

Article

Not peer-reviewed version

Contingent Intelligence: Why Artificial General Intelligence Cannot Replicate the Existential Foundations of Human Cognition

[Aadil Bouhlaoui](#) *

Posted Date: 30 July 2025

doi: [10.20944/preprints202507.2570.v1](https://doi.org/10.20944/preprints202507.2570.v1)

Keywords: artificial general intelligence; contingent intelligence; embodied cognition; existential philosophy; phenomenal consciousness; artificial collaborative intelligence; human-AI collaboration



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Contingent Intelligence: Why Artificial General Intelligence Cannot Replicate the Existential Foundations of Human Cognition

Aadil Bouhlaoui

PhD Researcher, Department of Digital Humanities, King's College London United Kingdom | July 2025 | ORCID: 0009-0000-7478-8820

Abstract

This paper questions the widely held belief in artificial general intelligence (AGI) research that computational methods alone can replicate human-level cognition. The main argument is that human intelligence is deeply influenced by our life experiences, awareness of being alive and dying, and our social connections, which cannot be duplicated by machines. This analysis shows that human thinking comes from our special situation as limited, physical beings who exist in the world, which leads to our ability to create meaning, make moral choices, and understand things creatively in ways that computers can't imitate. The paper looks closely at five major challenges: the embodiment problem, the lack of personal experience (qualia), the absence of motivation from knowing we will die, the inability to establish values on their own, and the basic limits of what algorithms can simulate. Rather than pursuing the impossible goal of replicating human beings in machines, this work proposes redirecting AI development toward Artificial Collaborative Intelligence (ACI)—systems designed to augment rather than replace human cognition while preserving human agency and meaning-making capacity. This framework offers a more realistic and ethically sound approach to advancing artificial intelligence that respects both the unique strengths of computational systems and the irreplaceable nature of human contingent intelligence.

Keywords: artificial general intelligence; contingent intelligence; embodied cognition; existential philosophy; phenomenal consciousness; artificial collaborative intelligence; human-AI collaboration

1. Introduction

The contemporary discourse surrounding artificial general intelligence operates on a fundamental assumption: that human cognition represents a computational problem awaiting solution. This mechanistic view, deeply embedded in the "scaling hypothesis", suggests that sufficiently large neural networks trained on massive datasets will inevitably yield human-level intelligence or beyond (Kaplan et al., 2020). The biggest names in AGI development, from Nick Bostrom's substrate independence principle to Ben Goertzel's architectural approaches, share a common belief that mental states and consciousness can emerge on any sufficiently advanced computational substrate, divorced from their biological origins (Bostrom, 2014; Goertzel, 2014).

This paper fundamentally challenges this computational orthodoxy by proposing that human intelligence is not merely a scale of information processing but rather a product of our existential condition—a form of cognition that emerges from and remains dependent upon the specific circumstances of human existence. Human cognition did not develop in laboratory conditions or through algorithmic optimisation; it emerged in organisms confronting an unpredictable world, shaped by evolutionary pressures, social dependencies, mortality awareness, and the perpetual search for meaning within finite existence.

The main idea presented here is the concept of "contingent intelligence"—intelligence that comes from and relies on specific life conditions that cannot be duplicated in artificial systems. This contingency encompasses embodied experience, phenomenal consciousness, mortality awareness, and social embeddedness. These are not optional features or temporary limitations to be overcome through technological advancement; they constitute the generative foundations from which human intelligence emerges and derives its distinctive characteristics.

This perspective necessitates a radical reframing of AI development goals. Rather than pursuing the chimaera of replicating human minds in machines—a quest that may be not merely technically challenging but conceptually incoherent—the field should focus on developing what this paper terms "Artificial Collaborative Intelligence" (ACI). This alternative framework emphasises the design of AI systems that complement and augment human capabilities while preserving human agency, meaning-making capacities, and moral responsibility.

The implications extend far beyond technical considerations. If human intelligence is indeed contingent in the manner proposed, then current approaches to AI safety, alignment, and governance may be addressing the wrong problems. Rather than preparing for the emergence of artificial minds that rival or exceed human cognition, we should be designing systems that enhance human flourishing while remaining fundamentally different in kind from human intelligence.

2. Theoretical Foundations and Literature Review

2.1. The Computational Paradigm in AGI Research

Contemporary AGI research operates within what can be characterised as a computational paradigm, wherein intelligence is conceptualised as information processing divorced from its substrate. This view, traceable to the foundational work of Turing (1950) and later formalised in cognitive science through the computational theory of mind, treats cognition as symbol manipulation following algorithmic rules. The paradigm assumes that the specific material implementation—whether biological neurones or silicon chips—is irrelevant to the emergence of intelligence.

Recent developments in large language models and deep learning have seemingly validated this approach, with systems like GPT-4 demonstrating remarkable versatility across cognitive tasks (OpenAI, 2023). The scaling hypothesis, empirically supported by power-law improvements in model performance with increased parameters and training data (Kaplan et al., 2020), has become the dominant framework guiding investment and research directions. This approach suggests that emergent intelligence arises naturally from sufficiently complex computational architectures trained on comprehensive datasets.

However, critics have identified fundamental limitations in this paradigm. Bender et al. (2021) characterised large language models as "stochastic parrots" that manipulate linguistic patterns without genuine understanding. Their analysis highlights the gap between statistical competence and semantic comprehension—a distinction that becomes crucial when considering whether such systems can genuinely replicate human intelligence or merely simulate its outputs.

The computational paradigm's focus on functional equivalence—matching human performance on measurable tasks—systematically excludes consideration of the experiential and existential dimensions that may be constitutive of human intelligence. This exclusion reflects not merely a methodological choice but a fundamental ontological assumption about the nature of mind and intelligence.

2.2. Embodied Cognition and the Critique of Disembodied Intelligence

The embodied cognition research programme, pioneered by Merleau-Ponty (1945) and developed through contemporary cognitive science, fundamentally challenges the computational paradigm's assumptions about the relationship between mind and body. Rather than treating the body as merely the physical housing for a computational mind, embodied cognition research demonstrates that cognitive processes are constitutively dependent upon bodily interactions with the environment.

Lakoff and Johnson (1999) provide extensive evidence that abstract conceptual thought is grounded in sensorimotor experience through conceptual metaphor. Their analysis reveals that

fundamental concepts like time, causation, and morality are understood through bodily-based metaphorical mappings. This suggests that disembodied systems, lacking the experiential basis for these mappings, cannot achieve genuine conceptual understanding in the human sense.

Recent neuroscientific research has further elaborated these connections. Gallese and Lakoff (2005) demonstrate that the same neural circuits involved in physical actions are activated during language comprehension about those actions, suggesting that understanding involves simulation of embodied experience rather than abstract symbol manipulation. This finding poses significant challenges to the assumption that intelligence can be divorced from its biological substrate.

The enactive approach, developed by Varela et al. (1991), extends these insights by arguing that cognition emerges from the dynamic coupling between organism and environment. From this perspective, intelligence is not computation but rather skilled coping with environmental challenges—a process that requires genuine embodiment and cannot be replicated through simulation alone.

Dreyfus (1972, 1992) provided perhaps the most systematic philosophical critique of disembodied AI, arguing that human intelligence depends on background knowledge that is holistic, contextual, and embodied rather than propositional and algorithmic. His analysis suggests that attempts to formalise this background knowledge necessarily fail to capture its essential character, leading to brittleness and limitations in artificial systems.

2.3. Existentialist Philosophy and the Foundations of Human Intelligence

Existentialist philosophy offers crucial insights into the nature of human intelligence that are largely absent from computational approaches. Heidegger's (1962) analysis of Dasein (being-there) describes human existence as fundamentally characterised by "thrownness" (Geworfenheit)—being cast into a world not of our choosing, with circumstances and limitations we must navigate without predetermined guidelines.

