Discussion Paper

Transformative Pathways towards Meaningful Collective Intelligence

-

AI in Service of People, Planet & Peace

Abstract

The development of AI is often regarded as inevitable, with public discourse primarily focused on managing its downstream risks rather than critically examining its intended beneficiaries and underlying purposes. This background paper argues that treating adoption of AI is a political choice, one that forecloses the more fundamental governance question the Club of Rome is positioned to insist on beyond AI safety guardrails: what is AI for, and in whose interest is it being built.

This paper examines AI as a set of non-neutral capabilities shaped by specific actors with specific interests, and traces how that shaping plays out across four domains: the unfulfilled promises of an earlier wave of digitalisation, whose techno-optimism AI now repeats; the political economy of AI, including data colonialism, labour displacement and gendered harms, and AI's entanglement with defence and surveillance infrastructure; the scale and unequal distribution of AI's ecological costs, from electricity and water use to e-waste; and the exclusion of Global Majority actors from the standards and institutions currently defining AI governance. Building on this analysis, the paper sets out governance principles oriented toward transformation rather than adaptation or harm mitigation, including data sovereignty as a collective right, digital commons as an institutional alternative to platform monopoly and state control, pluriversal epistemology as a design principle, and international precedent for coordinated restraint on the riskiest frontier development.

The paper refrains from treating AI's current trajectory as fixed. It calls for governance built on a prior democratic question of what it is needed for, and for whom, open to the answer that some AI, in some form, should not proceed. It closes by setting out the conditions under which AI could genuinely serve people, planet, and peace, rather than reinforcing the concentrations of power it was built within. This paper is meant to document the Club of Rome's call for actions and contributions: toward a fuller 2027 publication, and, in parallel, toward policy initiatives or interventions that members can begin to activate now.

1. The Prior Question: Do We Actually Need AI?

Most AI discourse treats adoption as inevitable, and the only real task as managing risks downstream. But the "train has already left the station" framing is itself a political choice, one that forecloses deliberation and delegitimises dissent before it can form. Asking *why* we need AI, and *for whom*, is not a Luddite move. It is the foundational governance question. Skipping it obscures the sense of human sovereignty and produces immediate dispossession. Every one of us is forced into living with the AI we know, even if it was not conceived with human and ecological needs in mind, but rather to increase control and extraction from humans and ecosystems. Is the technologically feasible instantly blessed with legitimacy? This question is not even asked today, while the 20th century has provided too many lessons on the perils of blind scientific and technological achievements, from

The Club of Rome
Lagerhausstrasse 9, 8400 Winterthur, Switzerland
Tel. +41 52 244 0808
www.clubofrome.org

eugenics as an instrument of supremacism to nuclear bombs and drones to continue putting wars at the centre of human societies. AI is not less serious than this.

There are genuine domains where AI[1] capabilities offer positive impact that would be difficult to achieve otherwise or at scale: pattern recognition in ecological monitoring, multilingual support for marginalised-language communities, diagnostic aid in under-resourced health systems, early disease detection from population-scale health data (University College London, 2026) and accelerating synthesis of fragmented scientific knowledge, to mention a few. AI's ability to process large, complex ecological datasets holds genuine promise for lifting barriers in conservation science and biodiversity monitoring, including through species recognition, ecosystem modelling, and evidence synthesis (Reynolds et al., 2025). Recent reviews have found that AI is currently being used in only two of the seven major biodiversity knowledge shortfall areas, leaving significant potential untapped in conservation science (McGill University, 2025).

At the same time, these gains are not guaranteed by the technology itself: there has been too much techno-optimism and insufficient attention to risks and broader systemic implications (Sandbrook, 2025). More generally, AI tends to reinforce the power structures it is deployed into rather than transform them, and even genuinely positive uses will not reach their full potential within a system that concentrates ownership, data and decision-making at the top. Precision agriculture illustrates the pattern: AI-enabled tools may reinforce existing power hierarchies and undermine local and smallholder knowledge systems rather than serve them (Sandbrook, 2025, citing Foster et al., 2023). Comparable dynamics recur wherever AI is layered onto already unequal systems.

The honest position is not "do we need AI" in the abstract, but: *need it for what, and who currently controls the answer to that question?* The AI we know is more than a technology with a singular trajectory. It is an underlying ideology that admits no questioning since it is supposed to produce higher efficiencies for economies that can get rid of humans, for war machines that can kill without risks, and for platforms that can accumulate virtual wealth at the top while siphoning the value of data from everyone else. The AI we know is a set of capabilities being shaped by specific actors with specific interests. The question of necessity has to be answered in relation to those power structures.

