



25
YEARS

IT for Change



Towards Regenerative AI: Frames for Inclusive, Indigenous, and Intentional Innovation





This report is a synthesis of the insights emerging from the two-day global dialogue ‘Towards Regenerative AI: Frames for Inclusive, Indigenous, and Intentional Innovation’, an official Pre-Summit Event of the India AI Impact Summit 2026. Held from 30 October to 1 November 2025 at the Infosys Science Foundation (ISF), Bengaluru, the event was organized by IT for Change in partnership with the Friedrich-Ebert-Stiftung, with support from the European Commission.



Credits

Conference design: Anita Gurumurthy, Nandini Chami, Shobhit S., and Sadhana Sanjay

Report authors: Shobhit S., Abhineet Nayyar, and Eshani Vaidya

Editorial guidance and review: Anita Gurumurthy

Copyediting: Susan Sreemala

Proofreading: Sadaf Wani and Susan Sreemala

Illustrations and publication design: Harikrishnan B.

Publication design guidance: Chinmayi Arakali

Conference support: Gautam B.K.

Rapporteurs: Abhineet Nayyar, Amoha Sharma, Eshani Vaidya, Marzia Ibrahim, Merrin Muhammad Ashraf, Nandini Chami, Saloni Mishra, Sherina Poyyail, and Susan Sreemala

Contents

1. Introduction	6
Context.....	6
Approach and methodology	9
2. Challenges: What’s Broken in the Current AI Paradigm	12
Structural concentration and power in the AI economy.....	12
Distributed harms across AI value chains and accountability gaps.....	14
Epistemic exclusion and the devaluation of plural knowledge systems	16
Economic fragmentation and constrained developmental pathways	17
Ecological blind spots and the lack of environmental safeguards	19
3. Toward a New Deal for the AI Economy.....	22
Meaningful and Dignified Work.....	23
Ensuring transparency and due process in algorithmic management	24
Recognizing invisible labor in data and AI value chains.....	25
Strengthening collective bargaining and worker voice in AI design and deployment	26
Diversified Economies.....	28
Addressing structural deficits in business models through antitrust law	29
Rethinking digital sovereignty through productive capacity.....	30
Expanding viable AI innovation pathways	31
Pluralistic Knowledge Societies	33
Building a communitarian understanding of knowledge production and ownership	34
Reorienting research institutions and knowledge infrastructures	35
Enabling communities to co-create and co-govern AI tools	36
Planetary Flourishing.....	37
Foregrounding restraint and intentional limits in AI innovation.....	38
Embedding environmental accountability across AI value chains.....	39
Advancing distributive justice in energy and AI transitions	39
4. Levers of Change for Just AI Futures	42
Techno-design choices	42
Discursive shifts and evidence collection	43
Law and institutions	43
Incentives and market shaping	44
Democratic participation and worker power	45
Levers of change for just AI futures	46


Participants

The individuals listed below participated in the Regen AI dialogue and contributed insights, critiques, and provocations that shaped this synthesis. Listing does not imply endorsement of every position articulated in the report.

1. Aditi Surie, Senior Consultant – Academics & Research, Institute for Human Settlements
2. Alexandre Costa Barbosa, Project Coordinator – UniPeriferias
3. Amoha Sharma, Program Coordinator – Center for Global Digital Justice
4. Amrita Sengupta, Research Fellow – Centre for Technology and Society, FGV Law School
5. Anuradha Ganapathy, PhD Student – University of Manchester
6. Balaji Parthasarathy, Professor – International Institute of Information Technology
7. Barbara Olive Maria Adams, Board Chair – Global Policy Forum
8. Bhagyashree Joshi, India Community Manager – Open Supply Hub
9. Brian Chen, Policy Director – Data & Society
10. Chee Yoke Ling, Executive Director - Third World Network
11. Deepti Bharthur - Independent Researcher
12. Ganesh Gopal, Co-founder – Centre for Development Initiatives (CDI) & Tinkerhub Foundation
13. Gaurav Godhwani, Executive Director & Co-founder – CivicDataLab
14. Gayathri Vasudevan, Chief Impact Officer – Sambhav Foundation
15. Jai Vipra, PhD student – Cornell University
16. James Henry Farrar, Director – Worker Info Exchange
17. José Renato Laranjeira de Pereira, PhD Researcher - Sustainable AI Lab, Universität Bonn
18. Laura Elizabeth Mann, Associate Professor, International Development Department – London School of Economics
19. Mandvi Kulshreshtha, Senior Program Adviser - Friedrich-Ebert-Stiftung, India Office
20. Maria Soledad Vogliano, Senior Researcher and Campaigner - ETC Group
21. Mirai Chatterjee, Chairperson – SEWA Cooperative Federation
22. Nachiket Udapa, Member – Mazdoor Kisan Shakti Sangathan (MKSS)
23. Padmini Ray Murray, Founder – Design Beku
24. Paola Ricaurte Quijano, Leader – Feminist AI Research Network, Latin America and the Caribbean
25. Piush Antony, Social Policy Specialist - Lucknow Field Office of UNICEF


26. Pooja Kapahi, Director for Youth, UNI Asia & Pacific Region – UNI Global Union
27. Poruri Sai Rahul, Chief Executive Officer – FOSS United
28. Prerakkumar Mukeshkumar Shah, Senior Project Manager – Vrutti, Catalyst Group
29. Rachada Buranasiri, Researcher – Institute of Asian Studies, Chulalongkorn University
30. Rafael Augusto Ferreira Zanatta, Director – Data Privacy Brasil
31. Rajiv Kumar, Project Manager – Rosa Luxemburg Stiftung
32. Ramya Chandrasekhar, Researcher – French National Centre for Scientific Research (CNRS), Paris
33. Revathi Sharma Kollegala, Technology Strategist – Center for International Forestry Research and World Agroforestry (CIFOR-ICRAF)
34. Roberto Remo Bissio Staricco, Coordinator– Social Watch
35. Ruchi Agarwal, Chief Executive Officer – National Insurance VimoSEWA Cooperative Ltd
36. Sabina Dewan, Founder and Executive Director – JustJobs Network
37. Shivangi Rai, Deputy Coordinator and Lead: Technology, Health & Society – Centre for Health Equity, Law & Policy, Indian Law Society
38. Shyam Krishna Raja Gopalan, Research Leader – RAND Europe
39. Sofia Beatriz Scasserra, Associate Researcher - TNI
40. Sreechand Tavva, Product Manager - Arghyam
41. Sreenath Namboodiri, Lawyer and Policy Researcher - Intellectual Property Law, Human Rights, and Public Policy
42. Suhel Bidani, Lead, Digital and AI - Gates Foundation
43. Titiksha Vashist, Co-founder, Lead Researcher - The Pranava Institute
44. Usha Raman, Professor - University of Hyderabad
45. V Abhishek, National Project Officer – UNESCO
46. IT for Change Team: Abhineet Nayyar, Anita Gurumurthy, Anusha Sharma, Chandra, Eshani Vaidya, Gurumurthy Kasinathan, Marzia Ibrahim, Merrin Muhammad Ashraf, Nandini Chami, Natasha Susan Koshy, Sadhana Sanjay, Saloni Mishra, Sherina Poyyail, Shobhit S., and Susan Sreemala

1. Introduction*



Context

The release of OpenAI’s ChatGPT – followed by a wave of competing models and downstream applications – marked a significant inflection point in the contemporary Artificial Intelligence (AI) paradigm. These systems have emerged within a political economy shaped by platform capitalism, where data maximization, resource capture, and top-down platformization have been naturalized as inevitable features of the digital economy. In this context, large language models (LLMs) are not just unsustainable in the long run; their growth relies on the relentless expansion of markets and their infrastructures. From the demand for ever more granular datasets and escalating compute capabilities, to rapid deployment in safety-critical areas like healthcare or education, the dominant trajectory has repeatedly favored extraction and appropriation over equitable distribution and public value. Platform capitalism has, in effect, normalized a wild-west economy in which scale, speed, and enclosure are treated as virtues instead of risks.




At the core of this trajectory lies a highly concentrated industrial structure. A small number of technology firms exercise near-unilateral control over key layers of the AI stack, including data, compute infrastructure, cloud platforms, and model development pipelines.

At the core of this trajectory lies a highly concentrated industrial structure. A small number of technology firms exercise near-unilateral control over key layers of the AI stack, including data, compute infrastructure, cloud platforms, and model development pipelines. This concentration enables data hoarding and reinforces winner-takes-all dynamics, where early advantages translate into almost insurmountable market power. Rather than markets efficiently selecting the “best” technologies, monopoly conditions shape what counts as innovation in the first place, narrowing the range of viable alternatives and locking in particular technological pathways.