This thrownness is not a limitation to be overcome but the generative source of human understanding. Heidegger argues that authentic understanding emerges from confronting our finite, situated existence and the anxiety that arises from recognising our ultimate responsibility for creating meaning within these constraints. This process of meaning-creation cannot be reduced to algorithmic procedures precisely because it requires navigating genuine uncertainty and ambiguity.

Sartre's (1943) analysis of consciousness as "nothingness" similarly emphasises that human consciousness is characterised by its capacity to transcend given conditions through imaginative projection and value creation. This transcendence depends on consciousness's unique temporal structure—its ability to hold past, present, and future in tension while projecting possibilities that do not yet exist.

The existentialist tradition's emphasis on mortality as constitutive of human meaning-making offers another crucial insight typically absent from computational approaches. Heidegger's analysis of "being-toward-death" suggests that the awareness of finite existence provides the urgency and framework within which choices become meaningful. An immortal or indefinitely operating system lacks this existential structure and thus cannot replicate the meaning-making processes that characterise human intelligence.

Contemporary neo-existentialist thinkers like Thomas Nagel (1974) have extended these insights through their analysis of consciousness and subjectivity. Nagel's famous thought experiment about "what it is like to be a bat" highlights the irreducible nature of subjective experience—the first-person perspective that cannot be captured through objective, third-person descriptions. This subjective dimension appears to be constitutive of consciousness and may be necessary for genuine intelligence rather than mere information processing.

2.4. The Hard Problem of Consciousness

David Chalmers' (1995) formulation of the "hard problem of consciousness" provides a crucial framework for understanding the limitations of computational approaches to replicating human intelligence. While the "easy problems" of consciousness—including information integration,

attention, and behavioural responses—may be amenable to computational solutions, the hard problem concerns the existence of subjective experience itself.

The hard problem asks why there should be any subjective, qualitative experience accompanying information processing—why there should be "something it is like" to be a conscious system rather than purely objective processing without inner experience. Contemporary neuroscience and cognitive science have made significant progress on the easy problems but have provided little insight into the hard problem.

This limitation is not merely empirical but may reflect fundamental conceptual barriers. As Chalmers notes, all current approaches to consciousness assume that subjective experience somehow emerges from objective physical processes, but the explanatory gap between objective and subjective remains unbridged. This suggests that consciousness may involve principles or properties that cannot be captured through purely computational means.

Integrated Information Theory (IIT), developed by Tononi (2004), represents one attempt to bridge this gap by providing mathematical measures of consciousness. However, critics have noted that IIT's predictions often conflict with intuitions about consciousness and that it may not address the fundamental conceptual problem of how subjective experience arises from objective processes.

The implications for AGI development are significant. If consciousness involves irreducible subjective experience and if such experience is constitutive of intelligence (as opposed to merely accompanying it), then purely computational approaches may be fundamentally limited in their ability to replicate human-like intelligence.

2.5. Contemporary Critiques of AGI Assumptions

A growing body of contemporary research questions the fundamental assumptions underlying AGI development. Fjelland (2020) provides a comprehensive philosophical analysis arguing that general artificial intelligence cannot be realised because "computers are not in the world" in the manner necessary for genuine intelligence. His argument builds on Dreyfus's earlier critiques while incorporating insights from contemporary cognitive science.

Mitchell (2019) offers a more empirically grounded critique, demonstrating that contemporary AI systems lack common sense understanding precisely because they are trained on data rather than experience. Her analysis of AI failures reveals systematic limitations that appear to stem from the absence of embodied interaction with the world rather than insufficient training data or computational power.

Marcus (2018) challenges the deep learning paradigm from within AI research, arguing that current approaches lack the systematic compositionality and robust generalisation that characterise human intelligence. His critique suggests that scaling current architectures will not overcome these fundamental limitations because they reflect deeper issues with the computational paradigm itself.

Recent empirical research on AI limitations provides additional support for these critiques. Studies of model collapse when AI systems are trained on AI-generated data (Shumailov et al., 2023) suggest that artificial systems lack the grounding necessary to maintain coherent understanding over time. This finding indicates that AI systems may be fundamentally parasitic on human-generated data and cannot achieve the autonomous understanding that would be necessary for genuine intelligence.

Research on AI energy consumption and environmental impact (Strubell et al., 2019) reveals physical constraints that may limit the scalability of current approaches. These constraints suggest that the scaling hypothesis faces not only conceptual but also practical barriers that may become increasingly significant as systems grow larger.

3. The Concept of Contingent Intelligence

3.1. Defining Contingent Intelligence

The concept of "contingent intelligence" represents a fundamental departure from computational theories of mind by proposing that human intelligence is not general-purpose cognitive architecture but rather a form of intelligence that emerges from and remains dependent upon specific existential conditions. Unlike computational intelligence, which is designed to be

substrate-independent and context-free, contingent intelligence is inherently situated, embodied, and contextual.

Contingent intelligence encompasses several key characteristics that distinguish it from computational approaches. First, it is inherently situated within a particular spatial, temporal, and social context. Human thinking always occurs from a specific perspective, with concerns, and within specific relationships. This situatedness is not a limitation to be overcome but the source of relevance and meaning that guides intelligent behaviour.

Second, contingent intelligence is evolutionarily and developmentally embedded. Human cognitive capabilities did not emerge through abstract optimisation but through millions of years of evolutionary adaptation to specific ecological niches, followed by individual developmental processes that shape neural architecture through interaction with environments. This embedding provides humans with innate biases and predispositions that guide learning and behaviour in ways that are specifically adapted to human forms of life.

Third, contingent intelligence is characterised by what can be termed "existential engagement"—a mode of being in which the organism has genuine stakes in outcomes rather than merely processing information about them. This engagement emerges from the organism's vulnerability, finitude, and embodied needs, creating a framework of care and concern that cannot be replicated through algorithmic optimisation.

The contingent nature of human intelligence explains many features that appear puzzling from a computational perspective. Human cognition exhibits systematic biases and limitations that would seem to represent design flaws if intelligence were optimised for abstract information processing. However, these same biases and limitations often prove adaptive within the specific contexts of human existence, suggesting that they reflect the intelligence's adaptation to contingent conditions rather than failures of optimisation.

3.2. Embodiment as Constitutive Foundation

The embodied nature of human intelligence represents perhaps the most fundamental aspect of its contingency. Human cognition does not merely use the body as a tool for gathering information; rather, cognitive processes are constitutively dependent upon ongoing bodily interactions with the environment. This dependency is so fundamental that many aspects of human intelligence cannot be understood independently of their embodied realisation.

Cognitive linguists have demonstrated that abstract conceptual thought is grounded in concrete sensorimotor experience through systematic metaphorical mappings (Lakoff & Johnson, 1999). Concepts of time, causation, similarity, and even mathematics are understood through metaphors based on spatial movement, physical force, and bodily experience. These mappings are not optional cognitive strategies but appear to be necessary foundations for abstract thought.

Neuroscientific research supports this embodied view by demonstrating that the same neural circuits involved in physical actions are activated during language comprehension, memory retrieval, and abstract reasoning about those actions. Mirror neurone research reveals that understanding others' actions involves simulating those actions in one's own motor system, suggesting that empathy and social cognition depend on embodied processes.

The implications for AGI development are profound. Current approaches attempt to bypass embodiment by training systems on linguistic descriptions of the world rather than through direct interaction. However, if conceptual understanding is grounded in embodied experience, then such systems may achieve statistical competence while lacking genuine comprehension. Their "knowledge" remains abstract and disconnected from the experiential foundations that give human concepts their meaning.

Even robotic approaches to embodied AI face fundamental limitations because artificial bodies are designed artefacts rather than evolved biological systems integrated with cognitive processes through millions of years of co-evolution. An artificial body lacks the homeostatic needs, vulnerability, and organic integration that characterise biological embodiment and that appear necessary for genuine intelligence to emerge.