Underlying all of this is a more fundamental confusion that has come to define mainstream AI discourse: the conflation of technological development with human development. The two are not the same, and treating them as equivalent has consequences. Technological capability expands what is *possible*; it says nothing about what is *desirable*, *just*, or *sustainable*. The assumption that more capable AI systems will naturally produce better human outcomes bypasses the ethical questions that should come first, questions about whose values are encoded, whose needs are prioritised, whose risks are discounted, and whose knowledge is rendered invisible. The Club of Rome has argued more broadly that development itself must be reframed around human and ecological wellbeing rather than the expansion of economic output; the same logic applies here. Technological development that is not preceded and constrained by ethical deliberation tends to reproduce and accelerate existing inequalities rather than disrupt them. This is not a novel observation, it has been made about every major general-purpose technology, but it is one that AI governance discourse consistently forgets in its urgency to "keep up." A holistic perspective of ethics, for humans and ecosystems, is not a layer to be added after the architecture is built. It is the condition under which the architecture should be designed from inception. This comes as a lesson learned from 40 intense years of digitalisation, along patterns that AI is not disrupting but reinforcing. Wilber (2000) and other researchers have pointed to frameworks for defining shared human values across different worldviews, but such work remains preliminary and cannot substitute for the ethical deliberation this paper calls for.

[1] "AI" spans narrow task-specific systems, general-purpose foundation models and increasingly autonomous "agentic" systems, each with different governance implications (OECD, 2022, 2024). Unless otherwise specified, "AI" here means the dominant paradigm: large-scale foundation-model-based systems, around which the political economy described in this paper has formed.

2. The Unfulfilled Promises of Digitalisation

AI does not come out of the blue, and it is not based on new breakthrough science. Its capacities are impressive and yet they are the result of the evolution of several factors combined, largely due to decades of digitalisation, with an inflection point happening 40 years ago, with the onset and generalisation of personal computing and ten years later of the internet. The AI we have would not have been possible without the increase by many orders of magnitude of computing power together with the accumulation of data in digital form. In this sense, AI is the consequence and an emerging feature of the larger wave of digitalisation, which in turn has been a key dimension of globalisation.

Then, it is possible to shed light on AI and its promises from the observation of four decades of digitalisation and what they led to. In the late 1990s, a techno-optimistic vision was pointing to the internet as an effective means of empowerment of individuals and of reinforcement of democracy. The creators of the internet were largely motivated by devolving power to the people. Wikipedia and other examples were mentioned as the proof of the emergence of higher levels of collective intelligence for the common good. And many identified the interference of governments as the main obstacle to the full development of the potential of the internet. John Barlow, founder of the Electronic Frontier Foundation (EFF), even proposed "A Declaration of Independence of the Cyberspace".

Thirty years later, Wikipedia still exists (and is still struggling to survive), but it is a drop of water in the ocean dominated by the digital platforms, a modern form of techno-feudalism that makes wealth accumulate at the top while most of us are captured and provide data and work for free. The recent unveilings of the political visions and engagements of the cyber-libertarian billionaires known as the Tech Bros such as Elon Musk or Peter Thiel do nothing but confirm that the beautiful aspirations of digitalisation have turned into regressive, anti-democratic and anti-human nightmares. After privatising the internet and AI for their own benefit, they are trying to privatise the state as well, at least in the USA. Any serious analysis about the promises or inevitability of the AI we know needs to address this dimension of who owns the trajectory of AI evolution

3. Power Structures and the Political Economy of AI: Who Owns the Trajectory

Current AI development is concentrated in a small number of large corporations, predominantly US and Chinese. This concentration reinforces and even expands further the already existing structural dependency for most of the world, one that mirrors prior asymmetries in pharmaceutical patents, agricultural seed ownership, and financial architecture, among others, where the benefits of a global system are distributed unequally while the costs fall disproportionately on the peripheries of power, within and across countries. Several interlocking features of this political economy deserve explicit attention.

Who owns the means and the debate. What is under way is a new process of enclosure. In the USA, infrastructures of a kind states once built – Amazon, Microsoft, Alphabet and Meta alone budgeted over $300 billion of capital spending in 2025, largely for data centres and chips – is being assembled in private hands and then let out for rent. The pattern is familiar from land, minerals and patents: whoever owns the asset holds a standing claim on the value everyone else creates with it, while the costs of building it – minerals, energy, water, land and underpaid data work – fall on people and places with no share in the returns. Governments appear as customers and hosts rather than owners. Wealth and decision-making power concentrate together, which is why AI governance is a distributional question with severe consequences for our social contract. Even in the public debate on AI, the presence of the big corporations is overwhelming, as if it was taken for granted that they should have the first roles in defining what type of governance is needed. The recent alarms about the possibility of an AI-driven extinction of humanity provide a bizarre paradox: no one is questioning the legitimacy of the corporations to discuss what to do, since they are supposed

to know better. This would be unimaginable if we were talking of private providers of atomic bombs or biological weapons.