* This introduction is authored by Anita Gurumurthy, Abhineet Nayyar, Shobhit S., and Eshani Vaidya, on behalf of IT for Change.

These dynamics are further reinforced by the dominance of speculative investment logics tied to monopoly capital. Capital allocation in AI overwhelmingly favors scale-driven models that promise market capture rather than demonstrable productivity gain or social value. As critics of platform monopolies have noted, claims surrounding technological superiority are often retrospective justifications for dominance achieved through concentration. Under these conditions, it is increasingly difficult for alternative, smaller-scale, or public interest-oriented digital innovations to emerge or survive. This distortion also spills over into public sectors, where AI systems are frequently adopted under conditions of vendor dependence and information asymmetry, embedding proprietary technologies into welfare delivery, criminal justice, or public administration, without adequate scrutiny or guardrails.

For Global South countries, these structural features of the contemporary AI economy translate into heightened asymmetries and constrained developmental choices. Lacking comparable data reserves, compute capacity, or bargaining power, many countries find themselves positioned as downstream adopters within AI value chains dominated by firms based in the US and China. Pathways to economic self-determination appear increasingly unclear when domestic digital infrastructures, standards, and innovation ecosystems are deeply dependent on external platforms. This dependency is not merely technological, but also geopolitical and economic, influencing what forms of regulation, industrial policy, and public investment are deemed legitimate and feasible.




The rapid build-out of data center farms and compute facilities has accelerated the exploitation of land, energy, and other natural resources by multinational firms, often aided by tax exemptions, fast-tracked approvals, and weak local accountability.

The material consequences of this dependency are already visible. The rapid build-out of data center farms and compute facilities has accelerated the exploitation of land, energy, and other natural resources by multinational firms, often aided by tax exemptions, fast-tracked approvals, and weak local accountability. At the same time, the deployment of opaquely developed AI systems without adequate sectoral integrity or cultural context exposes already marginalized cultures and knowledge systems to renewed forms of epistemic erasure, reproducing colonial power relations under the guise of technological progress.

Compounding these challenges, governments across much of the Global South face persistent capacity gaps and external pushback in regulating key domains such as competition, intellectual property, and taxation. These are precisely the policy levers required for such states to shape AI markets in line with domestic priorities, yet they remain constrained by geopolitical pressures, trade regimes, and asymmetries in technical expertise.

Against this backdrop, partial successes with digital public infrastructure (DPI)-based innovation in countries such as India and Brazil demonstrate that the extractive and monopolistic characteristics of today's AI economy are not inevitable. With coordinated and deliberate interventions, an alternative paradigm of 'Regenerative AI' (Regen AI) is both achievable and urgently necessary. Contrary to today's globalized, privatized, and market-led AI industry, Regen AI is rooted in the right of all societies to pursue flourishing futures.



With coordinated and deliberate interventions, an alternative paradigm of 'Regenerative AI' (Regen AI) is both achievable and urgently necessary. Contrary to today's globalized, privatized, and market-led AI industry, Regen AI is rooted in the right of all societies to pursue flourishing futures.

By conceptualizing AI as a socio-technical assemblage rather than a purely computational system, it opens space for alternative trajectories that:

- encourage diverse economic pathways over a universal global model;
- validate pluralistic lifeworlds and knowledges, rather than a singular ideal of progress;
- galvanize vibrant local livelihoods, instead of jobs that hollow out local capacity;
- contribute to public and social value in addition to private returns;
- enable local oversight and control in place of inscrutable systems; and
- promote ecological sustainability, without privileging global scale as an end in itself.

Importantly, Regen AI can be viewed as a reorienting prism built on three interrelated dimensions:



Inclusivity, i.e., centering democratic innovation pathways, through fair distribution of AI dividends inside and across borders.



Indigeneity, i.e., situating AI innovation in the context of local economies, through sovereign, communitarian, and worker-led models of data and AI stewardship, and a critical evaluation of ecological effects.



Intentionality, i.e., shaping AI ecosystems to serve the common good, through social and market regulation that aims to maximize public interest, including the articulation of limits, explicit no-go zones, and long-term sustainability measures.

Achieving such a departure from the current AI paradigm requires a collaborative effort towards re-imagining our collective social contract and transforming the political and economic institutions to govern digital technologies in the public interest and towards ecologically sustainable futures. Put differently, it necessitates a **New Deal for the AI Economy**.

Approach and methodology

This report emerges from a two-day dialogue held on 31 October and 1 November 2025 at the Infosys Science Foundation (ISF), Bengaluru, which brought together over sixty researchers, technology practitioners, labor organizers, and policy professionals from India, the European Union, Southeast Asia, Latin America and other regions. While acknowledging the centrality of existing critiques of the AI paradigm, the dialogue was designed to articulate the contours of an alternative, people-centered AI innovation framework, with a shared baseline understanding of the challenges.

The first day focused on identifying structural challenges posed by the current AI paradigm to inclusive societies and Indigenous economies. Participants were organized into five thematic Community Labs, aligned with their preferences and areas of expertise:

CL1: Addressing Data and AI Harms

CL2: Building South-led AI Economies

CL3: Designing AI Models for Local Context and Control

CL4: Tackling Big Tech's Market Dominance

CL5: Making the Right Choices for Sustainable AI Economies

Each Lab was guided by a common set of questions:

What is missing in the mainstream discourse?

What should be on the agenda?

What are the emerging areas of contestation?

What levers of influence can drive change?

Short catalyst presentations at the outset helped establish shared reference points and surface productive tensions across perspectives.

The Community Labs provided a diagnosis of the contemporary AI paradigm, foregrounding the dynamics shaping current trajectories of AI development and deployment.

The second day built on this foundation, shifting the focus from diagnosis to forward-looking pathways. Drawing on the insights generated during the first day, discussions were organized around four interrelated prongs as organizing lenses of a New Deal for the AI Economy:

Meaningful and Dignified Work

Diversified Economies


Pluralistic Knowledge Societies

Planetary Flourishing

The day began with a set of lightning talks offering conceptual provocations on each prong, followed by group discussions that explored social, political, economic, and technological pathways for change. Participants were asked to consider:

What would be the elements of an intentional innovation paradigm designed to advance meaningful work, diversified economies, pluralistic knowledge societies and respect for planetary boundaries?

This synthesis report captures the outcomes of the two-day dialogue as both a stocktaking exercise and a provocation for future action. It is structured in three parts: first, an articulation of the key structural challenges shaping the current AI paradigm; second, an elaboration of the four prongs of a New Deal for the AI Economy; and third, a set of cross-cutting levers of influence that could enable a transition toward a Regen AI paradigm grounded in inclusivity, indigeneity, and intentionality.



This synthesis report captures the outcomes of the two-day dialogue as both a stocktaking exercise and a provocation for future action. It is structured in three parts: first, an articulation of the key structural challenges shaping the current AI paradigm; second, an elaboration of the four prongs of a New Deal for the AI Economy; and third, a set of cross-cutting levers of influence that could enable a transition toward a Regen AI paradigm grounded in inclusivity, indigeneity, and intentionality.

The analysis presented in this report draws on insights from the Community Labs, plenary discussions, and lightning talks convened over the course of the dialogue. These inputs were synthesized through an iterative qualitative process undertaken by the IT for Change team, drawing on session notes, recordings, and rapporteur summaries.




Participants of the two-day global dialogue 'Towards Regenerative AI: Frames for Inclusive, Indigenous, and Intentional Innovation'

Rather than attributing positions to individual speakers or seeking consensus, the synthesis prioritized documentation of recurring themes, illustrative examples, and shared lines of concern, while also attending to points of divergence and tension across groups. This approach aimed to preserve the plurality of perspectives expressed during the dialogue while situating them within a coherent analytical framework relevant to policy debates on the AI economy.


This report is indebted to the participants of the dialogue as co-travellers in knowledge production. Their experiences, critiques, and situated insights constitute the analytical backbone of this work.

2. Challenges:

What's Broken in the Current AI Paradigm



This section captures the challenges that emerged most consistently across the Community Lab discussions, plenary sessions, and lightning talks during the dialogue. They do not constitute an exhaustive diagnosis of the AI ecosystem, nor do they imply that all participants agreed on every formulation presented here. Instead, these are to be seen as concerns and considerations that repeatedly surfaced across discussions, surrounding the political economy, governance arrangements, and material conditions shaping the current AI paradigm, especially as they affect Global South societies and economies.




