artificial throat with sound-sensing ability based on laser induced graphene. *Nat Commun*, 2017, 8: 14579
- 77 Wei Y, Qiao Y, Jiang G, et al. A wearable skinlike ultra-sensitive artificial graphene throat. *ACS Nano*, 2019, 13: 8639–8647
- 78 Yang Q, Jin W, Zhang Q, et al. Mixed-modality speech recognition and interaction using a wearable artificial throat. *Nat Mach Intell*, 2023, 5: 169–180
- 79 Li S, Jiang Q, Liu S, et al. A DNA nanorobot functions as a cancer therapeutic in response to a molecular trigger in vivo. *Nat Biotechnol*, 2018, 36: 258–264
- 80 Xing J, Yin T, Li S, et al. Sequential magneto-actuated and optics-triggered biomicrorobots for targeted cancer therapy. *Adv Funct Mater*, 2021, 31: 2008262
- 81 Gu H, Hanedan E, Boehler Q, et al. Artificial microtubules for rapid and collective transport of magnetic microcargoes. *Nat Mach Intell*, 2022, 4: 678–684
- 82 Tang X, Shen H, Zhao S, et al. Flexible brain-computer interfaces. *Nat Electron*, 2023, 6: 109–118
- 83 Hochberg L R, Bacher D, Jarosiewicz B, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 2012, 485: 372–375
- 84 Wang Z, Shi N, Zhang Y, et al. Conformal in-ear bioelectronics for visual and auditory brain-computer interfaces. *Nat Commun*, 2023, 14: 4213
- 85 Musk E. An integrated brain-machine interface platform with thousands of channels. *J Med Internet Res*, 2019, 21: 16194
- 86 Oxley T J, Opie N L, John S E, et al. Minimally invasive endovascular stent-electrode array for high-fidelity, chronic recordings of cortical neural activity. *Nat Biotechnol*, 2016, 34: 320–327
- 87 Choi Y S, Yin R T, Pfenniger A, et al. Fully implantable and bioresorbable cardiac pacemakers without leads or batteries. *Nat Biotechnol*, 2021, 39: 1228–1238
- 88 Li J, Dong Y, Park J H, et al. Body-coupled power transmission and energy harvesting. *Nat Electron*, 2021, 4: 530–538
- 89 Liu X, Wang Y, Wang G, et al. An ultrasound-driven implantable wireless energy harvesting system using a triboelectric transducer. *Matter*, 2022, 5: 4315–4331
- 90 Lu C, Jiang H, Cheng X, et al. High-performance fibre battery with polymer gel electrolyte. *Nature*, 2024, 629: 86–91
- 91 Merchant A, Batzner S, Schoenholz S S, et al. Scaling deep learning for materials discovery. *Nature*, 2023, 624: 80–85
- 92 Strubell E, Ganesh A, McCallum A. Energy and policy considerations for deep learning in NLP. 2019. ArXiv:1906.02243
- 93 Schwartz R, Dodge J, Smith N A, et al. Green AI. *Commun ACM*, 2020, 63: 54–63
- 94 Luccioni A S, Viguiers S, Ligozat A-L. Estimating the carbon footprint of BLOOM, a 176B parameter language model. 2022. ArXiv:2211.02001
- 95 Dodge J, Prewitt T, Tachet des Combes R, et al. Measuring the carbon intensity of AI in cloud instances. 2022. ArXiv:2206.05229
- 96 Perez E, Huang S, Song F, et al. Red teaming language models with language models. 2022. ArXiv:2202.03286
- 97 Bhardwaj R, Poria S. Red-teaming large language models using chain of utterances for safety alignment. 2023. ArXiv:2308.09662
- 98 Marcus G. The next decade in AI: four steps towards robust artificial intelligence. 2020. ArXiv:2002.06177
- 99 Srivastava A, Rastogi A, Rao A, et al. Beyond the imitation game: quantifying and extrapolating the capabilities of language models. 2022. ArXiv:2206.04615
- 100 Mitchell M, Wu S, Zaldivar A, et al. Model cards for model reporting. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019. 220–229
- 101 Bender E M, Gebru T, McMillan-Major A, et al. On the dangers of stochastic parrots: can language models be too big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021. 610–623
- 102 Sutton R S. The bitter lesson. 2019. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>
- 103 Frank M R, Autor D, Bessen J E, et al. Toward understanding the impact of artificial intelligence on labor. *Proc Natl Acad Sci USA*, 2019, 116: 6531–6539
- 104 Zaidan E S, Ibrahim I A. AI governance in a complex and rapidly changing regulatory landscape: a global perspective. *Humit Soc Sci Commun*, 2024, 11: 1121
- 105 Gasser U, Almeida V A F. A layered model for AI governance. *IEEE Internet Comput*, 2017, 21: 58–62
- 106 Tallberg J, Erman E, Furendal M, et al. The global governance of artificial intelligence: next steps for empirical and normative research. *Int Stud Rev*, 2023, 25: 040
- 107 Roberts H, Hine E, Taddeo M, et al. Global AI governance: barriers and pathways forward. *Int Aff*, 2024, 100: 1275–1286
- 108 Tabassi E. Artificial Intelligence Risk Management Framework. *AI RMF* 1.0. 2024
- 109 Bommasani R, Arora S, Chayes J, et al. Advancing science- and evidence-based AI policy. *Science*, 2025, 389: 459–461
- 110 Yun H. China's data sovereignty and security: implications for global digital borders and governance. *Chin Polit Sci Rev*, 2025, 10: 178–203
- 111 UNESCO. Recommendation on the ethics of artificial intelligence. 2021
- 112 World Health Organization. Ethics and governance of artificial intelligence for health. 2021
- 113 Davidsen Buhl M, et al. Safety cases for frontier AI. 2024. ArXiv:2410.21572
- 114 Falco G, Shneiderman B, Badger J, et al. Governing AI safety through independent audits. *Nat Mach Intell*, 2021, 3: 566–571
- 115 Raji I D, Xu P, Honigsberg C, et al. Outsider oversight: designing a third party audit ecosystem for AI governance. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2022. 557–571
- 116 Gebru T, Morgenstern J, Vecchione B, et al. Datasheets for datasets. *Commun ACM*, 2021, 64: 86–92
- 117 Ian M, Malone M J, Schwindt K, et al. Governance approaches to securing frontier AI. 2025
- 118 McGregor S. Preventing repeated real world AI failures by cataloging incidents: the AI incident database. 2020. ArXiv:2011.08512
- 119 Liang P, Bommasani R, Lee T, et al. Holistic evaluation of language models. 2022. ArXiv:2211.09110
- 120 ISO. Robots and Robotic Devices—Collaborative Robots. *ISO/TS 15066*: 2016
- 121 ISO. Information Technology—Artificial Intelligence—Management System. *ISO/IEC 42001*: 2023
- 122 Krauss J K, Lipsman N, Aziz T, et al. Technology of deep brain stimulation: current status and future directions. *Nat Rev Neurol*, 2021, 17: 75–87
- 123 Orsborn A L, Moorman H G, Overduin S A, et al. Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, 2014, 82: 1380–1393
- 124 Luckcuck M, Farrell M, Dennis L A, et al. Formal specification and verification of autonomous robotic systems: a survey. *ACM Comput Surv*, 2019, 52: 100
- 125 Law H, Ho L. Can a dual mandate be a model for the global governance of AI? *Nat Rev Phys*, 2023, 5: 706–707