Data dispossession and sovereignty. The same logic of enclosure extends to data: the traces of everyday life – what people say, buy, search for and where they go – are captured as raw material for AI, and a growing body of scholarship describes this appropriation as a new form of colonialism. Couldry and Mejias (2019) argue that data colonialism constitutes "an emerging order for appropriating and extracting social resources for profit through data", a process through which human life itself is annexed to capital, with the Global Majority serving as both a primary site of extraction and a testing ground for unregulated forms of datafication. This manifests through material and digital forms of extraction and exploitation that sustain conditions of dependency, concentrated wealth accumulation, and power asymmetries between a very small group of happy few and everyone else. Ricaurte (2019) extends this critique further, analysing data-centric epistemologies as expressions of the coloniality of power, manifested as the "violent imposition of ways of being, thinking, and feeling that leads to the expulsion of human beings from the social order, denies the existence of alternative worlds and epistemologies, and threatens life on Earth." Understanding this as a regime of rationality, she argues, is only the first step. Through a reductionistic exercise, the complexity of living systems is assumed to be replaceable by explicit streams of quantifiable and decontextualised data. And then, economies based on the prominence and privatisation of data foster extractive models of resource exploitation, implying and ignoring the violation of human rights, cultural exclusion, and ultimately ecocide. When everything is assumed to be a data source, when life itself is understood as nothing more than a continuous flow of data, alternative forms of being, thinking, and sensing are crowded out. (Ricaurte, 2019). A historical, decolonial perspective reveals how uneven distributions of power materialise as digital colonisation, algorithmic violence, gender violence, class divides, and racism.

This dynamic is not abstract. "AI colonialism is a central concern – data potentially extracted from the Global South might be forwarded predominantly to data centres in the Global North for training AI models, followed by AI-driven mandates being issued to the Global South on how land and resources should be managed" (Reynolds et al., 2025). That this concern is being named within a mainstream conservation journal is itself significant: the extractive logic of AI development is now visible even in fields that are otherwise enthusiastic about AI's potential.

Displacement of human labour, gender imbalances and care. Acemoglu, Autor and Johnson (2023) argue that the currently predominant direction of AI innovation emphasises automation, displacement of skilled labour, and diminished worker voice due to stepped-up monitoring and surveillance. Countries where formal service-sector employment is most exposed to automation, and where workers have the least structural power to shape the direction of technology, will absorb these displacements most harshly. The evidence on aggregate effects is mixed, though: some studies find AI augmenting rather than replacing many tasks, at least so far (Noy & Zhang, 2023; Eloundou et al., 2024), and outcomes vary considerably by sector, occupation and how automation is deployed. And indeed, automation cannot be applied to the tasks that we do not know yet. The transformation of economies and societies towards regeneration, a balance between equitable human development and ecosystemic boundaries, would create many new jobs, and this creation could actually be enabled by AI, if properly framed.

But this is not where we are now, at least not yet. What is more consistent across the evidence is the distributional pattern: even where AI creates net gains, the workers and countries least able to shape its deployment are disproportionately exposed to its costs. Financing and governance frameworks that treat "AI for development" as a purely technical or fiscal question risk deepening this gap rather than closing it (Bretton Woods Project, 2026). More generally and in the way it is presently framed, AI becomes an instrument of massive dispossession, especially in the light of its duality: considering it both inevitable and an existential threat for humanity leaves almost everyone in a state of helplessness, by wish we relinquish our agency into the hands of the very few immensely benefitting from its privatisation.

AI, as framed in a platform economy that is increasingly recognised as “techno-feudalism”, intensifies extraction, based not only on data capture without retribution but also on low-paid data-labelling and content-moderation work, performed disproportionately in the Global Majority and reproducing rather than shifting gender imbalances (Gray & Suri, 2019). This pattern extends into the AI workforce itself: globally, women hold roughly a fifth of AI-related jobs, and the gap is more acute in developing economies, where women are also stereotyped into the domestic and caregiving roles that AI and robotics are now marketed to automate (Kumar & Choudhury, 2022). Even the public representation of AI in movies extends or aggravates the gender imbalances: women make less than 10% of AI roles (Cave, Dihal, Drage & McInerney, 2023).

A feminist political economy lens makes clear that AI is not a neutral tool: it constitutes and reproduces the social relations that decide whose labour counts and gets paid. Those same relations leave other labour unpaid, unseen and taken for granted (Lee & Dederichs, 2025). Furthermore, it can be argued that automated care technologies quietly appropriate and devalue the relational labour of care rather than replacing the need for it (Rhee, 2023). Overall, a transformative agenda for AI has to name, explicitly, the labour that subsidises AI's efficiencies and the costs it pushes onto people it doesn't count.