Issues of concentration of power, extractive value chains, epistemic exclusion, constrained developmental pathways, and ecological blind spots were identified as central features of the dominant AI trajectory that undermine democratic innovation, local control, meaningful work, and environmental sustainability.

Taken together, these challenges help clarify what a shift toward a Regen AI paradigm would need to contend with. Issues of concentration of power, extractive value chains, epistemic exclusion, constrained developmental pathways, and ecological blind spots were identified as central features of the dominant AI trajectory that undermine democratic innovation, local control, meaningful work, and environmental sustainability. These challenges represent the systemic limits of the prevailing AI paradigm, and delineate the terrain on which alternative digital futures must be built. The section that follows builds on this diagnosis by outlining four prongs of a New Deal for the AI Economy, proposed as forward-looking responses to these limits and failures.

Structural concentration and power in the AI economy

A central challenge shaping the contemporary AI economy is the structural concentration of power. This concentration extends beyond firm size or market dominance to encompass control across layers of the AI stack, including compute infrastructure, capital, platforms, high-value datasets, and agenda-setting power, which influence the trajectory of AI development and deployment.



Concentration in the current AI economy extends beyond firm size or market dominance, to encompass control across layers of the AI stack, including compute infrastructure, capital, platforms, high-value datasets, and agenda-setting power.

The extreme concentration of compute and hardware capacity emerged as a defining feature of the current AI paradigm. Companies such as NVIDIA, whose market capitalization now exceeds the GDP of many countries, illustrate how control over semiconductor design and supply chains has become a geopolitical chokepoint. While advanced semiconductor manufacturing has shifted from North America to parts of East and Southeast Asia, this has not translated into a more equitable distribution of control over global AI supply chains. Instead, it is reinforcing global hierarchies, with resource-intensive manufacturing concentrated in specific regions, while strategic control and value remain elsewhere.

This infrastructural concentration is closely tied to speculative investment narratives, particularly around Artificial General Intelligence (AGI). **Massive capital expenditure on data centers and compute infrastructure is driven by projected returns from large-scale models**, even as their long-term viability and social value remain uncertain. Such bubbles risk deepening dependency by locking states and institutions, particularly in the Global South, into long-term infrastructural dependencies that narrow policy autonomy and limit alternative development pathways.

These dynamics **reproduce colonial patterns of value extraction**. Countries rich in primary commodities, including land, water, energy, and minerals, are increasingly positioned as sites for hyperscale data centers and extractive supply chains, while higher-value activities such as model development and monetization are concentrated in the Global North.

In response, state-led **localization and digital sovereignty strategies often prioritize provider substitution over structural control**, leaving underlying questions of governance and value capture unresolved. Examples such as ‘Indic AI models’ were cited to illustrate this concern. While intended to reduce dependence on the ‘Magnificent Seven’, many such initiatives focus primarily on training language models on training Indian datasets or deploying them through domestically branded platforms, even as they remain reliant on cloud infrastructure, compute resources or proprietary architectures controlled by Big Tech. As a result, they risk reproducing existing power relations unless accompanied by changes in ownership, governance and business models. Similarly, moves to promote domestic platforms were cited as instances where dependence on Big Tech is replaced by domestic monopolies, without strengthening public value creation, accountability, or commons-based infrastructures.

Participants stressed that **business models matter as much as ownership**. The dominance of targeted advertising and data profiling was cited as a core driver of platform power and service degradation, or ‘enshittification.’ From this perspective, participants agreed that breaking up large firms without dismantling extractive business models was inadequate.

Further, as AI systems are increasingly integrated into public services, including welfare, health and education, **states are becoming heavily reliant on privately controlled digital infrastructures for core public functions.** Public-private partnerships structured around non-disclosure clauses were cited as weakening transparency and public oversight, while national security exemptions and narratives further restrict disclosure by placing AI systems, datasets, and procurement decisions beyond civil society scrutiny.

Dominant firms and advanced economies also exercise agenda-setting power in shaping how global AI governance is framed. Debates tend to prioritize technical safety and ethical principles, while sidelining questions of ownership structures, business models and value capture. Voluntary standards and ethics frameworks were widely criticized as insufficient, particularly when they divert attention from more structural interventions relating to competition, public procurement, and infrastructure ownership.

“


Under the assumption of AGI, a lot of capital expenditure is being undertaken, primarily to benefit the ‘Magnificent Seven’ that have seen soaring market valuations over the last few years. Even if a bubble exists, there is a chance that it may not burst, and would just allow for user experience to get gradually worse. This is even likelier as unilaterally deployed and implemented AI takes over, further enshittifying core digital services.

Jai Vipra, PhD student at Cornell University

Distributed harms across AI value chains and accountability gaps

A second major challenge concerns the way AI-related harms are produced across fragmented and opaque value chains, while mechanisms for accountability and redress remain weak. Harms are generated across stages of data extraction, model development, infrastructure provisioning, and use, often spanning multiple jurisdictions and contractual layers. This makes responsibility difficult to assign and systemic harms easy to obscure.

At the upstream end of AI value chains, data extraction practices represent a significant source of harm. Data, including personal, community-based, cultural, or behavioral, is **frequently scraped without meaningful consent, clear purpose limitation, or benefit sharing.** Contributors whose data feeds AI systems often lack visibility into how it is aggregated, (re)used, or monetized.



Harms are generated across stages of data extraction, model development, infrastructure provisioning, and use, often spanning multiple jurisdictions and contractual layers. This makes responsibility difficult to assign and systemic harms easy to obscure

This constitutes a form of enclosure of the commons, where **social data and knowledge are converted into private assets**, while the individuals and communities that produce them remain unrepresented and uncompensated. This was repeatedly linked to asymmetries of power between data subjects and data-extracting firms.

Further downstream, **algorithmic systems increasingly mediate access to work, welfare, and public services** through opaque decision-making processes. In the platform economy, (semi)-automated systems determine task allocation, pricing, incentives, discipline, and dismissal. Participants cited cases of drivers and delivery workers facing account deactivation or pay change justified solely through algorithmic decisions. These systems enable wage suppression, intensify work through surveillance, and shift economic risk onto workers, particularly in contexts of labor oversupply. Similar automated decision-making in welfare delivery, for example, denies individuals and communities access to social security, despite them meeting eligibility criteria.

Labor harms extend well beyond platform-based work. **Outsourced labor underpinning AI systems, including data annotation and content moderation, is performed under precarious conditions**, disproportionately in the Global South. These forms of work are often poorly remunerated, psychologically taxing, and excluded from labor protections. Labor exploitation is also embedded in the material supply chains of AI, such as mining and electronics manufacturing, where workers face especially hazardous environments. Across these contexts, productivity gains generated by AI systems are captured by firms, and workers bear the social, physical, and psychological costs.

Public sector procurement and deployment decisions further exacerbate accountability gaps. AI systems are frequently introduced as efficiency-enhancing solutions, without adequate assessment of whether they are necessary or appropriate.

Public authorities often lack the technical or institutional capacity to scrutinize vendors' claims or meaningfully involve affected communities in decision-making. This produces a persistent gap in which technologies are deployed rapidly, while evidence and public oversight lag far behind.

Despite the scale and diversity of these harms, **mechanisms for accountability and redress remain underdeveloped.** Liability regimes are fragmented and weakened by contractual arrangements and trade secrets.

Harms are typically addressed only at the point of impact, through individual complaints or litigation, rather than through systemic accountability across the AI value chain. Even where impact assessments or audits are instituted, these are often limited to technical performance or narrow compliance metrics excluding labor conditions, social impacts and ecological costs.


“

The dictum ‘data is the new oil’ allows it to be fetishised and indiscriminately collected. As a result of this fetishisation, industrial policy is being targeted at setting up of data centres, and manufacturing of semi-conductors and other core elements of the value chain. Unfortunately, the capital-to-labour ratio in these industries remains highly skewed towards the former.

Balaji Parthasarathy, International Institute of Information Technology, Bangalore

Epistemic exclusion and the devaluation of plural knowledge systems

A third challenge concerns the ways in which contemporary AI systems reproduce epistemic exclusion by privileging certain forms of knowledge and marginalizing others. AI systems, far from being ‘neutral’, encode specific assumptions about what counts as valid knowledge, and whose perspectives are legible and actionable.



AI systems, far from being 'neutral', encode specific assumptions about what counts as valid knowledge, and whose perspectives are legible and actionable.