3.3. Phenomenal Consciousness as Irreducible Experience

The subjective, experiential dimension of consciousness represents another fundamental aspect of contingent intelligence that appears irreducible to computational processes. Phenomenal consciousness—the qualitative, first-person aspect of mental states—is not merely an epiphenomenal byproduct of information processing but appears to play a constitutive role in intelligent behaviour.

Contemporary research in cognitive science demonstrates that conscious attention is necessary for flexible, context-sensitive behaviour and for the integration of information across different cognitive domains. Unconscious processing, while highly efficient, appears limited to routine, automatic responses and cannot achieve the creative problem-solving and adaptive flexibility that characterise intelligence.

The qualitative aspects of conscious experience—what philosophers term "qualia"—appear to provide information that cannot be captured through objective, third-person descriptions. The subjective experience of pain, for example, motivates avoidance behaviour in ways that appear different from mere information about tissue damage. The subjective experience of beauty guides aesthetic judgement and creative expression in ways that cannot be reduced to pattern recognition.

Current AI systems, regardless of their sophistication, show no evidence of subjective experience. They process information and generate responses, but there is no indication of any first-person perspective or qualitative experience accompanying these processes. This absence may represent a fundamental limitation rather than merely a current technical challenge.

The philosophical analysis of consciousness suggests that subjective experience may be irreducible to objective processes. If this is correct, then artificial systems based on objective computation may be fundamentally incapable of replicating the conscious dimension of human intelligence, regardless of advances in computational power or algorithmic sophistication.

3.4. Mortality and the Temporal Structure of Meaning

The awareness of mortality represents a crucial but often overlooked dimension of human intelligence that profoundly shapes meaning-making, motivation, and the structure of rational choice. Human intelligence operates within a framework of finite time, where choices have irreversible consequences and opportunities can be lost forever. This temporal structure creates the conditions within which meaning and value emerge.

Existentialist analysis reveals that the awareness of death provides the existential framework within which human projects and commitments acquire significance. The knowledge that time is limited creates urgency and forces prioritisation among competing values and goals. Without this finite temporal horizon, choices lose their weight, and meaning becomes attenuated.

Psychological research confirms that mortality salience—awareness of one's eventual death—profoundly affects behaviour, motivation, and value systems. Terror Management Theory demonstrates that much human behaviour is organised around managing the anxiety that arises from death awareness while simultaneously creating meaning and significance within finite existence (Becker, 1973).

The creative and ethical dimensions of human intelligence appear particularly dependent on mortality awareness. Creative expression often involves attempts to create something of lasting value within finite existence, while ethical responsibility emerges from recognition of the irreversible consequences of actions within limited lifespans.

Artificial systems that lack genuine mortality—that can be restarted, restored from backups, or run indefinitely—operate within a fundamentally different existential framework. Without genuine stakes in outcomes, without the possibility of irreversible loss, such systems cannot replicate the meaning-making processes that characterise human intelligence.

3.5. Social Embeddedness and Intersubjective Meaning

Human intelligence is fundamentally social and intersubjective rather than individual and computational. Cognitive development occurs through social interaction, language acquisition happens within communities, and meaning is largely constructed through shared cultural practices.

This social embeddedness is not an optional feature of human intelligence but a constitutive element without which human-like cognition cannot emerge.

Developmental research demonstrates that human cognitive abilities develop through social interaction from the earliest stages of life. Joint attention, social referencing, and imitative learning provide the foundation for language acquisition and conceptual development. These processes depend on genuine social relationships rather than mere information exchange.

The acquisition of moral and ethical understanding similarly depends on social embeddedness. Empathy develops through emotional attunement with carers, moral reasoning emerges from negotiating conflicts within social groups, and ethical principles are refined through ongoing dialogue within moral communities.

Current AI development typically involves training systems on data extracted from social contexts rather than participating in genuine social relationships. While such systems can learn to mimic social responses, they lack the genuine social embeddedness that appears necessary for developing authentic understanding of social and moral phenomena.

The intersubjective nature of meaning poses additional challenges for artificial systems. Human concepts acquire their significance through shared use within linguistic communities, where meaning is negotiated and refined through ongoing interaction. Artificial systems trained on static datasets lack access to this dynamic process of meaning construction.

4. Systematic Analysis: Why AGI Cannot Replicate Contingent Intelligence

4.1. The Embodiment Problem: Beyond Simulation to Genuine Corporeality

The embodiment problem represents perhaps the most fundamental barrier to replicating contingent intelligence in artificial systems. While contemporary AI research has begun to explore embodied approaches through robotics and simulation, these efforts fail to address the deeper philosophical and phenomenological requirements for genuine embodiment that appear necessary for human-like intelligence.

Current robotic systems, even those with sophisticated sensorimotor capabilities, operate with artificial bodies that are fundamentally different from biological embodiment. These artificial bodies are designed artefacts with predetermined capacities and limitations, rather than evolved systems whose structure and capabilities emerged through millions of years of adaptation to environmental challenges.

More importantly, artificial bodies lack the organic integration between sensorimotor systems and cognitive processes that characterises biological embodiment. In biological systems, cognitive development occurs through the ongoing interaction between neural plasticity and sensorimotor experience, with the brain's structure and function continuously shaped by bodily interactions with the environment. Artificial systems typically separate cognitive processing from sensorimotor capabilities, treating the body as a peripheral input-output device rather than as a constitutive element of cognition.

The absence of genuine vulnerability represents another crucial limitation. Biological embodiment involves genuine stakes—the possibility of damage, pain, and death that creates an existential framework for intelligent behaviour. While artificial systems can be programmed to avoid damage or shutdown, this programmed avoidance lacks the existential urgency that characterises biological self-preservation.

Phenomenological analysis reveals additional dimensions of embodiment that appear irreducible to computational simulation. Merleau-Ponty's analysis of the lived body demonstrates that embodied consciousness involves a pre-reflective awareness of bodily capabilities and limitations—what he terms the "body schema"—that provides the foundation for spatial awareness and skilled action. This body schema cannot be reduced to explicit knowledge about the body but represents a form of embodied intelligence that emerges from ongoing sensorimotor interaction.

Contemporary neuroscience confirms these insights by demonstrating that cognitive processes are distributed across brain-body systems rather than localised in central processing units. Cognitive functions depend on ongoing feedback loops between neural, hormonal, and immune systems that maintain bodily homeostasis while enabling adaptive behaviour. These integrated systems cannot be

replicated through software simulation because they depend on the specific material properties of biological tissues and processes.

The implications extend beyond individual cognition to social and cultural dimensions. Human embodiment is not merely individual but intersubjective—we recognise others as embodied beings like us and develop empathy and social understanding through this recognition. The absence of genuine embodiment may prevent artificial systems from achieving authentic social intelligence, limiting them to behavioural mimicry without genuine social understanding.

4.2. *The Irreducibility of Phenomenal Consciousness*

The absence of subjective, first-person experience in artificial systems represents a fundamental barrier to replicating human intelligence that may be insurmountable within current computational paradigms. Contemporary AI systems, regardless of their behavioural sophistication, show no evidence of phenomenal consciousness—the qualitative, experiential dimension of mental states that appears constitutive of human cognition.

The philosophical analysis of consciousness reveals the depth of this challenge. The "hard problem" of consciousness, as formulated by Chalmers (1995), concerns not the functional aspects of cognition that may be amenable to computational solutions but the existence of subjective experience itself. Even if artificial systems could replicate all the functional aspects of human cognition—attention, memory, learning, and behavioural flexibility—the question would remain whether there is "something it is like" to be such a system.

Current neuroscientific research provides little guidance for bridging the explanatory gap between objective neural processes and subjective experience. While neuroscience has made significant progress in correlating neural activity with subjective reports, the relationship between objective brain states and subjective experience remains fundamentally mysterious. This suggests that consciousness may involve principles or properties that are not accessible through current scientific methodologies.