Armaments, Surveillance, Intelligence War Efficiencies. The relationship of AI with war is not collateral. Since the 1970s, substantial efforts have been invested to intensify the technological level of military action with the deliberate intention to reduce the involvement of humans. Besides higher efficacy, this pursued a political goal: the dramatic reduction of human losses on the more tech-savvy side in order to legitimise again the use of war, without compromising the support by society at large as happened with the Vietnam War. A detailed analysis of this relationship between war and technology goes beyond the purpose of this paper.

At the same time, let us emphasise that the technologies now marketed as civilian AI have deep roots in defence, energy and intelligence funding, and that orientation persists and is intensifying: companies including Palantir, Google, Microsoft, AWS and Meta are now heavily involved in defence contracting, supplying, for instance, the cloud infrastructure and classified-network access behind current Pentagon AI operations; venture capital flowing into defence technology alone nearly doubled between 2024 and 2025, from roughly $27 billion to $49 billion (Yeung, 2026). This matters for how AI is framed. Rather than being simply a civilian technology, to a significant degree it is an instrument of war, developed and funded to make surveillance, targeting and armed conflict more efficient. An account of AI’s trajectory that treats war efficiencies as a secondary function understates how much of the underlying infrastructure, capital and political will has been organised around exactly that purpose.

The production of standards. Technical standards, safety frameworks, and governance norms are being written predominantly in Global Minority institutions. Specific areas of concern for Global Majority actors include infrastructural and regulatory monopolies, harms associated with the labour and material supply chains of AI infrastructure, and commercial exploitation, while countries that arrive late to governance tables do not adopt neutral rules, but rules written by competitors with different interests (Png, 2022). Concerns remain that the explosion in soft law instruments professing lofty ethical principles is largely toothless in practical terms, leaving the Global Majority with little means to hold big tech companies and their governments accountable for unethical practices (Vijayakumar, 2024).

The standards game must also be understood through the lens of coloniality. Png (2022) argues that the concentration of harms and benefits of AI systems across the Global South and North reflects the continuation of historical processes of dispossession, extraction, and exploitation that shaped the modern world, within countries and across countries. International laws and institutions underpinning global AI governance were themselves constructed by and rooted in institutions of geopolitical

domination, meaning that even well-intentioned efforts toward inclusive governance risk reproducing structural injustices if they lack a macro-historic analysis of how those institutions operate.

4. Ecological Consequences of AI: Huge and Unequally Distributed

The infrastructure underpinning AI carries ecological costs that are both staggering in scale and deeply uneven in distribution, as documented in a recent report by Aczel et al. (2026) for the United Nations University Institute for Water, Environment and Health. In 2025, data centres consumed an estimated 448 TWh of electricity, a level that would rank them eleventh globally if data centres were a country. On current trajectories, that demand could roughly double to 945 TWh by 2030, nearly triple the combined annual electricity use of Pakistan, Bangladesh, and Nigeria, together home to more than 650 million people. Producing that much electricity would generate a carbon footprint of 399 million tonnes of CO2, requiring 6.7 billion trees grown over ten years to offset, roughly twice the number of trees in the United Kingdom. The associated water footprint of 9.3 trillion litres would equal the annual domestic water needs of all 1.3 billion residents of Sub-Saharan Africa, and the land footprint would exceed 14,500 km^2, nearly ten times the size of Mexico City. AI workloads alone accounted for around 20% of data centre electricity use in 2025; if that share rises to 40% by 2030 as projected, AI's electricity demand could reach roughly 378 TWh, enough to meet the residential electricity needs of the entire population of Sub-Saharan Africa for over two years. The International Energy Agency and other researchers have pointed to potential energy savings through applications such as AI-optimised power grids and more efficient industrial processes, but these savings depend on deliberate deployment and are not guaranteed by the technology itself (International Energy Agency, 2025).

The United Nations University Institute for Water, Environment and Healthreport further states that the environmental burden does not end with energy. AI hardware depends on critical minerals whose extraction and processing cause environmental and social harms concentrated disproportionately in the Global Majority, often in places with weak regulatory oversight. At end of life, poorly managed e-waste exposes frontline communities to hazardous substances. By 2030, AI infrastructure could generate up to 2.5 million metric tonnes of e-waste annually, roughly equivalent to discarding 250 Eiffel Towers every year. These patterns confirm that responsible AI governance requires full value-chain oversight, from mineral sourcing through to recycling and safe disposal, and that the costs of the current trajectory fall hardest on those with the least power to resist them (Aczel et al., 2026).