Dominant AI models are built around data-intensive and abstraction-driven epistemologies that misrepresent **contextual, relational and embodied forms of knowledge**. This is particularly problematic in settings where knowledge is transmitted orally or embedded in social practices or lived experiences. Examples cited included Indigenous agricultural practices, artisanal production, and community-based knowledge systems, which are often flattened into data inputs when incorporated into AI systems. In this process, meaning is lost and community-based values are obscured.

‘Localization’ and ‘decentralization’, often offered as fixes, do not automatically address such epistemic injustice. While decentralized or local systems are often framed as alternatives to global models, local contexts themselves are structured by power and hierarchy. As a result, AI systems designed to be locally relevant may still reproduce dominant local perspectives while marginalizing others.

Academic and research institutions emerged as another site of epistemic exclusion. Knowledge production in AI remains heavily concentrated in Western institutions, shaped by funding structures, publication incentives, and proprietary research models. Local knowledge systems are often treated as data sources rather than epistemic authorities, reinforcing extractive research relationships. International collaborations framed around ‘global challenges’ channel outputs into IP regimes that benefit firms and institutions in high-income countries, while limiting local learning and ownership.

The dominance of a small set of global languages in AI training data further constrains the ability of models to represent diverse worlds. Translation is an inherently political process, where concepts embedded in one cultural context cannot be seamlessly mapped onto another. Attempts to universalize knowledge through AI therefore reproduce older imperial logics, on which difference is erased in the name of standardization and scale.

Co-creation and participatory design emerged as responses to epistemic exclusion. Examples were shared of technology projects where communities participated not only as data providers but as co-designers with decision-making power over problem articulation, system design, and use. At the same time, such approaches are time-intensive and often in tension with prevailing funding models and innovation timelines.

“

People’s subjective experiences are not easy to understand, and harder to codify as an AI might require. However, academic training, at present, hardly includes aspects of socio-technical capacity building or human-centered design methods — which is especially critical for stakeholders like engineers and developers who create these tools in the first place and end up siding with the corporations’ thinking.

Rachada Buranasiri, Institute of Asian Studies,
Chulalongkorn University

Economic fragmentation and constrained developmental pathways

Current AI trajectories are reshaping developmental pathways in ways that fragment economies and constrain domestic value creation, particularly in the Global South. While AI is often framed as a general-purpose technology capable of driving productivity and growth, its prevailing modes of development and use are reinforcing structural asymmetries rather than enabling broad-based domestic economic transformation.

“

While digital sovereignty and South-led economies are important, the more important question is if these economies are state-driven or capital-driven. Whether they are coop-owned, producer-owned, or worker-owned, we will have to think in terms of socialised ownership and imagine an ecosystem that is built to support such initiatives.

Rajiv Kumar, Rosa Luxemburg Stiftung


A key concern is the **geographic separation between contribution and value capture** in AI value chains. Countries in the South contribute essential inputs to AI systems – data, labor, land, energy, and minerals – while higher-value activities such as model development and monetization remain concentrated in the North. This pattern limits opportunities for domestic industrial upgradation and fiscal returns.

This fragmentation is particularly concerning for small and medium enterprises (SMEs), cooperatives, and local producers. **Dominant AI ecosystems privilege scale, proprietary data access and capital-intensive infrastructure**, which systematically disadvantages smaller actors. As a result, domestic firms, especially in the Global South, are positioned primarily as downstream users of proprietary AI tools rather than as developers of technologies aligned with local sectoral needs.

The **absence of sufficient public investment in foundational infrastructures** arose as another key constraint. Inadequate public investment in shared compute facilities and open research infrastructures has increased dependence on private cloud providers and foreign platforms. Examples cited included public sector AI deployments reliant on vendor-controlled infrastructure or outsourced components of digital public infrastructure, limiting opportunities for domestic learning and institutional capacity building.

Participants also pointed to the **disconnect between AI skills development and broader industrial strategies**. Training and upskilling initiatives are often pursued in isolation from demand-side policies. Examples include specialized training programs producing graduates who struggle to find relevant employment because domestic firms lack the capacity or incentives to absorb those skills. In this context, skills development risks serving global labor markets more than domestic economic transformation.

Financial and investment dynamics further compound these challenges. **AI investment is profoundly shaped by venture-driven logics that favour rapid scaling and infrastructure-heavy projects**. This is crowding out alternative innovation models, including cooperative platforms, public-interest technologies, and sector-specific applications that may deliver social or developmental value.



AI investment is profoundly shaped by venture-driven logics that favour rapid scaling and infrastructure-heavy projects. This is crowding out alternative innovation models, including cooperative platforms, public-interest technologies, and sector-specific applications that may deliver social or developmental value.

Crucially, trade and investment regimes are also constraining developmental policy space. Digital trade agreements and intellectual property regimes can limit states' ability to regulate data use, impose local content or technology transfer requirements, or leverage public procurement to support domestic firms. These constraints are particularly acute for countries pursuing Foreign Direct Investment (FDI)-led development strategies, where the imperative to attract investment narrows the scope for asserting public-interest conditions or long-term industrial goals. There was broad convergence around the view that economic fragmentation is being produced by policy choices that prioritize integration into global AI markets over domestic capability building. Without targeted public investment in shared infrastructure, strategic procurement, and sustained support for SMEs and cooperatives, AI adoption is likely to entrench low-value participation rather than enable diversification.

Ecological blind spots and the lack of environmental safeguards

The absence of consideration of environmental impacts of AI was highlighted as a major blind spot in current governance frameworks. While AI is frequently framed as a tool for climate action, its own ecological footprint remains poorly disclosed and largely externalized to the Global South. This reflects a failure to treat AI as a materially grounded industrial system, dependent on energy, water, land, minerals, and global hardware supply chains, rather than as an abstract digital technology.

“

The critical mineral value chain is necessary for the current AI paradigm, which makes it important not just to counter BigTech dominance, but also to heavily involve actors like the state and environmental groups, as well as hitherto ignored constituencies like farmers and workers. We need a green economy, yes, but we also must make a just economy.

Chee Yoke Ling, Third World Network

A recurring concern is the **rapid expansion of hyperscale data centers**, particularly in regions already facing land, water, and energy stress. Data centers are being promoted as development projects in countries such as India, Brazil, Kenya, and Argentina, often without environmental impact assessments or meaningful community consultation. These facilities compete directly with agriculture, housing, and public utilities for scarce resources, while delivering limited local employment or long-term economic benefits. Environmental approvals are often fragmented or procedural, failing to account for cumulative impacts.

Energy demand emerged as a central concern. **Data center operations and large-scale model training require enormous electricity supply**, placing pressure on national grids. Estimates highlighted that data center electricity consumption is comparable to that of entire countries. Participants warned that AI-driven demand risks locking in dependency on fossil fuels, even as governments commit to energy-transition targets.

Water use was identified as another largely invisible impact. **Processes of data center cooling and semiconductor fabrication consume vast quantities of freshwater**, often in drought-prone regions. These withdrawals are rarely disclosed in a transparent manner, and regulatory oversight is weak. Water competition linked to AI infrastructure risks intensifying local conflict, particularly where communities lack access to grievance redress mechanisms.


“

The material reality of AI is such that a boost in its current trajectory alters the society's metabolism in relation to nature. Hyperscale data centers compete with agricultural land use and protected territories, their energy needs compete with social requirements, and everything from mineral extraction to data center cooling increases the AI economy's carbon footprint.

Maria Soledad Vogliano, ETC Group

Upstream extraction of critical minerals was repeatedly raised as an integral yet under-acknowledged component of AI's environmental footprint. The mining of lithium, cobalt, copper, and rare earths is essential to AI hardware supply chains, with documented links to deforestation, toxic pollution, biodiversity loss, and displacement of Indigenous and rural communities. Moreover, in some contexts, mining for 'critical minerals' is exempted from standard environmental safeguards or consultation requirements.

Participants were critical of existing environmental standards and safeguards. **Efficiency standards, carbon disclosures, and voluntary Environmental, Social, and Governance (ESG) commitments are largely performative**, functioning as reputational tools rather than meaningful constraints. Such standards risk normalizing continued growth by making AI appear "greener". In particular, carbon accounting mechanisms tend to narrow environmental impact to emissions, while ignoring water and land use, material extraction, toxic chemicals and e-waste.



Extant governance frameworks rarely investigate whether specific AI applications justify their ecological cost.

Another gap identified was the **absence of purpose-based environmental scrutiny**. Extant governance frameworks rarely investigate whether specific AI applications justify their ecological cost. Calls were made to explicitly identify non-use or low-use cases, especially where AI deployment delivers marginal social benefit. This included questioning the necessity of large-language models and advocating for smaller, task-specific and localized systems where appropriate.