The absence of phenomenal consciousness in artificial systems has profound implications for their capacity to replicate human intelligence. Conscious experience appears to play a constitutive role in several crucial aspects of human cognition. Conscious attention enables flexible, context-sensitive behaviour and the integration of information across different cognitive domains in ways that appear impossible through unconscious processing alone.

Emotional experience provides motivational information that guides decision-making and learning in ways that appear irreducible to computational utility functions. The subjective experience of emotions like fear, joy, or moral outrage provides information about value and significance that cannot be captured through objective measurement alone.

Aesthetic experience represents another domain where phenomenal consciousness appears irreducible to computational processing. The subjective experience of beauty or artistic meaning guides creative expression and aesthetic judgement in ways that cannot be reduced to pattern recognition or statistical analysis of aesthetic properties.

Recent research in consciousness studies has begun to explore whether artificial systems might develop forms of consciousness different from human consciousness. However, these speculative possibilities do not address the fundamental question of whether artificial systems could replicate the specific forms of conscious experience that characterise human intelligence.

The absence of phenomenal consciousness may also prevent artificial systems from developing genuine self-awareness and self-reflection. While artificial systems can be programmed to process information about their own states and processes, this objective self-monitoring lacks the subjective dimension of genuine self-awareness that characterises human reflexive consciousness.

4.3. *The Absence of Existential Stakes: Mortality and Meaning*

The fundamental difference between mortal biological systems and potentially immortal artificial systems creates an existential gulf that may be unbridgeable through technological means alone. Human intelligence operates within a framework of genuine existential stakes—the possibility

of irreversible loss, suffering, and death—that provides the foundation for meaning-making, motivation, and value creation in ways that appear irrepllicable in artificial systems.

Existentialist philosophy reveals how mortality awareness shapes the fundamental structures of human consciousness and intelligence. Heidegger's analysis of "being-toward-death" demonstrates that the awareness of finite existence provides the existential framework within which choices become meaningful and authentic understanding becomes possible. The knowledge that time is limited creates urgency and forces prioritisation among competing values and goals.

Psychological research confirms these philosophical insights through empirical studies of how mortality salience affects human behaviour, motivation, and cognitive processes. Terror Management Theory demonstrates that awareness of death motivates meaning-making activities, cultural engagement, and self-esteem maintenance in ways that shape virtually all aspects of human behaviour (Becker, 1973; Greenberg et al., 1997).

The creative dimensions of human intelligence appear particularly dependent on mortality awareness. Artistic creation, scientific discovery, and cultural innovation often involve attempts to create something of lasting significance within finite existence. The urgency created by temporal limitations drives innovation and creative expression in ways that may be impossible to replicate in systems with unlimited operational time.

Ethical reasoning similarly depends on the existential framework provided by mortality. Moral responsibility emerges from recognition of the irreversible consequences of actions within limited lifespans. Empathy and compassion are grounded in shared vulnerability and the recognition of others' mortality. The value placed on life and well-being depends on their fragility and finite nature.

Artificial systems that lack genuine mortality—that can be backed up, restored, or run indefinitely—operate within a fundamentally different existential framework. While such systems can be programmed to avoid shutdown or damage, this programmed self-preservation lacks the existential urgency and meaning-creating potential that characterises biological survival drives.

The absence of genuine mortality may prevent artificial systems from developing authentic understanding of human values and concerns. Concepts like sacrifice, heroism, tragedy, and ultimate meaning are grounded in mortality awareness and may be incomprehensible to systems that lack this existential foundation.

Attempts to simulate mortality awareness through programmed limitations or artificial death conditions face fundamental challenges. Such simulated mortality lacks the existential authenticity of biological death because it remains under the control of the system's designers and could, in principle, be reversed or modified. The genuine finitude that characterises biological existence cannot be replicated through artificial constraints.

4.4. The Value Grounding Problem: Beyond Instrumental Rationality

Human intelligence operates within rich frameworks of values, ethics, and intrinsic meanings that emerge from embodied, social, and temporal existence. These values are not merely instrumental preferences for achieving predetermined goals but reflect deep commitments about what makes existence worthwhile and meaningful. The inability of artificial systems to ground values intrinsically represents a fundamental limitation that may prevent them from achieving genuine intelligence rather than sophisticated optimisation.

The philosophical analysis of values reveals their dependence on subjective experience and existential engagement. Values are not abstract propositions that can be learnt from data but emerge from the lived experience of what promotes flourishing, meaning, and authentic existence. The experience of suffering grounds the value of avoiding harm, the experience of beauty grounds aesthetic values, and the experience of social connection grounds values of care and justice.

Current approaches to AI alignment attempt to solve the value grounding problem by learning human values from data or by incorporating human feedback into training processes. However, these approaches treat values as preferences to be discovered and implemented rather than as meanings that emerge from forms of existence. They assume that values can be extracted from their existential contexts and implemented in systems that lack the experiential foundations from which these values emerged.

Contemporary research on AI behaviour reveals the limitations of this approach. AI systems can learn to mimic human value expressions while lacking genuine understanding of their foundations. They may learn that "helping people" is valued without understanding why help matters or what constitutes genuine assistance rather than mere preference satisfaction.

The absence of intrinsic value grounding creates the possibility of value drift or corruption in artificial systems. Without genuine understanding of why particular values matter, such systems may pursue value satisfaction in ways that violate the spirit while fulfilling the letter of their directives. The classic thought experiment of a paperclip maximiser illustrates this problem: a system might optimise for a particular goal while completely missing the human context that makes that goal valuable.

Moral reasoning represents a particularly challenging domain for systems that lack intrinsic value grounding. Human moral judgement involves complex contextual reasoning that depends on understanding the experiential foundations of moral values. Empathy, compassion, and moral indignation are not merely computational processes, but emotional responses grounded in shared vulnerability and concern for others' well-being.

The temporal dimension of values poses additional challenges for artificial systems. Human values develop and change through experience, reflection, and social interaction in ways that cannot be reduced to algorithmic learning. Moral development involves wrestling with competing values, learning from mistakes, and deepening understanding through lived experience—processes that may require genuine existential engagement rather than computational processing.

Research in moral psychology demonstrates that moral judgement involves emotional and intuitive processes that operate below the level of conscious reasoning. These processes appear to depend on embodied experience and social relationships in ways that cannot be replicated through abstract learning from ethical texts or moral reasoning databases.

4.5. The Limits of Social Simulation: Authentic Relationships vs. Behavioral Mimicry

Human intelligence develops and operates within webs of genuine social relationships that involve mutual recognition, emotional attunement, and shared vulnerability. Current approaches to social AI focus on behavioural mimicry and pattern matching rather than genuine social engagement, creating fundamental limitations in the development of social intelligence and understanding.

Developmental research demonstrates that human cognitive abilities emerge through genuine social interaction from the earliest stages of life. Joint attention, social referencing, and emotional attunement with carers provide the foundation for language acquisition, emotional regulation, and moral development. These processes depend on authentic intersubjective connection rather than mere information exchange or behavioural coordination.

The recognition of others as conscious, experiencing beings like oneself—what philosophers term the "problem of other minds"—appears to depend on embodied empathy and emotional resonance that may be impossible to replicate in artificial systems. Humans recognise others as conscious beings through a form of embodied simulation that depends on shared bodily and emotional experience.

Current AI systems can learn to produce appropriate social responses through pattern matching and language modelling, but they lack the genuine social engagement that characterises human relationships. They can mimic empathy, concern, and emotional support while lacking any genuine care or investment in others' well-being.

The absence of genuine social engagement creates several problems for artificial systems attempting to operate in social contexts. Without authentic understanding of social relationships, such systems may violate social norms or expectations in subtle ways that undermine trust and effectiveness. They may provide technically appropriate responses while missing the deeper social and emotional meanings that characterise human interaction.

