

THE HUMAN EDGE

Global Perspectives on
AI and Prosperity



© 2026 Global Federation of Competitiveness Councils

This publication may not be reproduced, in whole or in part, in any form beyond copying permitted by sections 107 and 108 of the U.S. copyright law and excerpts by reviewers for the public press, without written consent from the publishers.

The GFCC is a nonprofit, 501(c)(3) organization as recognized by the U.S. Internal Revenue Service. The GFCC's activities are funded by contributions from its member organizations.

To learn more about the GFCC visit our home page at thegfcc.org

Table of Contents

A Letter from the GFCC Leadership	4
Learning, Adapting, and Building the Answers	6
Executive Summary by Prof. Aleks Subic, Chairman of the Global Task Force on Human and AI	8
Chapter 1 - The Moment We're In: The State of AI in 2026	15
Chapter 2 - How the World Sees the Human-AI Co-evolution: What Leaders Across 32 Countries Told Us	23
Chapter 3 - Voices of the Transition: Perspectives from the Global Community	32
Chapter 4 - The Human Edge: What Humans Bring to the Age of AI	41
Chapter 5 - The Inclusion Gap: AI's Unequal Footprint	54
Chapter 6 - What Good Looks Like: A Gallery of Human-AI Integration That Is Working	64
Chapter 7 - The Work Begins Now: Turning Insight into Action	84
Annex I - Survey Methodology Overview	100
Annex II - GFCC Members	102
Annex III - GFCC Fellows	103

A Letter from the GFCC Leadership

We are at a tipping point. Artificial intelligence (AI) has already changed our relationships with work, life, and learning. Beyond that, it is reshaping the foundations of how economies grow, how knowledge is created, and how opportunity is distributed. The question before us is not whether AI will impact our lives and be part of our future. It is whether we will shape it wisely and benefit from it, or be shaped by it passively.

The Global Federation of Competitiveness Councils (GFCC), in partnership with Torrens University Australia, is pleased to present *The Human Edge: Global Perspectives on AI and Prosperity*, a comprehensive report capturing the views of leaders from 32 countries, representing a wide range of sectors, industries, and institutions, at the forefront of the human–AI transition.

The insights outlined here draw from a pioneering global survey exploring the intersection of human and AI co-evolution and provide a unique global outlook highlighting how organizations and societies are adapting in the face of the institutional and capability challenges posed by AI. More importantly, they illuminate what it will take to ensure that technological advancement translates into inclusive and shared prosperity across societies and economies.

The findings are both encouraging and urgent. At the center of many of the opinions shared is the need to enable talent development through an institutional overhaul across businesses, educational systems, and government programs—an agenda that has long been one of GFCC's defining priorities. For years, we have championed the idea of talent as the foremost driver of future competitiveness.

In 2021, we launched *Future Skills: Lessons and Insights from a Review of Innovative Skills Development Initiatives*, a groundbreaking report mapping the skills, strategies, and initiatives needed to power talent development and boost competitiveness in a changing global landscape.

It is not by chance that some of the projects highlighted in the Future Skills report are also featured here as real-world examples of successful human-AI integration. For years, we have been studying these examples, sharing best practices, and engaging our global community in discussions about the drivers and enablers necessary to make future economies work for everyone.

As the new findings demonstrate, the opportunity ahead is immense, but leading the path forward requires intentionality and deliberate action. We need to invest in human capability. We must reimagine the institutions and governance frameworks for the world we are entering rather than the one we are leaving behind. The window to act is open now, not indefinitely.

The GFCC is committed to convening, connecting, and championing the efforts required to ensure that human capability leads the AI era. We are deeply grateful to Professor Aleks Subic—Vice-Chancellor at Torrens University Australia, co-author of this report, and Chair of the GFCC task force that steered its development—for the exceptional intellectual leadership that made this report possible.

We hope you find insightful lessons here and join this global conversation.