CHAPTER SUMMARY


Structural challenges shaping the AI economy

The challenges identified across the dialogue point to a set of interlocking structural distortions shaping the contemporary AI economy. These challenges underscore that the dominant AI paradigm is not neutral or inevitable, but the outcome of political, economic, and institutional choices that demand structural rethinking rather than incremental correction.

- **Structural concentration of power across the AI stack** — including compute, capital, platforms, data, and agenda-setting — has enabled a small number of firms to shape innovation trajectories, extract value, and constrain alternative models of development and governance.
- **Distributed harms across fragmented value chains** mean that AI-related harms are produced across data extraction, model development, infrastructure provisioning, and deployment, while responsibility, liability, and redress remain weak, opaque, and difficult to assign.
- **Epistemic exclusion and the devaluation of plural knowledge systems** arise when dominant AI models privilege data-intensive, standardized epistemologies, rendering Indigenous, local, relational, and embodied forms of knowledge invisible or extractable.
- **Economic fragmentation and constrained developmental pathways**, particularly in the Global South, reflect a widening gap between contribution and value capture, where data, labor, land, and energy are supplied locally while higher-value activities and rents accrue elsewhere.
- **Ecological blind spots in AI governance** obscure the material and environmental costs of AI systems—including energy use, water consumption, mineral extraction, and land competition—while sustainability narratives are often mobilized to justify continued expansion rather than enforce ecological limits.



3. Toward a New Deal for the AI Economy



This section builds on the challenges outlined in the previous section to articulate a constructive response. The challenges cumulatively point to the conclusion that the dominant AI paradigm is structurally misaligned with goals of social justice, democratic governance, and ecological sustainability. These failures are not simply outcomes of rapid technological change, but the result of specific political, economic, and institutional choices that shape how AI is developed, deployed, and governed.

The four prongs of the New Deal for the AI Economy respond to this diagnosis by shifting the focus from harm mitigation to structural transformation. They distil a set of normative and institutional directions that emerged across the dialogue as necessary conditions for a Regen AI paradigm. Each prong — Meaningful and Dignified Work, Diversified Economies, Pluralistic Knowledge Societies, and Planetary Flourishing — offers a distinct but interrelated lens for rethinking AI governance, anchoring technological development in human dignity, local capacity, epistemic plurality, and ecological limits.

The four prongs of the New Deal for the AI Economy

Meaningful & Dignified Work	
Diversified Economies	
Pluralistic Knowledge Societies	
Planetary Flourishing	



**Meaningful and
Dignified Work**

Meaningful and dignified work emerged across the dialogue as a foundational pillar of a New Deal for the AI Economy. Rather than a downstream concern to be addressed through reskilling programs, employability schemes or compensatory welfare measures, it is a core organizing principle for how AI systems are designed, deployed, and governed. The organization of work provides a critical vantage point to assess many of the structural failures of the current AI paradigm. In this context, meaningful and dignified work refers to work organized in ways that sustain dignity, predictability, agency, and collective power in an economy increasingly structured by data-driven decision making and work management. It foregrounds livelihoods and social reproduction, rather than abstract productivity gains or efficiency metrics.

Discussions highlighted that labor precarity in AI-mediated economies is not a technological inevitability, but a ‘political project’. What is often framed as ‘disruption’ reflects deliberate institutional and political choices, surrounding labor classification, deregulation, platform business models, and the systematic weakening of collective bargaining. From this perspective, meaningful work functions as a governing lens for AI adoption itself, rather than as an impact area to be managed after technological decisions have already been made.

“

“Precarity is a political project, and the current political economy is hostile to worker organizing. Particularly in the age of authoritarianism, there is only so far vulnerable workers can go before they meet the law, and this will continue to be an issue for worker organizing.”

James Farrar, Worker Info Exchange

Ensuring transparency and due process in algorithmic management


Algorithmic management is expanding rapidly across platform-mediated work such as ride-hailing, delivery, and logistics, and is growingly also embedded in public services including welfare, health, and education. Automated systems now routinely determine pricing, task allocation, incentives, penalties, and dismissal, often with little visibility into how these decisions are made. Workers are frequently unable to understand how their pay is calculated, why access to work fluctuates, or on what grounds accounts are suspended or terminated.

This opacity functions as a mechanism of control. When decision-making processes are hidden behind proprietary systems, workers face a structural asymmetry of information that undermines their ability to contest outcomes. Uncertainty itself becomes disciplinary, shaping behavior in the absence of clear or contestable rules.

Participants challenged the framing that algorithmic management is the necessary price of flexibility for workers. Flexibility was recognized as valuable, particularly for women workers balancing care responsibilities or informal livelihoods. However, flexibility is routinely mobilized to justify unpaid waiting time, income volatility, and the erosion of procedural safeguards. Regulatory interventions such as wage floors for ride-hailing workers in New York were cited as partial regulatory responses, but as participants noted, platforms often respond by increasing worker onboarding, thereby diluting earnings through oversupply.

Successful legal challenges to automated dismissal in some parts of the Global North underline that claims of algorithmic complexity cannot override basic principles of labor due process. Accepting automated decision-making as inherently inscrutable erodes long-established labor rights, including the right to know and negotiate how wages are set and the right to contest disciplinary action.

From this perspective, meaningful work in AI-mediated environments requires transparency in pay and task allocation, limits on automated decisions affecting income or termination, enforceable rights to explanation and appeal, and protections that delink flexibility from precarity. These measures are best understood as extensions of labor law into algorithmically managed workplaces, rather than exceptional safeguards.



Meaningful work in AI-mediated environments requires transparency in pay and task allocation, limits on automated decisions affecting income or termination, enforceable rights to explanation and appeal, and protections that delink flexibility from precarity.


Recognizing invisible labor in data and AI value chains

Meaningful work in the AI economy must extend beyond visible platform-mediated gig work to include the often invisible work that underpins AI systems. Data annotation, content moderation, and model training are foundational forms of labor that remain structurally precarious, low-paid, and weakly protected. This work is disproportionately concentrated in the Global South and frequently organized through short-term contracts and outsourcing arrangements that obscure employer responsibility.

Content moderation represents an especially harmful form of data work, involving prolonged exposure to violent, disturbing, or abusive material, often without adequate mental health support or recognition as hazardous work. Despite the psychological toll, such labor is rarely covered by occupational health and safety frameworks or social protection systems.

Labor invisibility is compounded by data extraction practices that treat workers' activity as free inputs. Labor-generated data, such as behavioral traces, performance metrics, and interaction logs, are routinely collected and repurposed to optimize platforms or train AI systems, often without workers' knowledge or consent. The principle of "no data without representation" was invoked to articulate this concern: workers whose labor generates data should have a say in how that data is used and how the resulting value is distributed.

This concern centers less on questions of individual ownership or compensation, and more on recognition, agency and participation at work. Meaningful work requires recognition of workers not merely as data sources, but as rights-bearing participants in data-driven systems.



A New Deal for the AI Economy must entail explicit recognition of data and AI-related work as labor. This means that labor law and social protection must be extended to data workers, along with enforceable standards for health, safety, and mental well-being and participatory governance over labor-generated data.

A New Deal for the AI Economy must therefore entail explicit recognition of data and AI-related work as labor, extension of labor law and social protection to data workers, enforceable standards for health, safety, and mental well-being, and participatory governance over labor-generated data.


Strengthening collective bargaining and worker voice in AI design and deployment

Digitally mediated workplaces are characterized by systematic erosion of workers' collective power. Surveillance, fragmented employment relationships, legal intimidation, and restrictions on protests function as structural barriers to organizing. In this context, individual grievance mechanisms and complaint-based remedies are insufficient for addressing harms produced by algorithmic systems that operate at scale.

Algorithmic management constitutes a collective issue rather than an individual one. Decisions about pricing, ratings, deactivation, or work allocation affect entire groups of workers simultaneously, making individualized remedies inadequate. Meaningful work therefore depends on rebuilding collective worker voice, not only in negotiating wages and conditions, but in shaping how AI systems are designed, deployed, and governed.


This includes extending collective bargaining into the technical domain – bargaining over algorithmic management practices, access to information about system functioning, and the right to contest changes to models or metrics that affect livelihoods.

Given the cross-border nature of platform economies and AI value chains, transnational coordination among worker organizations emerges as a necessary condition for effective bargaining.