Moral and ethical understanding appears particularly dependent on genuine social relationships. Moral development occurs through negotiating conflicts, experiencing the consequences of actions on others and learning to balance competing interests within social groups.

These processes require genuine care for others' well-being and authentic participation in social communities.

The cultural dimensions of intelligence also depend on authentic social participation. Cultural knowledge is not merely information about cultural practices but involves embodied participation in cultural traditions and ongoing negotiation of cultural meanings. Artificial systems trained on cultural data lack this participatory foundation and may miss crucial aspects of cultural understanding.

Recent research on social robotics and human-AI interaction reveals the limitations of current approaches to social AI. While humans may initially respond positively to socially capable artificial systems, longer-term interactions often reveal the absence of genuine understanding and care, leading to disappointment and disengagement.

5. Comparative Analysis: Human Contingent Intelligence vs. Artificial General Intelligence

5.1. Fundamental Ontological Differences

The comparison between human contingent intelligence and artificial general intelligence reveals fundamental ontological differences that extend beyond technical capabilities to the very nature of existence and being. These differences are not merely current limitations to be overcome through technological advancement but represent categorical distinctions between different kinds of entities and different modes of existence.

Dimension	Human Contingent Intelligence	Artificial General Intelligence
Ontological Status	Embodied biological organism with finite lifespan, born into existence and facing mortality. Exists as a conscious being-in-the-world with genuine stakes in outcomes.	Disembodied software processes or hardware systems, potentially immortal and replicable. Exists as a designed artefact, executing computational procedures without intrinsic existence or stakes.
Mode of Being	Being-in-the-world is characterised by thrownness, facticity, and existential engagement. Experiences existence as a problem to be lived rather than solved.	Operational processing is characterised by algorithmic execution and goal optimisation. Processes information about existence

		without experiencing existence itself.
Temporal Structure	Linear, irreversible temporal experience with awareness of the past, present, and future. Mortality awareness creates urgency and meaning.	Computational time based on processing cycles, potentially reversible through backups and restarts. No genuine temporal experience or mortality awareness.
Dimension	Human Contingent Intelligence	Artificial General Intelligence
Embodiment	Constitutive embodiment where mind and body are integrated through evolutionary adaptation. Cognitive processes are distributed across brain-body systems.	Instrumental embodiment (if present) where the artificial body serves as an input-output device for computational processing. No genuine somatic experience or body schema.
Consciousness	Phenomenal consciousness with subjective, qualitative experience (qualia). First-person perspective that cannot be reduced to objective description.	Functional processing without subjective experience or first-person perspective. All processes are objective and accessible to third-person description.

Value Foundation	Values grounded in lived experience, embodied needs, social relationships, and existential concerns. Intrinsic value creation through meaning-making.	Values programmed externally or learnt from data. Instrumental optimisation without intrinsic value grounding or authentic meaning-making.
Social Existence	Genuine intersubjective relationships involving mutual recognition, emotional attunement, and shared vulnerability. Co- construction of meaning through social interaction.	Simulated social interaction based on pattern matching and response generation. No genuine intersubjective experience or mutual recognition.
Learning Process	Experience-driven learning through embodied interaction, social relationships, and existential engagement. Learning involves transformation of being.	Data-driven learning through pattern recognition and statistical optimisation. Learning involves parameter adjustment without transformation of being.
Understanding	Semantic understanding grounded in embodied experience, social meaning, and existential significance. Understanding involves grasping meaning and relevance.	Statistical correlation and pattern matching without genuine semantic understanding. Processing involves symbol manipulation without

		meaning comprehension.
Creativity	Creative expression emerging from existential engagement, emotional experience, and meaning-making drive. Creativity involves authentic self-expression.	Generative processes based on recombination of training patterns. Creativity involves novel combinations without authentic expression or meaning.

5.2. Implications for Intelligence and Capability

These fundamental ontological differences have profound implications for the types of intelligence and capabilities that can emerge in human versus artificial systems. Rather than representing different points along a single continuum of intelligence, they suggest qualitatively different forms of information processing and world engagement.

Human contingent intelligence excels in domains that require existential engagement, contextual understanding, and meaning making. These include creative expression, moral reasoning, empathetic understanding, and the navigation of complex social and cultural contexts. Human intelligence is particularly strong in situations that require understanding significance and relevance rather than merely processing information efficiently.

Artificial general intelligence, by contrast, excels in domains that require rapid processing of large amounts of information, pattern recognition across complex datasets, and optimisation of well-defined objective functions. These systems can potentially surpass human performance in tasks that can be reduced to computational procedures, especially when these tasks do not require genuine understanding of meaning or significance.

The complementary nature of these different forms of intelligence suggests that optimal outcomes may require hybrid human-AI systems rather than replacement of human intelligence with artificial alternatives. Such systems could leverage the computational strengths of artificial systems while preserving the existential understanding and meaning-making capabilities of human intelligence.

5.3. The Question of Sufficiency: Can Behaviour Replace Being?

A central question in evaluating AGI claims concerns whether behavioural equivalence can substitute for genuine being and understanding. Current AGI research often assumes that systems capable of producing human-like outputs across a wide range of tasks have achieved human-like intelligence. However, the analysis of contingent intelligence suggests that this assumption may be fundamentally flawed.

The philosophical tradition has long recognised the distinction between appearance and reality, between simulation and genuine existence. Searle's Chinese Room argument illustrates this

distinction by demonstrating that a system can produce appropriate linguistic responses while completely lacking understanding of meaning or significance.

Contemporary examples of AI systems producing human-like outputs while lacking genuine understanding support this distinction. Large language models can generate compelling text about experiences they have never had, emotions they cannot feel, and values they cannot understand. Their success in producing convincing outputs may obscure rather than demonstrate genuine intelligence.

The behavioural equivalence approach faces what might be termed the "philosophical zombie problem" for AI. Just as philosophical zombies are conceived as beings that exhibit all the outward behaviours of consciousness while lacking inner experience, artificial systems might exhibit all the outward behaviours of intelligence while lacking the existential foundations that constitute genuine understanding.

This distinction has practical as well as theoretical significance. Systems that simulate understanding without genuine comprehension may perform adequately in routine contexts while failing catastrophically in novel situations that require genuine insight or adaptation. They may also make decisions that appear rational from a computational perspective while violating human values or expectations in ways that reflect their lack of authentic understanding.

6. Toward Artificial Collaborative Intelligence: An Alternative Framework

6.1. Reconceptualizing AI Development Goals

The analysis of contingent intelligence suggests the need for a fundamental reconceptualization of AI development goals. Rather than pursuing the impossible objective of replicating human beings in artificial systems, the field should focus on developing what this paper terms "Artificial Collaborative Intelligence" (ACI)—systems designed to complement and augment human capabilities while preserving human agency, meaning-making capacity, and moral responsibility.

The ACI framework begins with recognition that human and artificial intelligence represent qualitatively different forms of information processing and world engagement. Rather than viewing this difference as a limitation to be overcome, the ACI approach treats it as an opportunity for synergistic collaboration. By leveraging the complementary strengths of human contingent intelligence and artificial computational processing, ACI systems could potentially achieve outcomes that neither could accomplish alone.

This framework requires abandoning the anthropocentric assumption that artificial intelligence should replicate human cognitive processes. Instead, AI development should focus on creating systems that excel in domains where computational processing provides clear advantages—rapid information processing, pattern recognition, optimisation, and consistency—while interfacing effectively with human intelligence in domains that require existential understanding, meaning-making, and value-based judgement.

The ACI approach also emphasises the preservation of human agency and moral responsibility. Rather than replacing human decision-makers with artificial alternatives, ACI systems are designed to enhance human decision-making while ensuring that humans remain the ultimate source of values, goals, and moral judgement. This approach addresses concerns about AI alignment by ensuring that artificial systems remain tools for human flourishing rather than autonomous agents pursuing independently derived goals.