Dauvergne (2022) has further shown that AI efficiency gains in supply chains tend to be reinvested into more production and consumption, ultimately accelerating rather than reducing ecological pressure. Again,all these environmental and social costs are matters of concern by themselves and in addition they are not evenly shared: communities in the Global Majority often bear the brunt of carbon emissions, water withdrawals, e-waste, and hidden labour that make AI possible, without sharing in its benefits. While the relationship between levels of inequality and the overall ecological footprint of humanity has been repeatedly documented, AI as we know it is not contributing to reduce neither inequality nor the footprint.

5. AI Governance for People, Planet & Peace

Dealing with the limitations and risks of AI as we know it is not simply a matter of adding more voices to existing frameworks. It requires reorienting the foundational questions that governance is asked to answer. It is important to highlight that “Global Majority” means here more than the “Global South”: the vast majority of citizens in countries of the Global North are also being dispossessed due to the ways that AI is being implemented. Global Majority actors can play three substantive roles:

challenging exclusionary governance mechanisms; providing legitimate expertise in the interpretation and localisation of risks through whole-systems, contextual and historical views; and providing sources of alternative governance mechanisms, including South-South solidarity, co-governance, and democratic accountability (Png, 2022). Several principles follow.

Data sovereignty as collective right. This is the governance response to the pattern of data colonialism described above: if data extraction concentrates wealth and power in the Global Minority, data sovereignty is what allows the Global Majority to set the terms. The dominant model of data governance, centred on individual consent, corporate self-regulation, and standards set by Global Minority institutions, has been increasingly contested by alternative frameworks emergingfrom the Global Majority. These frameworks share a common orientation: treating data as a collective asset requiring governance rooted in community rights, historical redress, and developmental priorities beyond a resource to be extracted and monetised. The African Union's Continental Data Policy Framework (2022) offers one of the most substantively developed articulations of this alternative, establishing that data justice extends beyond individual privacy rights to social and economic rights, and that governance must actively address the ways in which datafication risks perpetuating historical injustices and structural inequalities (African Union, 2022). Building on this, the AU Continental AI Strategy (2024) proposes a five-year implementation framework oriented around harnessing AI's benefits for people and institutions, building continental capabilities, and fostering cooperation on Africa's own terms (African Union, 2024). These are not isolated regional experiments. They form part of a broader and still-nascent architecture of Global Majority AI governance, connected to analogous processes in Latin America through the 2024 Santiago Declaration and in Southeast Asia through ASEAN's AI governance guidelines, in which the Majority begins to set standards rather than receive them.

Digital commons as institutional alternative. A recurring theme in Club of Rome member discussions of this paper is that existing governance vocabulary, built for a world of states and private firms, has no natural home for AI as a shared resource. One institutional answer, developed independently in internet governance and in recent digital-infrastructure practice, is to govern computational infrastructure explicitly as a commons: a resource held and managed collectively, with rules for access and contribution set by its users rather than by a monopoly owner or by the state alone. Tim Berners-Lee, reflecting on the web's capture by extractive platforms, has called for exactly this kind of re-collectivisation through the Contract for the Web's public-good principles (Berners-Lee / World Wide Web Foundation, 2019) and through the Solid protocol (Sambra et al., 2016), which aims to give individuals control over their own data rather than surrendering it to platforms. A body of recent case-study research backs this up empirically: projects such as Barcelona's Decidim platform, Germany's Center for Digital Sovereignty, the DHIS2 health-data platform, and the European Open Science Cloud already function as digital commons providing public infrastructure, with public institutions holding the protocols while private and civic actors build on top of them (Open Future, 2024). Whether such models can scale from these sector-specific examples to AI at large is untested, but they demonstrate that "infrastructure as commons" is a working alternative to both the platform-monopoly model and the state-control model.

Pluriversal epistemology as design principle. AI systems encode assumptions about what counts as knowledge, what relationships matter, and how causality works. These assumptions are not culturally neutral; they embody preconceptions of what counts and what not, derived from the rationalistic focus on individualities and explicit measurements, and therefore largely ignoring the substance of relationships that cannot be described explicitly or measured. Mignolo's concept of "border thinking" captures how marginalised epistemologies resist Western dominance by asserting their unique insights, proposing pluriversal ethics as a framework in which multiple, equally valid worldviews coexist rather than function as deviations from a Western norm (Mignolo, 2012). Mohamed, Png and Isaac (2020) draw on Mignolo's concept of a "pluriversal epistemology of the future" to argue that decolonising AI requires dismantling harmful power asymmetries and turning toward a framework that, unlike universalisms, acknowledges and supports a wider radius of socio-political, ecological, cultural, and economic needs. Sethy (2026) argues that indigenous knowledge systems embody centuries of cumulative experiential wisdom deeply rooted in

ecological, spiritual, and communal contexts, and that as AI becomes integral to global systems of governance, development, and resource allocation, there is a pressing need to interrogate its epistemological underpinnings rather than allow them to pass unremarked. Santos (2014) calls for an "epistemology of the South" to recover silenced knowledges and challenge the assumption that what is produced in the Global North is necessarily universal. This is directly relevant to the kind of work The Club of Rome and the Fifth Element undertake on ecological knowledge and learning systems that situate human communities within rather than above the natural world. The knowledge embedded in those systems should shape AI.