Training and upskilling initiatives that focus solely on immediate employability fail to address structural power imbalances. Instead, a shift from skills to capabilities is required – building workers’ long-term ability to understand, question, and collectively shape digital systems.

Discussions further emphasized the limits of narrow skills-based responses to technological change. Training and upskilling initiatives that focus solely on immediate employability fail to address structural power imbalances. Instead, a shift from skills to capabilities is required – building workers’ long-term ability to understand, question, and collectively shape digital systems. Trade unions, worker organizations, and civil society are therefore critical sites for this capability-building, alongside education systems that prioritize critical digital literacy and democratic participation at work.





Diversified Economies

‘Diversified economies’ refers to a deliberate effort to reshape AI markets so that incentives, financing, and governance structures actively nurture locally grounded, highly contextualized, and collectively driven development and deployment of AI tools. Diversification foregrounds and attempts to expand the possibilities of who builds AI, for whom and under what economic logic—so that AI development supports domestic capacity building, sectoral upgrading and resilient regional ecosystems.

A diversified AI economy would be characterized by multiple centers of innovation operating at different scales, including public institutions, SMEs, cooperatives and community-embedded actors, rather than a narrow concentration of value creation around hyperscaler platforms. This orientation is critical for aligning AI innovation with local developmental priorities, employment generation, and public value creation across different economic contexts.



A diversified AI economy would be characterized by multiple centers of innovation operating at different scales, including public institutions, SMEs, cooperatives and community-embedded actors, rather than a narrow concentration of value creation around hyperscaler platforms.

Realizing such a transformation demands intentional market design. This requires fiscal, industrial and innovation policies that reward public AI value creation, enable knowledge and technology transfer, and lower entry barriers for local innovators. Diversification is thus a product of coordinated public investment, procurement, and regulation, which together set the terms for market competition.

Addressing structural deficits in business models through antitrust law

The current AI landscape is shaped by network effects, first-mover advantages, and capital-intensive infrastructures that favor concentration across multiple layers of the AI stack, limiting the space available for alternative models. Diversification therefore falls squarely within the remit of competition and market-shaping policy. Antitrust is not only a corrective to excessive concentration, but a proactive tool for opening space for alternative organizational forms, business models, and innovation pathways.

In many Global South contexts, antitrust enforcement remains largely ex-post, responding to market abuses after they have already skewed markets. This points to the need for anticipatory approaches that set out baseline conditions for market participation, such as mandatory data sharing, interoperability and transparency, before dominance consolidates. Examples from the EU, China, and Brazil illustrate how ex-ante frameworks can shape market structure upstream.

This perspective is closely linked to the view that business models matter as much, if not more, than ownership structures. Interventions limited to firm-break up or market-share thresholds are inadequate if extractive revenue models, such as targeted advertising or large-scale profiling, remain intact. Diversification thus requires reconfiguring incentives so that AI innovation oriented towards public services, local production and social infrastructure becomes economically viable and institutionally supported.

“

“In Brazil, there is a new policy thinking that focuses not just on competition law, but aims to make structural changes to counter market externalities like price management. Restricting transfer of data between subsidiaries or big data analyses for specific purposes in specific markets, can help tackle power asymmetries.”

Rafael Zanatta, Data Privacy Brasil

Institutional capacity and independence are essential for effective competition policy. Risks of regulatory capture, including revolving-door appointments and indirect influence through research funding and standard setting processes, undermine effective competition policy. Strengthening the autonomy of competition authorities, through statutory safeguards, transparent appointments and budgetary independence, is necessary to support long-term market shaping and to limit the capacity of powerful firms to formulate the very rules meant to constrain them.


Finally, the transnational nature of AI supply chains calls for a cross-border perspective on digital competition. Cross-border coordination is critical to prevent regulatory arbitrage and races to the bottom, including through alignment on competition principles, data governance norms and conditions attached to AI-related investment and trade.

Rethinking digital sovereignty through productive capacity

While the language of ‘digital sovereignty’ features prominently in policy discourses, sovereignty framed narrowly around local deployment or linguistic adaptation does not, on its own, translate into economic autonomy. User-facing services may appear localized while remaining dependent on externally controlled standards, APIs, compute infrastructure and research ecosystems.

A more substantive approach to sovereignty must be grounded in productive capacity rather than symbolic control. Many Global South economies face material constraints, including limited access to compute, fragile energy systems, and restricted fiscal space. In this context, strategies centered on creating national champions without addressing these shared structural constraints risks reproducing dependency in new forms.

Diversification requires policy scaffolding that supports a wide range of domestic actors, including SMEs, cooperatives, workers and public institutions. Measures discussed include preferential procurement of domestically produced AI solutions, targeted subsidies for green data-center construction, tax credits to support open-source AI contributions, and institutional arrangements that reduce exposure to predatory pricing or lock-in. Such interventions enable participation in AI value creation, rather than insulation of domestic actors from global markets.



Measures discussed include preferential procurement of domestically produced AI solutions, targeted subsidies for green data-center construction, tax credits to support open-source AI contributions, and institutional arrangements that reduce exposure to predatory pricing or lock-in.

At a macro level, fiscal and redistributive mechanisms, including global taxation measures on AI-related profits, form part of a long-term strategy to finance public interest AI research and shared infrastructures. These measures should operate in conjunction with domestic industrial and innovation policy, forming part of a broader policy orientation.

Expanding viable AI innovation pathways

Diversification depends on expanding the range of viable AI innovation trajectories, particularly those geared towards sectoral and public interest applications. Publicly funded and mission-driven research plays a critical role, especially where it addresses concrete local challenges in developmental domains such as agriculture, health, language access and public administration.

This points to the importance of realigning academic and research incentives away from narrow time-to-market metrics towards longer-terms problem solving rooted in local contexts. Greater coordination between universities, public research institutes, and community organizations can enable AI systems that are socially relevant, technically appropriate and institutionally embedded.

Clear policy conditions also shape innovation pathways. Expectations around technology transfer, transparency and sector-specific limits on deployment ('redlines'), particularly in safety-critical or high-risk domains, can function as steering mechanisms for socially productive innovation rather than as constraints on innovation per se.


“

“A small number of companies have established dominant positions in the AI economy by growing their network power and their market power, bolstered by high entry barriers for late comers.”

Laura Mann, London School of Economics

Finally, diversification requires challenging narratives of inevitability surrounding large-scale, general-purpose AI models. The current configuration is the result of historical choices: early public-funded research in the United States, aggressive IP regimes, and the rise of ‘academic capitalism.’ Diversified economies thus emerge as contingent outcomes of policy choices, institutional design and collective action.





Pluralistic Knowledge Societies

‘Pluralistic knowledge societies’ refers to an AI ecosystem that recognizes and actively sustains multiple ways of producing and governing knowledge. Built with deep learning techniques, today’s AI solutions depend on massive, granular, and highly contextualized datasets, with LLMs requiring billions of labelled examples that capture subtle linguistic, behavioral, and environmental cues. However, the digital divide shaped by historical geopolitical fault lines means that the data most readily available for AI training is overwhelmingly rooted in Western economic interests, legal frameworks, and cultural norms.

Indigenous, local, and traditional knowledge systems—often transmitted orally, embedded in social practices, stored in community rituals, or encoded in place-based relationships—remain underrepresented or entirely absent from the pipelines that power contemporary AI models. When AI technologies are built on such a skewed substrate, they reproduce and amplify the cultural assumptions embedded in their source material. The consequence is more than misrecognition; it creates fertile ground for epistemic erasure, where entire ways of knowing are rendered invisible, devalued, or overwritten.

Across discussions, participants emphasized that AI systems are not neutral repositories of information, but infrastructures that actively shape whose knowledge becomes legible, authoritative, and valuable. From this perspective, pluralism must extend beyond representation and diversity to function as a necessary condition for democratic innovation.

“

"In Spanish, we use the term ‘knowledge’ in plural, whereas it is seen as singular in English, indicating that a perfect cultural translation between cultures is impossible. So merely changing the business model of AI to be less extractive may not resolve the issue of monocultural impositions, as Western cultures continue to find newer ways of dispossession, control, and violence."

Paola Ricaurte Quijano, Feminist AI Research Network, Latin America and the Caribbean


Building a communitarian understanding of knowledge production and ownership

Contemporary knowledge production is framed around the figure of the solitary researcher, publishing in peer-reviewed journals, securing grants, and claiming IP through patents and licenses. This individualist conception aligns neatly with dominant IP regimes, where ownership, licensing and revenue streams are allocated to identifiable persons or corporate entities. However, many Indigenous, local and community-based knowledge systems operate on fundamentally different principles.