Deborah L. Wince-Smith
Founder and President,
GFCC



Charles O. Holliday Jr.
Chairman, GFCC

Learning, Adapting, and Building the Answers

We can assure you that the ideas in this report are ours, although we would be remiss not to acknowledge that the text itself has benefited from AI-powered copy-editing. In that sense, this report reflects the reality that artificial intelligence is already shaping the work and lives of all of us, including those who contributed to it.

The rise of AI and particularly agentic AI is at the heart of this discussion. It is not a new topic for the GFCC community. We began exploring it in 2023 during our Annual Meeting in London, when Gallup hosted us at The Shard. At the time, some of the ideas being discussed seemed far-fetched. They were not. Today, agentic AI is part of our reality, and the scale of the transformation is impossible to ignore.

We are all hearing about it, reading about it, thinking about it, and increasingly experiencing it firsthand. Yet important questions remain. How are leaders across the world perceiving these changes? What opportunities do they see? What risks concern them most? And above all, how can we ensure that humans remain at the center of progress in a world where technological capabilities are advancing at unprecedented speed?

These questions are particularly important at a time when many of the gaps that challenge societies are not narrowing, but widening. Economic divides, educational disparities, unequal access to technology, and differences in institutional capacity continue to shape how communities experience change. What can we do to ensure that AI becomes a force for opportunity rather than a source of deeper fragmentation?

The GFCC is uniquely positioned to explore these questions. Our community brings together leaders from dozens of countries and diverse backgrounds across business, government, academia, and civil society. When we undertook this survey and report under the leadership of Professor Aleks Subic, we expected to gather valuable perspectives and generate thoughtful discussion. What we received surpassed our expectations — in the depth of perspective, the candor of the responses, and the surprising areas of consensus that emerged across very different contexts. But the contributions that informed this report were not limited to the survey itself.

Moreover, the content presented here also benefits from discussions of our preliminary findings with survey respondents and the GFCC community. We hosted two online sessions that gave participants the opportunity to react to, challenge and expand upon the findings and opinions surfaced through the survey. These discussions allowed us to deepen the analysis.

The insights presented in this report reveal not only how leaders around the world are thinking about AI, but also where their views converge and where important differences remain. In a world where the pace of change is outrunning the capacity of institutions and policies to keep up, those differences are neither surprising nor unwelcome. On the contrary, they help illuminate the challenges that demand our attention and provide a clearer understanding of where action is needed.

On behalf of the entire GFCC project team, we would like to express our sincere gratitude to Professor Aleks Subic for his leadership, intellectual guidance, and partnership throughout this project. His true commitment was instrumental in transforming this initiative into a truly global conversation.

The publication of this report is not the end of the journey. It is the beginning of the next phase. Across the GFCC community, we are already working to translate these insights into action, developing initiatives, partnerships, and practical recommendations that can help societies, institutions, and individuals navigate the opportunities and challenges of the AI era.

The future of AI has not yet been written. Its impact will depend on the choices we make, the values we uphold, and our willingness to ensure that technological progress advances human flourishing. Like everyone else, we are learning, adapting, and building the answers as we go. We invite you to join us on that journey.



Dr. Roberto Alvarez
Executive Director, GFCC



Dr. Elaine Rodriguez
Director of Operations,
GFCC



Executive Summary

We are living through a decisive moment in the evolution of human capability, institutional competitiveness, and inclusive growth. **Artificial intelligence is no longer an emerging technology at the edge of the economy.** It has entered classrooms, workplaces, boardrooms, laboratories, public services, and homes with a speed and scale that few institutions were designed to absorb.



Prof. Aleks Subic

Vice-Chancellor Torrens University
Australia and Chairman of the Global
Task Force on Human and AI

The central message from the GFCC Human and AI work is clear: the question is no longer whether AI will transform economies and societies. It already is. The real question is whether we will shape that transformation deliberately, humanely, and inclusively – or allow it to widen the gaps that already exist between people, institutions, sectors, and nations.

The reports and survey findings point to a powerful but uncomfortable conclusion. AI capability is accelerating faster than institutional readiness. The technology is advancing rapidly, investment is surging, and adoption is spreading across sectors. Yet education systems, workforce models, governance frameworks, and social protections are not keeping pace. This is the defining competitiveness challenge of the AI transition.