International precedent for coordinated restraint. A coordinated pause on the riskiest frontier AI development is often dismissed as unrealistic. The historical record is more mixed than that dismissal suggests. Scientists have imposed this kind of restraint on themselves before: in 1974, a group of leading molecular biologists led by Paul Berg called for a voluntary moratorium on recombinant DNA research until its risks could be assessed, and the moratorium was universally observed until the 1975 Asilomar Conference set out the safety conditions under which the research could resume (Berg et al., 1974; Berg et al., 1975). The 1987 Montreal Protocol achieved near-universal ratification to phase out ozone-depleting substances despite significant short-term costs to signatories; by 2022, it had eliminated nearly 99% of those substances and put the ozone layer on track for full recovery by mid-century (World Meteorological Organization et al., 2022), and it has itself been analysed as a precedent for coordinating on other technology-driven catastrophic risks, including AI's environmental externalities (Chin, 2025).

A separate case-study review of nuclear, chemical, and biological weapons agreements similarly finds that states can agree to forgo a militarily and commercially useful technology when the case for collective restraint is clear, though such agreements typically take years of verification-building and are vulnerable to non-compliance (Wasil et al., 2024). None of the treaties themselves were built for a technology like AI, and each took years of prior scientific consensus-building before agreement was reached. But they undercut the claim that a coordinated slowdown is inherently utopian. The harder questions concern verification and enforcement, and how to prevent a pause from simply freezing existing power asymmetries in place. Bringing to the fore proposals of coordinated restraint would also help to question who is legitimate to shape processes that affect the common good, moving them from the capture by private hands to the agency of governments and citizens at large.

A deeper question. Furthermore, a deeply philosophical question needs to be raised. The rhetoric of AI, as with previous instances of digitalisation, is largely based on the promise to change the world for the better and to improve our capacity to understand how life works and hence, potentially, to develop responses for a regenerative approach to human wellbeing on a healthy planet. But could AI and more computing power, as promised by quantum technologies, provide the capacity to “crack the code of life?” Many scientific problems that are now intractable could indeed become tractable. But many others are undecidable problems that cannot be solved with more computing power. They come from the irreducible complexity of life, and this comes from free will (which also applies to non-humans). For the Club of Rome, which historically sits at the intersection of science and humanism, this matter is crucial: what if the expansion of technology was actually reducing, instead of increasing, our perception and attunement to life? What if, by adopting with enthusiasm the technical effectiveness of AI, our conception of the world would be reduced to mechanistic approaches, not only dehumanised but even devoid of the substance of life?

6. Conditions for Transformation

The transformative possibility of AI in service of people and nature is real. However, it is not automatic, and the dominant developments of AI are not following this trajectory currently. The regulatory and governance responses that have been adopted up to now either put the pressure on humans to adapt, as if the AI we have was a natural phenomenon as inevitable as the rotation of the Earth, or try to minimise some of the bad consequences (f.i. on privacy), without questioning the legitimacy of current developments. What is missing is a third type of approach, the transformative

one, which radically questions AI and at the same time creates the conditions for its use for the common good of humanity and the planet.

A caution belongs here about the terms on which this debate is being conducted. Much of the current urgency around AI is shaped by a period of financial exuberance: the Bank for International Settlements' 2026 Annual Economic Report compares today's AI capital expenditure to the canal, railway and dot-com manias, and notes it already outpaces the earnings of the firms making it (Bank for International Settlements, 2026). This matters for how this debate is framed. Reacting to every announcement from a handful of well-resourced AI companies risks amplifying their influence and mistaking market narrative for inevitability. The governance questions here are urgent regardless of where AI valuations go. What is contingent is the path AI deployment takes, since it rests on a financing structure that may itself prove unstable. None of this changes what the transformative approach requires.

Getting there requires naming the political economy clearly: the near-total monopolisation of the computational infrastructure required to train and deploy AI systems, the extractive logic of data collection, the anti-labour direction of automation, and the standards game being played without most of the world at the table. It requires building and resourcing alternative governance architectures. And it requires insisting on the prior question, *need it for what?* as a democratic one, open to a genuine range of answers, including "not this kind of AI, not here, not now."