Knowledge in these contexts is often produced collectively, refined over generations through shared practice, oral transmission and communal stewardship, with value residing in its collective custodianship rather than exclusive ownership.

This tension becomes especially visible when conventional IP frameworks attempt to parcel out rights over cultural artefacts, traditional songs, agricultural practices, or medicinal formulae. By assigning exclusive rights, the law often strips communities of the ability to control how their heritage is used, commercialized, or transformed. These frictions amplify in the AI context, where LLMs and other large-scale systems ingest vast volumes of text and media—including folklore, oral histories and community-generated content—without meaningful consent, attribution or recognition of provenance. Outputs derived from these culturally specific insights are then presented as generic ‘responses’, erasing social relations, context and responsibilities through which that knowledge was produced.

Pluralistic knowledge societies therefore require a shift toward communitarian approaches to knowledge production and governance. Instead of extractive pipelines that treat community knowledge as raw material, this would entail community-led curation, co-authorship of datasets and participatory governance arrangements that extend over the lifecycle of AI systems and AI-derived artefacts. By establishing bottom-up feedback loops that allow communities to exert authority over how their knowledge is represented and circulated, plural epistemic orders can be protected against subsumption into dominant epistemologies.



Instead of extractive pipelines that treat community knowledge as raw material, this would entail community-led curation, co-authorship of datasets and participatory governance arrangements that extend over the lifecycle of AI systems and AI-derived artefacts.

Reorienting research institutions and knowledge infrastructures

Between community-held knowledge and deployed AI systems sit powerful mediating institutions, including universities, research labs, standards bodies, funding agencies, and digital repositories. These institutions function as sites where epistemic hierarchies are reproduced—or could be transformed. Current research ecosystems, shaped by publication incentives, competitive funding models, and proprietary outputs, systematically privilege knowledge that is easily codified, standardized, and scaled, while undervaluing slower, relational, or situated knowledge practices. Even efforts to include Indigenous or local knowledge often remain extractive when institutional incentives remain unchanged. Community knowledge is frequently treated as data input rather than as an epistemic system with its own authority, governance norms, and ethical conditions. International research collaborations framed around “solving global challenges” can reproduce this asymmetry, by extracting data and contextual expertise from the South, while intellectual property, reputational capital and technical capacity accrue to institutions in the North.

Therefore, reorienting research institutions and infrastructures toward co-creation is central to fostering pluralistic knowledge societies. This includes funding models that support long-term partnerships with communities; evaluation criteria that value social relevance, reciprocity, and care; and institutional arrangements that recognize communities as knowledge authorities rather than stakeholders or beneficiaries. Public repositories, data trusts, and research infrastructures can serve as sites for embedding collective governance, provided they are designed with explicit accountability to knowledge-producing communities.

Enabling communities to co-create and co-govern AI tools

Pluralistic knowledge societies reject the assumption of universal AI solutions, on the understanding that the very definition of “need” varies across contexts. Some communities benefit from low-tech, offline-first tools that support records-keeping or resource monitoring, while others require sophisticated systems, such as natural-language interfaces to enable access to legal documents or administrative information in local dialects. In all cases, affected communities must be involved at the outset, before technical design decisions are made.

Early and sustained participation can enable collective deliberation over the purpose, scope, and ethical implications of AI interventions—including whether AI is appropriate at all. In domains such as welfare administration, where eligibility criteria are already codified in law, algorithmic systems can introduce unnecessary complexity, opacity, and risk. In contrast, in domains involving context-dependent judgement—such as agricultural decision-making in micro-climates—AI systems can at times complement human expertise, provided their limitations are openly discussed and governance arrangements remain transparent.

Although participatory approaches are time- and resource-intensive, they yield tangible benefits. For example, in certain cases, locally sourced vocabularies have resulted in improvement of a system’s performance on dialect recognition, and community-validated metadata has helped reduce land-use misclassification. Involving communities in problem definition and constraint-setting can also help avoid misdiagnosis and prevent the allocation of AI resources into complex systems where simpler interventions are more appropriate.

Sustaining such approaches requires deliberate institutional support. This includes targeted capacity-building to enable communities to engage critically with AI systems; strengthening community-based knowledge rights to formalize collective stewardship arrangements; and addressing extractive commercial practices, such as planned obsolescence and tax evasion, which undermine pathways for long-term community control and epistemic autonomy.



Planetary Flourishing

Prevailing narratives of ‘AI for sustainability’ often function as legitimizing narratives for continued expansion and extraction across AI value chains. While AI is increasingly promoted as a tool for climate mitigation, biodiversity monitoring, or energy optimization, this framing frequently obscures the growing ecological costs of AI systems themselves. Rather than treating sustainability as a “use-case” for AI, ecological constraints must function as a governing condition for when, where, and whether AI systems should be developed.

Governing AI within planetary flourishing requires unpacking the materiality of AI and rejecting the assumption that digital systems are abstract or immaterial. AI systems are embedded in resource-intensive infrastructures dependent on energy, water, land, minerals, and global hardware supply chains. From this perspective, planetary flourishing must operate as material limits that set the outer conditions of AI innovation.

“

“The question here is not whether AI has potential. It is whether its current trajectory aligns with a habitable planet. The logic of “bigger models, more data, more compute” clashes directly with the physics of a finite planet. Without intentional limits, the growth curve of AI becomes a growth curve of extraction and emissions.”

Maria Soledad Vogliano, ETC Group

Foregrounding restraint and intentional limits in AI innovation

A defining feature of the current AI paradigm is the dominant growth logic, embedding the assumption that “bigger models, more data, and more compute” constitute progress. This logic systematically sidelines questions of necessity, proportionality, and ecological costs, especially in public and developmental contexts where AI is increasingly deployed as a default solution.

Governance bounded by planetary limits requires governance frameworks that enable restraint. This includes systematic assessment of whether AI is necessary for a given task, and explicit consideration of low-use or non-use pathways where ecological costs outweigh potential social benefits. Such assessments reframe innovation as a question of purpose, not inevitability.

Purpose-based assessment must include deliberation with environmental groups, small producers, workers and other stakeholders across the AI value chain, over whether AI systems are appropriate at all.

In domains such as welfare administration, healthcare triage, or public surveillance, AI deployment can introduce significant ecological and social costs. In contrast, the use of smaller-scale, task-specific, offline-first, or non-Machine Learning (ML) alternatives may deliver comparable outcomes with far lower resource intensity – an especially salient consideration in Global South contexts marked by acute energy and water constraints.

Boundary-enforcing mechanisms—such as express moratoria, bans, and non-use provisions—are therefore integral for aligning AI development with ecological limits. Such mechanisms establish ecological considerations as ethical baselines for innovation, rather than exceptional interventions reserved for worst-case scenarios.

Embedding environmental accountability across AI value chains

Ecological accountability cannot be limited to narrow metrics such as operational carbon emissions. Meaningful oversight necessitates lifecycle approaches that account for energy use, water consumption, land competition, mineral extraction, toxic chemicals, and e-waste across the AI value chain.

Mandatory disclosure and assessment regimes are particularly critical for hyperscale data centers, semiconductor fabrication, and critical mineral extraction. Environmental impact assessments must be cumulative, place-based, and participatory, especially in regions where AI infrastructure is promoted as a development strategy. Procedural or fragmented approvals are woefully inadequate for capturing long-term and cross-sectoral ecological impacts.

Regulatory tools already available to states, including public procurement conditions, antitrust interventions, and community-based monitoring mechanisms, can be leveraged to embed environmental standards upstream. Voluntary ESG commitments and efficiency standards, while often foregrounded in regulatory discourses, remain insufficient in the absence of binding obligations, transparency and enforcement.

Advancing distributive justice in energy and AI transitions

The ecological costs of AI expansion are unevenly distributed. Communities in the Global South bear disproportionate burdens associated with energy demand, water stress, mineral extraction, and land use, while the economic benefits of AI innovation remain highly concentrated.

In this context, a key question is who gets to decide what levels of computation, infrastructure expansion, and resource use are ecologically and socially acceptable, and on whose behalf these decisions are made.

Efficiency gains generated by AI cannot be allowed to accrue privately while ecological costs are socialized, particularly through weakened environmental review processes justified by narratives of national growth or global competitiveness.

Just AI transitions require fiscal and redistributive instruments designed to reallocate value. These include taxation of AI-related profits, financing mechanisms to support renewable energy infrastructure and public digital infrastructures, ecological restoration, and compensation for affected communities. ESG frameworks may play a complementary role, but only where translated into binding standards, drawing on precedents from environmental, safety and appliance regulations that have successfully translated normative commitments into enforceable market rules.