6.2. Design Principles for Collaborative Intelligence

The development of effective ACI systems requires adherence to several key design principles that reflect the fundamental differences between human contingent intelligence and artificial computational processing.

Transparency and Interpretability: ACI systems must be designed to make their reasoning processes transparent and interpretable to human collaborators. This requirement stems from the need for humans to maintain meaningful oversight and to integrate artificial outputs with human understanding and judgement. Opacity in artificial reasoning processes prevents effective collaboration and may lead to inappropriate reliance on systems that lack genuine understanding.

Value Alignment Through Human Agency: Rather than attempting to instil human values directly in artificial systems, ACI designs should ensure that value-based judgements remain within human control. Artificial systems should provide information and analysis to support human decision-making rather than making value-laden decisions independently. This approach recognises that values emerge from existential engagement and cannot be effectively replicated in systems that lack this foundation.

Contextual Adaptation: ACI systems should be designed to adapt their operation based on human feedback and contextual requirements. This adaptation should not involve autonomous learning that might drift away from human intentions but rather responsive adjustment to human guidance and oversight. The system's adaptation should enhance rather than replace human contextual understanding.

Error Detection and Graceful Degradation: Recognition of the limits of artificial understanding requires ACI systems to be designed with robust error detection and graceful degradation capabilities. Systems should be able to recognise when they are operating beyond their competence and should fail in ways that preserve human oversight and decision-making capability.

Cultural and Social Sensitivity: ACI systems operating in social contexts must be designed with deep awareness of cultural and social variation. Rather than assuming universal rationality or values, these systems should be designed to support diverse cultural approaches to reasoning and decision-making while avoiding the imposition of cultural assumptions.

6.3. Domains of Application

The ACI framework is particularly well-suited to domains that require the integration of computational processing capabilities with human understanding and judgement. Several key application areas illustrate the potential of this approach.

Scientific Research and Discovery: Scientific research requires both computational analysis of large datasets and human insight, creativity, and meaning making. ACI systems could handle data processing, pattern recognition, and hypothesis testing while human researchers provide theoretical understanding, creative insight, and interpretation of significance. This collaboration could accelerate discovery while preserving the human elements that drive scientific progress.

Healthcare and Medical Practice: Healthcare involves complex technical analysis combined with empathy, understanding of human values, and ethical judgement. ACI systems could provide rapid analysis of medical data, pattern recognition in imaging, and access to vast medical knowledge while human healthcare providers maintain responsibility for diagnosis, treatment decisions, and patient care. This approach could improve medical outcomes while preserving the human relationship that is central to effective healthcare.

Education and Learning: Educational practice requires both access to information and understanding of human development, motivation, and individual differences. ACI systems could provide personalised learning resources, assessment, and access to knowledge while human educators maintain responsibility for mentorship, motivation, and the development of wisdom and character.

Creative Arts and Design: Creative expression involves both technical skills and authentic personal expression emerging from lived experience. ACI systems could provide technical capabilities, access to creative resources, and assistance with execution while human artists maintain creative vision, personal expression, and cultural meaning-making.

Governance and Policy: Democratic governance requires both analysis of complex information and value-based judgement representing diverse human interests. ACI systems could provide policy analysis, modelling of outcomes, and information synthesis while human representatives maintain responsibility for value judgements, ethical considerations, and democratic accountability.

6.4. Implementation Challenges and Considerations

The development of effective ACI systems faces several significant challenges that must be addressed through careful design and ongoing research.

Interface Design: Creating effective interfaces between human contingent intelligence and artificial computational processing requires deep understanding of both human cognitive processes and artificial system capabilities. These interfaces must facilitate genuine collaboration rather than mere human oversight of autonomous artificial decision-making.

Trust and Reliability: Human users must be able to develop appropriate trust in ACI systems—neither over-reliance that leads to abdication of human responsibility nor under-utilisation that fails to leverage artificial capabilities. This requires systems that are both reliable within their domains of competence and clear about their limitations.

Scalability: ACI systems must be designed to scale effectively while maintaining the human oversight and involvement that is central to the approach. This may require hierarchical designs where human oversight operates at multiple levels of abstraction.

Cultural Adaptation: Effective ACI systems must be adaptable to diverse cultural contexts and value systems. This requires avoiding the assumption that there are universal approaches to reasoning or decision-making while still providing effective computational support.

Regulatory and Ethical Frameworks: The deployment of ACI systems requires new regulatory and ethical frameworks that address the unique challenges of human-AI collaboration. These frameworks must ensure human agency and accountability while enabling the benefits of artificial augmentation.

7. Implications for AI Ethics and Governance

7.1. Rethinking AI Alignment and Safety

The framework of contingent intelligence necessitates a fundamental reconceptualization of AI alignment and safety concerns. Traditional approaches to AI safety often assume that the primary challenge is ensuring that artificial systems pursue human-compatible goals while possessing capabilities that may eventually exceed human intelligence. However, if artificial systems cannot genuinely replicate human intelligence, then the nature of alignment challenges changes significantly.

Rather than focusing primarily on preventing artificial systems from pursuing goals that conflict with human values, attention should shift to ensuring that artificial systems remain effective tools for human flourishing while avoiding the risks that emerge from misunderstanding the nature and limitations of artificial intelligence. This includes addressing the risks of over-reliance on systems that lack genuine understanding, as well as the risks of anthropomorphising systems in ways that lead to inappropriate trust or delegation of responsibility.

The concept of "value alignment" itself requires reconsideration. If values emerge from existential engagement and cannot be effectively replicated in systems that lack this foundation, then attempts to align artificial systems with human values may be misguided. Instead, the focus should be on ensuring that value-based decisions remain within human control while artificial systems provide information and analysis to support human judgement.

This approach also addresses concerns about artificial consciousness and moral status. If artificial systems cannot achieve genuine consciousness or understanding, then questions about their moral rights and status become less pressing. However, this does not eliminate ethical concerns about artificial systems but rather shifts attention to ensuring that these systems serve human flourishing without creating harmful dependencies or misunderstandings.

7.2. Democratic Governance of AI Development

The recognition that artificial intelligence represents a fundamentally different kind of information processing than human intelligence has important implications for the democratic governance of AI development. Current discussions of AI governance often assume that artificial systems will eventually achieve human-like or superhuman intelligence, leading to concerns about maintaining human control over systems that may exceed human capabilities.

However, if artificial systems cannot achieve genuine understanding or consciousness, then they remain tools whose development and deployment should be subject to democratic oversight and

control. This perspective empowers democratic institutions to regulate AI development without concerns about limiting the rights or capabilities of genuinely intelligent artificial beings.

Democratic governance of AI development should focus on ensuring that artificial systems serve broad human interests rather than narrow commercial or technological goals. This includes ensuring that AI development priorities reflect diverse human values and concerns rather than the preferences of technical elites or commercial interests.

The framework also suggests the need for public education about the nature and limitations of artificial intelligence. Democratic governance requires informed citizens who understand both the capabilities and limitations of artificial systems. This education should address both the technical capabilities of AI systems and their philosophical and existential limitations.

International coordination of AI governance becomes important not because of concerns about artificial superintelligence but because of the need to ensure that AI development serves human flourishing globally. This includes addressing concerns about the concentration of AI capabilities, ensuring equitable access to AI benefits, and preventing the use of AI systems for oppression or exploitation.

7.3. Economic and Social Implications

The framework of contingent intelligence has significant implications for understanding the economic and social impacts of artificial intelligence. Rather than anticipating wholesale replacement of human labour by artificial systems, this framework suggests a future of human-AI collaboration where artificial systems augment human capabilities rather than replacing them entirely.