The Club of Rome's tradition of systems thinking is well-positioned to make this argument with rigour and without technological fatalism. The conversation about AI governance is happening now. The window for shaping it toward genuinely different ends needs to be wide open.

References

Acemoglu, D., Autor, D. & Johnson, S. (2023). How AI can become pro-worker. *VoxEU/CEPR*, October 2023. https://cepr.org/voxeu/columns/how-ai-can-become-pro-worker

Aczel M., Chamanara S., Matin M., Farsi A., Marwala T., Madani K. (2026) Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints, United Nations University Institute for Water, Environment and Health (UNU-INWEH), Richmond Hill, Ontario, Canada, doi:10.53328/INR26RMA002

African Union (2022). *Continental Data Policy Framework*. African Union Commission. https://au.int/en/documents/20220728/au-data-policy-framework

African Union (2024). *Continental Artificial Intelligence Strategy*. African Union Commission. https://au.int/en/pressreleases/20250404/african-union-champions-ai-africas-socioeconomic-transformation

Bank for International Settlements (2026). Annual Economic Report 2026. Basel: BIS, released 28 June 2026. https://www.bis.org/publ/arpdf/ar2026e.htm

Berners-Lee, T. / World Wide Web Foundation (2019). Contract for the Web. https://contractfortheweb.org/

Berg, P., Baltimore, D., Boyer, H. W., Cohen, S. N., Davis, R. W., Hogness, D. S., Nathans, D., Roblin, R., Watson, J. D., Weissman, S. & Zinder, N. D. (1974). Potential biohazards of recombinant DNA molecules. Science, 185(4148), 303. https://doi.org/10.1126/science.185.4148.303

Berg, P., Baltimore, D., Brenner, S., Roblin, R. O. & Singer, M. F. (1975). Summary statement of the Asilomar conference on recombinant DNA molecules. Proceedings of the National Academy of Sciences, 72(6), 1981–1984. https://doi.org/10.1073/pnas.72.6.1981

Berners-Lee, T. (2019, March 12). "30 Years On, What's Next #ForTheWeb?" World Wide Web Foundation. Reproduced via Library Journal's infoDOCKET (original webfoundation.org publication no longer live, following the World Wide Web Foundation's 2025 closure): https://www.infodocket.com/2019/03/12/30-years-on-whats-next-fortheweb-by-sir-tim-berners-lee/

Bretton Woods Project (2026, April). Financing caring economies beyond extraction: Feminist perspectives on tax justice, AI and development. https://www.brettonwoodsproject.org/2026/04/financing-caring-economies-beyond-extraction-feminist-perspectives-on-tax-justice-ai-and-development/

Cave, S., Dihal, K., Drage, E. & McInerney, K. (2023). Who makes AI? Gender and portrayals of AI scientists in popular film, 1920–2020. Public Understanding of Science, 32(6), 745–760. https://doi.org/10.1177/09636625231153985

Chin, Z. S. (2025). Dimensional Characterization and Pathway Modeling for Catastrophic AI Risks. arXiv:2508.06411. https://arxiv.org/abs/2508.06411

Couldry, N. & Mejias, U.A. (2019). *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford University Press. Preface and Chapter 1 available at: https://law.unimelb.edu.au/__data/assets/pdf_file/0008/3290381/Couldry-and-Mejias-Preface-and-Ch-1.pdf

Dauvergne, P. (2022). Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy*, 29(3), 696–718. https://doi.org/10.1080/09692290.2020.1814381

Eloundou, T., Manning, S., Mishkin, P. & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of large language models. *Science*, 384(6702), 1306–1308. https://doi.org/10.1126/science.adj0998

Foster, C. et al. (2023). Cited in Sandbrook, C. (2024) on AI and smart agriculture reinforcing power hierarchies among smallholder farmers.

Gray, M. L. & Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.

International Energy Agency. (2025). Energy and AI. IEA. https://www.iea.org/reports/energy-and-ai

Kumar, S. & Choudhury, S. (2022). Gender and feminist considerations in artificial intelligence from a developing-world perspective, with India as a case study. Humanities and Social Sciences Communications, 9, 31. https://www.nature.com/articles/s41599-022-01043-5

Lee, M. & Dederichs, B. (2025). Viewing AI from the bottom: A feminist political economic critique of a technocentric pedagogy. Feminist Pedagogy, 6(3), Article 5. https://digitalcommons.calpoly.edu/feministpedagogy/vol6/iss3/5

Leverhulme Centre for the Future of Intelligence (n.d.). Gender & Technology Research Project. University of Cambridge. https://www.lcfi.ac.uk/research/project/ai-and-gender

Mignolo, W.D. (2012). *Local Histories/Global Designs: Coloniality, Subaltern Knowledges, and Border Thinking*. Princeton University Press.