Regenerative AI: A New Deal for the AI Economy

Meaningful and Dignified Work

Countering precarity, invisibility, and algorithmic control of workers

- Re-centering labor as an organizing principle for AI adoption, not a downstream corrective.
- Recognizing data and AI labor as a foundational building block of digitalizing economies.
- Extending labor rights and due process to all work managed algorithmically in local-to-global value chains.
- Nurturing worker-led platform models to democratize data value.

Diversified Economies

Countering concentration of data value and predatory market structures

- Shifting AI markets away from hyperscaler dominance toward multiple centers of innovation.
- Grounding AI development in domestic capabilities, sectoral specificities, and public value creation.
- Evolving industrial policy, public procurement frameworks, and competition law to strengthen local value generation.

Pluralistic Knowledge Societies

Countering knowledge extraction and epistemic exclusion

- Making knowledge pluralism a precondition for democratic innovation, not a representational fix.
- Acknowledging communities as epistemic authorities, not merely data sources.
- Reorienting knowledge ecosystems, intellectual property regimes, and AI design toward collective co-creation, stewardship, and governance.


Planetary Flourishing

Countering ecological harm and growth-first logics

- Treating the environment as a baseline consideration for AI innovation, not externality.
- Adopting redlines, restraint, and proportionality in decisions about AI development and deployment.
- Embedding lifecycle accountability, distributive justice, and just transition into AI governance.



4. Levers of Change for Just AI Futures



This section synthesizes the strategies identified during the dialogue into a set of levers for advancing a Regen AI paradigm. The dialogue reaffirmed that the transition toward such a paradigm requires coordinated interventions across discourse, institutions, markets, democratic processes and technological design. The levers outlined in the next section distil the strategies identified by participants for shifting power relations, incentives and decision-making authority across the AI lifecycle. While presented in distinct buckets for analytical clarity, these levers are mutually reinforcing and must be pursued in combination by a range of societal institutions and actors to facilitate durable change.

Techno-design choices

AI futures are shaped by design choices, addressing questions of scale, architecture, ownership, and purpose, which are often treated as technical inevitabilities rather than political decisions. Key directions include:

- Institutionalizing restraint-led innovation, including explicit consideration of non-use and low-use cases where AI offers limited social benefit relative to its costs.
- Promoting frugal and appropriate technologies, such as small and modular language models, offline-first and low-resource systems, and sector-specific tools, particularly in public services.
- Embedding co-design and participatory governance into AI development, ensuring that communities and workers co-determine problem definitions, limits, and success metrics.
- Supporting commons-based approaches, including open-source, cooperative and community-owned AI systems, backed by social licences, public funding and procurement.
- Recognizing the political economy of hardware and infrastructure, including how choices around compute and capital allocation shape downstream innovation pathways.
- Encouraging upstream public research models, including collaborative and non-proprietary approaches to foundational AI research, inspired by CERN-style collaboration for AI.

Discursive shifts and evidence collection

Prevailing narratives surrounding AI, centered on inevitability, efficiency and scale, constrain political imagination and weaken public accountability. A critical lever of change lies in reshaping how AI is discussed and legitimized in public and policy spaces. Key directions include:

- Challenging narratives of inevitability and neutrality by foregrounding AI as the outcome of political, economic and institutional choices rather than an autonomous technological force.
- Unpacking the materiality of AI by conceptualizing it as a socio-technical assemblage grounded in infrastructures, labor, data, and political-economic relations.
- Making harms visible and legible, including labor exploitation, ecological costs, epistemic exclusion and enclosure, particularly where these are obscured through technical complexity or business secrecy.
- Strengthening evidence through situated research that documents how existing AI systems operate in specific contexts, instead of relying on abstract projections.
- Building public repositories and shared knowledge bases on AI harms, failures and contested deployments, drawing on regulatory disclosures, worker testimony, community audits, and litigation filings.
- Expanding critical interdisciplinary education and awareness at all levels, including ethics and accountability training for engineers, public officials and enterprises.
- Grounding ‘localization’ and ‘sovereignty’ narratives in the histories, struggles and political imaginaries of popular social movements and collective mobilizations.
- Instituting protections for whistleblowers within large technology firms.

Law and institutions

Law and regulation are among the most evident tools for shaping AI markets and innovation trajectories, provided they are aligned with public purpose and supported by institutional capacity. Key directions include:

- Using regulation as an ex-ante market-shaping tool, to align with developmental, industrial and social policy objectives, instead of simply responding to harm.
- Resisting regulatory imitation and developing context-specific regulatory frameworks that reflect locally salient norms, power structures and social hierarchies.
- Instituting red lines and harm thresholds, including explicit prohibitions on AI applications that pose unacceptable social, political or environmental risks.

- Reversing the burden of proof, so that firms must demonstrate public benefit and safety before deployment.
- Reforming intellectual property regimes to institutionalize collective ownership structures.
- Strengthening accountability across AI value chains, including mechanisms for collective redress, liability beyond point-of-use harms, and transparency obligations.
- Expanding access to information, including extending freedom-of-information principles to AI systems deployed in public services, procurement or infrastructure.
- Embedding social, environmental and algorithmic audits into governance frameworks, with community participation and intersectional representation, including independent ombudspersons.
- Engaging international and multilateral institutions, such as the ILO and the UN' Working Group on Business and Human Rights, to recognize corporate liability and support transnational accountability for AI firms.
- Building South-South coordination in international trade, data governance, and AI diplomacy, while actively avoiding dominance by large emerging economies.

Incentives and market shaping

AI trajectories are fundamentally shaped by incentive structures embedded in fiscal policy, procurement decisions, and investment priorities. Changing market outcomes therefore requires deliberate redesign of these mechanisms. Key directions include:


- Leveraging public procurement to prioritize public value, domestic integration, worker protections, transparency and sustainability.
- Moving beyond price/efficiency-based evaluation to include labor standards, explainability, interoperability, environmental impact, and local capacity building.
- Embedding safety certification, mandatory disclosures and technology assessment into procurement and deployment cycles, enabling a more precautionary approach.
- Taxing AI value chains and large technology firms, with revenues directed towards public goods, shared infrastructure, ecological repair and domestic research and development.
- Investing in shared and common-pool infrastructures, including compute, data commons, and digital public infrastructures, tied to public-benefit access conditionalities.
- Shifting from FDI-led to demand-led innovation models, anchoring AI adoptions in domestic productive capacity and public needs.
- Mandating technology transfer and capacity building as conditions for market access, especially where firms rely on public data, land or subsidies.

- Financing digital public goods and data trust models that support community ownership, cooperative platforms and sector-led innovation.

Democratic participation and worker power

Democratic participation was repeatedly identified as both a goal and a lever of change in AI governance. Participants stressed that meaningful participation must extend beyond consultation to include collective power and enforceable rights. Key directions include:

- Establishing universal labor guarantees, regardless of employment status, including pay for all working time, social protection, and enforceable health and safety protections.
- Curbing algorithmic management, including limits on fully automated decisions relating to pay, discipline and dismissal, and enforceable rights to explanation and appeal.
- Strengthening collective bargaining, extending it to cover data use, algorithmic management and technology-driven change.
- Creating institutional mechanisms for participation, such as social audits, community consultations, and participatory oversight in AI deployment, particularly in public services.
- Recognizing labor-generated data as a collective resource, and formulating clear demands around its use in AI development and deployment.
- Supporting worker-owned technologies, including counter-surveillance tools to deter union-busting by technology firms.
- Developing worker-centric just transition frameworks for sectors such as automation, including autonomous vehicles.
- Supporting transnational organizing and solidarity, including actions at strategic choke points in AI value chains.
- Building alliances across workers, civil society, researchers and affected communities to counter fragmentation and regulatory capture.



Techno-design choices

Addressing design choices around scale, architecture, ownership, and purpose, as political decisions rather than technological inevitabilities to shaping AI futures.




Law and institutions

Aligning law and regulation for shaping AI markets and innovation trajectories with public purpose and supported by institutional capacity.



Discursive shifts and evidence collection

Reshaping how AI is discussed and legitimized in public and policy spaces.



Democratic participation and worker power

Extending meaningful participation beyond consultation to include collective power and enforceable rights.



Incentives and market shaping

Redesigning incentive structures embedded in fiscal policy, procurement decisions, and investment priorities that AI trajectories to change market outcomes.

Levers of change for just AI futures