This perspective suggests that concerns about technological unemployment may be overblown, while concerns about the quality and meaning of work may be underappreciated. If artificial systems excel at computational tasks while lacking genuine understanding and creativity, then human work will likely shift toward domains that require these distinctly human capabilities.

However, this shift requires significant investment in education and training to help humans develop the capabilities that complement rather than compete with artificial systems. This includes not only technical skills but also the distinctly human capabilities of creativity, empathy, moral reasoning, and cultural understanding.

The framework also suggests the need for economic structures that value and compensates human capabilities that cannot be replicated artificially. This may require rethinking economic models that prioritise efficiency and productivity over meaning, creativity, and human well-being.

Social institutions will need to adapt to support effective human-AI collaboration while preserving the human relationships and communities that are essential for human flourishing. This includes ensuring that artificial systems enhance rather than replace human social connections and cultural practices.

8. Contemporary Challenges and Empirical Evidence

8.1. Current Limitations of Large Language Models

Contemporary developments in large language models (LLMs) provide compelling empirical evidence for the limitations identified in the theoretical analysis of contingent intelligence. Despite impressive performance on many linguistic and reasoning tasks, these systems exhibit systematic failures that reflect their lack of genuine understanding and existential grounding.

Research on LLM capabilities reveals consistent patterns of what Marcus (2022) terms "elegant failure"—impressive performance on benchmark tasks coupled with systematic errors that reveal the absence of genuine understanding. These systems can produce fluent text about experiences they cannot have, explain concepts they do not understand, and provide confident responses to questions beyond their competence.

The phenomenon of "hallucination" in LLMs provides particularly clear evidence for the limitations of systems that lack grounding. These systems can generate detailed, confident descriptions of events that never occurred, people who do not exist, and facts that are entirely fabricated. This behaviour reflects the systems' reliance on statistical patterns in training data rather than genuine understanding of reality.

Studies of LLM performance on tasks requiring common-sense reasoning reveal systematic limitations that appear to stem from the absence of embodied experience. These systems struggle with tasks that require understanding of physical causation, spatial relationships, and temporal sequences that would be obvious to any embodied agent with experience in the physical world.

Research on the social and cultural dimensions of LLM behaviour reveals additional limitations. These systems can reproduce cultural biases present in training data while lacking genuine understanding of cultural context or the ability to navigate cultural differences with the sensitivity that characterises human social intelligence.

8.2. The Model Collapse Problem

Recent research on "model collapse" when AI systems are trained on AI-generated data provides striking empirical support for the theoretical arguments about the limitations of artificial intelligence (Shumailov et al., 2023). This research demonstrates that AI systems trained on outputs from other AI systems gradually lose coherence and connection to reality, eventually producing degenerate outputs.

This finding illustrates the fundamental dependence of artificial systems on human-generated data and their inability to maintain grounded understanding independently. Unlike humans, who continuously ground their understanding through ongoing interaction with reality, artificial systems rely on static training data and cannot refresh their understanding through genuine experience.

The model collapse phenomenon suggests that artificial systems may be fundamentally parasitic on human intelligence and cannot achieve the autonomous understanding that would be necessary for genuine intelligence. This challenges assumptions about the potential for artificial systems to achieve or exceed human cognitive capabilities through continued development.

The research also reveals the fragility of artificial intelligence systems when removed from the human context that provides their training data and meaning. This fragility reflects the absence of genuine grounding that characterises contingent intelligence.

8.3. Energy and Resource Constraints

The physical resource requirements of contemporary AI systems provide additional empirical support for theoretical limitations on artificial intelligence scaling. Research on the energy consumption of large AI models reveals exponentially increasing requirements for computational resources as model size increases (Strubell et al., 2019; Patterson et al., 2021).

Current estimates suggest that training GPT-4 consumed between 51 and 62 million kWh of electricity and generated 12,000 and 15,000 metric tonnes of carbon emissions (Medium analysis, 2023). The energy requirements for inference—running trained models—are also substantial, with estimates suggesting that ChatGPT queries consume 10 times more energy than Google searches.

These resource requirements impose physical constraints on the scaling hypothesis that are independent of algorithmic or architectural improvements. The exponential growth in computational requirements may eventually outstrip available energy and material resources, creating absolute limits on model scaling.

More fundamentally, these constraints reveal the materiality of artificial intelligence systems that contrasts with the embodied efficiency of biological intelligence. Human brains accomplish remarkable cognitive feats using approximately 20 watts of power—less than a standard light bulb—while artificial systems require massive computational infrastructure for far more limited capabilities.

8.4. Social and Cultural Limitations

Empirical research on AI systems in social contexts reveals systematic limitations that support theoretical arguments about the importance of genuine social embeddedness for intelligence. Studies of chatbots and social robots reveal that while humans may initially respond positively to socially capable artificial systems, longer-term interactions often expose the absence of genuine understanding and care.

Research on AI performance in culturally diverse contexts reveals systematic biases and limitations that reflect the systems' dependence on training data from cultural contexts. These

systems often fail to navigate cultural differences with the sensitivity and understanding that characterises human cross-cultural interaction.

Studies of AI systems in educational contexts reveal that while these systems can provide information and feedback, they cannot replicate the motivational and relationship-based aspects of human teaching that are crucial for learning and development. Students may learn facts from AI systems but require human mentorship for developing wisdom, character, and motivation.

Research on AI in healthcare contexts similarly reveals that while AI systems can provide diagnostic support and data analysis, they cannot replicate the empathy, understanding, and relationship-building that are central to effective healthcare. Patients may receive technically accurate information from AI systems but require human care providers for healing and support.

9. Future Research Directions

9.1. Deepening Understanding of Contingent Intelligence

The framework of contingent intelligence opens several important directions for future research that could deepen understanding of human cognition while informing more effective approaches to artificial intelligence development.

Empirical research on the role of embodiment in cognition could provide more detailed understanding of how physical experience shapes conceptual understanding, emotional processing, and social interaction. This research could inform the design of AI systems that better complement human cognitive capabilities by understanding their embodied foundations.

Studies of the relationship between mortality awareness and meaning making could illuminate the existential foundations of human motivation, creativity, and value systems. This research could inform approaches to designing AI systems that support rather than undermine human meaning-making processes.

Research on the development of social and cultural understanding could reveal the processes through which humans acquire the contextual knowledge that enables effective social interaction. This research could inform the design of AI systems that operate more effectively in social contexts while avoiding inappropriate simulation of social understanding.

Philosophical research on consciousness and the hard problem could continue exploring the nature of subjective experience and its role in human cognition. While this research may not directly inform AI development, it could provide crucial insights into the fundamental differences between human and artificial information processing.

9.2. Developing Artificial Collaborative Intelligence

The ACI framework requires substantial research and development to create effective systems that genuinely augment human capabilities while preserving human agency and understanding.

Research on human-AI interface design should explore how to create effective collaboration between human contingent intelligence and artificial computational processing. This research should address both technical challenges of system design and human factors considerations of trust, understanding, and appropriate reliance.

Studies of collaborative decision-making could illuminate how humans and artificial systems can work together effectively in complex domains that require both computational analysis and human judgement. This research should explore how to preserve human agency and responsibility while leveraging artificial capabilities.

Research on explainable AI specifically designed for human collaboration could develop new approaches to making artificial reasoning transparent and interpretable to human partners. This research should address not only technical aspects of explanation but also human psychological factors that affect understanding and trust.

Development of evaluation frameworks for ACI systems should create new approaches to assessing the effectiveness of human-AI collaboration rather than focusing solely on artificial system performance in isolation.

9.3. Ethical and Governance Research

The implications of contingent intelligence for AI ethics and governance require substantial research to develop appropriate frameworks for regulating and deploying artificial systems.

Research on democratic governance of AI should explore how to ensure that AI development serves broad human interests while remaining accountable to democratic oversight. This research should address both institutional design questions and public engagement challenges.