Mohamed, S., Png, M.T. & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659–684. https://doi.org/10.1007/s13347-020-00405-8

Noy, S. & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. https://doi.org/10.1126/science.adh2586

OECD (2022). OECD Framework for the Classification of AI Systems. OECD Digital Economy Papers, No. 323. Paris: OECD Publishing. https://doi.org/10.1787/cb6d9eca-en

OECD (2024). Explanatory Memorandum on the Updated OECD Definition of an AI System. OECD Artificial Intelligence Papers. Paris: OECD Publishing, 5 March 2024. https://www.oecd.org/en/publications/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_623da898-en.html

Open Future (2024). Digital Commons as Providers of Public Digital Infrastructure. https://openfuture.eu/wp-content/uploads/2024/11/241113_Digital-Commons-as-Providers-of-Public-Digital-Infrastructures.pdf

Png, M.-T. (2022). At the Tensions of South and North: Critical Roles of Global South Stakeholders in AI Governance. *FAccT '22: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, 1434–1445. https://doi.org/10.1145/3531146.3533200

Reynolds, S.A. et al. (2025). The potential for AI to revolutionize conservation: a horizon scan. *Trends in Ecology & Evolution*, 40(2), 191–207. https://doi.org/10.1016/j.tree.2024.11.013

Rhee, J. (2023). From ELIZA to Alexa: Automated Care Labour and the Otherwise of Radical Care. In J. Browne, S. Cave, E. Drage & K. McInerney (Eds.), Feminist AI: Critical Perspectives on Algorithms, Data, and Intelligent Machines (Ch. 10). Oxford University Press. https://academic.oup.com/book/55103

Ricaurte, P. (2019). Data epistemologies, the coloniality of power, and resistance. *Television & New Media*, 20(4), 350–365. https://doi.org/10.1177/1527476419831640

Sandbrook, C. (2025). Beyond the hype: Navigating the conservation implications of artificial intelligence. Conservation Letters, 18(1), e13076. https://doi.org/10.1111/conl.13076

Santos, B. de S. (2014). *Epistemologies of the South: Justice Against Epistemicide*. Paradigm Publishers.

McGill University (2025). AI has untapped potential to advance biodiversity conservation. *ScienceDaily*, 6 May. Reporting on Pollock, L.J. et al. (2025), *Nature Reviews Biodiversity*, 1(3), 166. https://doi.org/10.1038/s44358-025-00022-3. Press release available at: https://www.sciencedaily.com/releases/2025/05/250506170938.htm

Sambra, A. V., Mansour, E., Hawke, S., Zereba, M., Greco, N., Ghanem, A., Zagidulin, D., Aboulnaga, A. & Berners-Lee, T. (2016). Solid: A Platform for Decentralized Social Applications Based on Linked Data. MIT CSAIL & Qatar Computing Research Institute. https://cs.brown.edu/courses/csci2390/2023/readings/solid.pdf

Sethy, M. (2026). Artificial intelligence and epistemic justice: a decolonial turn through indigenous knowledge systems. *AI & Society*, 41, 5401–5410. https://doi.org/10.1007/s00146-026-02936-8

Solid Project (n.d.). Solid: Take back control of your data. Inrupt / MIT / World Wide Web Consortium initiative, founded by Tim Berners-Lee. https://solidproject.org/

University College London. (2026, June 4). Blood protein signature can predict lung cancer risk years earlier [Press release; reporting on a study in Cell]. https://www.ucl.ac.uk/news/2026/jun/blood-protein-signature-can-predict-lung-cancer-risk-years-earlier

Vijayakumar, A. (2024). *AI Ethics for the Global South: Perspectives, Practicalities, and Policy*. RIS Discussion Paper No. 296. Research and Information System for Developing Countries. https://www.ris.org.in/sites/default/files/Publication/DP-296-Anupama-Vijayakumar.pdf

Wasil, A. R., Barnett, P., Gerovitch, M., Hauksson, R., Reed, T. & Miller, J. W. (2024). Governing dual-use technologies: Case studies of international security agreements and lessons for AI governance. arXiv:2409.02779. https://arxiv.org/abs/2409.02779

Wilber, K. (2000). Integral psychology: Consciousness, spirit, psychology, therapy. Shambhala.

World Meteorological Organization, United Nations Environment Programme, NOAA, NASA & European Commission (2022). Scientific Assessment of Ozone Depletion: 2022, Executive Summary. GAW Report No. 278. https://csl.noaa.gov/assessments/ozone/2022/executivesummary

Yeung, B. (2026, July 23). The rise of the military-technology complex. *Bulletin of the Atomic Scientists*. https://thebulletin.org/2026/07/the-rise-of-the-military-technology-complex/