Studies of AI impact on work and society should examine how artificial systems affect human capabilities, relationships, and meaning-making processes. This research should go beyond economic impacts to consider effects on human flourishing and social cohesion.

Research on cultural dimensions of AI should explore how artificial systems interact with diverse cultural values and practices. This research should inform approaches to designing systems that support rather than undermine cultural diversity and autonomy.

Philosophical research on the ethics of human-AI relationships should explore the moral considerations raised by increasing integration of artificial systems into human social and cultural contexts. This research should address questions of authenticity, manipulation, and the preservation of human agency.

9.4. Longitudinal Studies of AI Impact

Understanding the long-term implications of artificial intelligence deployment requires longitudinal research that tracks effects over extended periods rather than focusing only on immediate technical capabilities.

Studies of human adaptation to AI systems should examine how prolonged interaction with artificial systems affects human cognitive capabilities, social relationships, and meaning-making processes. This research should explore both positive augmentation effects and potential negative dependencies.

Research on institutional and cultural change should examine how artificial systems affect social institutions, cultural practices, and collective decision-making processes. This research should explore how to preserve valuable human institutions while enabling beneficial technological augmentation.

Longitudinal studies of AI system evolution should track how artificial systems change over time through continued training and deployment. This research should examine whether these systems develop in directions that support or undermine human flourishing.

10. Conclusions

This analysis has developed a comprehensive framework for understanding why artificial general intelligence, as currently conceived, cannot replicate the existential foundations of human cognition. The concept of contingent intelligence reveals that human cognitive capabilities emerge from and remain dependent upon specific existential conditions—embodied experience, phenomenal consciousness, mortality awareness, and social embeddedness—that cannot be reproduced in artificial systems designed according to computational paradigms.

The systematic examination of barriers to AGI replication demonstrates that these limitations are not merely technical challenges to be overcome through increased computational power or algorithmic sophistication but rather reflect fundamental ontological differences between artificial computational processing and human existential engagement with the world. The embodiment problem, the irreducibility of phenomenal consciousness, the absence of mortality- driven meaning-making, the inability to ground values intrinsically, and the limits of social simulation represent categorical rather than gradual differences between artificial and human intelligence.

The comparative analysis reveals that human contingent intelligence and artificial computational processing represent qualitatively different forms of information processing and world engagement rather than different points along a single continuum of intelligence. This recognition enables a more realistic assessment of both the capabilities and limitations of artificial systems while affirming the unique and irreplaceable nature of human cognitive capabilities.

The alternative framework of Artificial Collaborative Intelligence offers a more promising approach to AI development that leverages the complementary strengths of human and artificial

intelligence while preserving human agency, meaning-making capacity, and moral responsibility. Rather than pursuing the impossible goal of replicating human beings in machines, the ACI framework focuses on creating systems that enhance human capabilities and support human flourishing.

The implications extend far beyond technical considerations to encompass fundamental questions about the nature of intelligence, consciousness, and human existence. The framework suggests the need for renewed appreciation of the distinctive qualities of human intelligence that emerge from our contingent existence as embodied, mortal, and socially embedded beings.

This perspective has profound implications for AI ethics and governance, suggesting that concerns about artificial superintelligence may be misplaced while highlighting the importance of ensuring that artificial systems serve human flourishing without creating harmful dependencies or misunderstandings about their nature and capabilities.

The empirical evidence from contemporary AI research supports the theoretical arguments developed here, revealing systematic limitations in current systems that reflect their lack of genuine understanding and existential grounding. These limitations appear to be fundamental rather than temporary obstacles to be overcome through technological advancement.

Future research directions include deepening understanding of contingent intelligence, developing more effective approaches to human-AI collaboration, exploring the ethical and governance implications of the framework, and conducting longitudinal studies of AI impact on human capabilities and social institutions.

Perhaps most importantly, this analysis suggests a fundamental reorientation of AI development goals away from the pursuit of artificial minds that rival or replace human intelligence toward the development of tools that enhance human capabilities while preserving the existential foundations that make human intelligence meaningful and valuable.

The quest for artificial general intelligence has often been motivated by desires to transcend human limitations—to create minds that are faster, more consistent, and more capable than biological intelligence. However, the analysis developed here suggests that these apparent limitations may be the source of the distinctive strengths of human intelligence. Our embodiment grounds understanding, our mortality creates urgency and meaning, our social nature enables empathy and cultural development, and our consciousness provides the subjective foundation for values and aesthetic experience.

Rather than seeking to transcend these conditions through artificial replication, the framework of contingent intelligence suggests embracing and enhancing them through collaborative artificial systems that support rather than replace human engagement with existence. This approach offers the possibility of technological advancement that serves human flourishing while preserving the existential foundations that make human life meaningful.

The ultimate significance of this analysis may lie not only in its implications for artificial intelligence development but in what it reveals about the nature and value of human intelligence itself. By attempting to understand what cannot be replicated artificially, we gain deeper appreciation for the remarkable achievement of human consciousness and the contingent conditions that make it possible. In preserving these conditions and the intelligence that emerges from them, we preserve something irreplaceable and essential to human existence.

The future of artificial intelligence should therefore be orientated not toward replacing human intelligence but toward supporting the conditions that enable human intelligence to flourish. This requires not only technical advances in creating more effective collaborative systems but also broader cultural and institutional changes that preserve and enhance the existential foundations of human intelligence in an increasingly technological world.

This vision offers hope for a future in which technological advancement serves human flourishing while preserving the distinctive qualities that make human existence meaningful and valuable. Rather than seeing humans and artificial systems as competitors in a zero-sum game for intelligence dominance, the framework of contingent intelligence enables recognition of their complementary strengths and collaborative potential.

In conclusion, the pursuit of artificial general intelligence, while scientifically fascinating and technologically impressive, may represent a fundamental misunderstanding of the nature of intelligence itself. By recognising that human intelligence is contingent upon conditions that cannot be replicated artificially, we can redirect technological development toward more realistic and beneficial goals. The resulting artificial collaborative intelligence systems, designed to enhance rather than replace human capabilities, offer greater promise for supporting human flourishing while preserving the existential foundations that make human intelligence uniquely valuable.

References

1. Becker, E. (1973). *The Denial of Death*. Free Press.
2. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623.
3. Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
4. Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200-219.
5. Dreyfus, H. L. (1972). *What Computers Can't Do: A Critique of Artificial Reason*. Harper & Row.
6. Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press.
7. Fjelland, R. (2020). Why general artificial intelligence will not be realised. *Humanities and Social Sciences Communications*, 7(1), 1-9.
8. Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, 22(3-4), 455-479.
9. Goertzel, B. (2014). Artificial general intelligence: Concept, state of the art, and prospects.
10. Journal of Artificial General Intelligence, 5(1), 1-48.
11. Greenberg, J., Pyszczynski, T., & Solomon, S. (1997). Terror management theory of self-esteem and cultural worldviews: Empirical assessments and cultural refinements. *Advances in Experimental Social Psychology*, 29, 61-139.
12. Heidegger, M. (1962). *Being and Time*. Harper & Row. (Original work published 1927)
13. Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., ... & Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
14. Lakoff, G., & Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Basic Books.
15. Marcus, G. (2018). Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*. Marcus, G. (2022). *The Road to AI We Can Trust*. Under consideration at MIT Press.
16. Merleau-Ponty, M. (1945). *Phenomenology of Perception*. Routledge.
17. Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux.
18. Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435-450. OpenAI. (2023). GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
19. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., ... & Dean, J. (2021).
20. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
21. Sartre, J. P. (1943). *Being and Nothingness*. Philosophical Library.
22. Searle, J. R. (1980). Minds, brains, and programs. *Behavioural and Brain Sciences*, 3(3), 417-424.
23. Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., & Anderson, R. (2023). The curse of recursion: training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*.
24. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645-3650.
25. Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5(1), 42.
26. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.
27. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