The GFCC survey of 113 leaders across 32 countries and five sectors captures this moment well. Leaders are not broadly fatalistic. Most see AI as transitional rather than purely disruptive, and there is cautious optimism that human-AI co-evolution can be positive. But that optimism is conditional. It depends on whether governments, businesses, universities, and civil society are willing to redesign the systems around AI, rather than merely inserting AI into existing models of work, education, and governance.

This conclusion is strongly supported by contemporary research. Stanford HAI's 2026 AI Index reports that global corporate AI investment more than doubled in 2025, reaching US\$581.7 billion, while McKinsey's 2025 global survey found that 88% of organizations now use AI in at least one business function, even though most remain at experimentation, piloting, or early scaling stages.

The World Economic Forum projects that by 2030, 170 million jobs may be created and 92 million displaced, producing a net increase of 78 million jobs but also major labor-market churn. The IMF estimates that almost 40% of global employment is exposed to AI, rising to around 60% in advanced economies. PwC's 2025 Global AI Jobs Barometer shows that workers with AI skills already command a 56% wage premium, signaling that the AI divide is not a future risk; it is already emerging in labor markets.

The opportunity is immense. AI can lift productivity, accelerate scientific discovery, improve decision-making, personalize education, expand healthcare access, support small businesses, and democratize expertise in ways that were previously unimaginable. But the same technology can also concentrate power, displace routine work, administrative and "white-collar" roles, disrupt traditional apprenticeship or graduate pathways, deepen digital divides, amplify bias, and reward those already closest to capital, compute, data, and advanced skills.

That is why the core issue is not technological adoption alone. It is human and institutional adaptation.

Key Insights



1

The human edge is real, but it must be deliberately developed

One of the strongest findings from the GFCC survey is that leaders do not see technical AI expertise alone as the decisive differentiator. The capabilities most valued in the AI economy are creativity, critical thinking, ethics, adaptability, lifelong learning, digital and AI literacy, complex problem-solving, collaboration, and entrepreneurial mindset. Technical AI and engineering expertise matters, but it ranked well below these broader human capabilities in the survey.

This is a profoundly important insight. It challenges the narrow assumption that preparing people for AI simply means producing more coders, data scientists, or AI engineers. Those skills are essential, but they are not sufficient. The greatest value will come from people who can frame problems, exercise judgement, ask better questions, work across disciplines, understand context, build trust, and use AI to create new forms of value.

In other words, AI raises the premium on distinctly human capability. But the report is right to insist that the human edge is not automatic. Creativity, judgement, and adaptability are not equally cultivated across society. They depend on access to quality education, meaningful work, digital infrastructure, psychological safety, mentoring, time to think, and opportunities to learn by doing. Those conditions are unevenly distributed.

This has major implications for education and workforce policy. We need to move beyond static content, fixed degrees, and episodic training interventions. Education in the AI-era must become more modular, adaptive, applied, and lifelong. It must integrate AI literacy with critical thinking, ethical reasoning, systems thinking, entrepreneurship, and professional readiness. It must also protect and reimagine the apprenticeship and graduate pathways through which people build experience and ability, particularly as AI begins to perform more entry-level cognitive work.

2

The central risk is not that AI replaces humans, but that it rewards some humans and excludes others

The survey shows that leaders see AI as an amplifier. It can amplify productivity, creativity, and access. But it can also amplify existing inequality. The groups most exposed are those already facing structural disadvantage: the Global South, low-income communities, workers in routine-based roles susceptible to automation, older workers, women in administrative and clerical roles, young graduates entering professions where AI may displace entry-level work, linguistic minorities, rural communities, and people without reliable digital access.

This matters because inclusive AI transformation will not happen by accident. Market forces alone are unlikely to distribute benefits evenly. Left to itself, AI will tend to concentrate value around those with data, compute, capital, advanced skills, and institutional capacity.

The implication is that inclusion must be designed into AI transformation from the outset. That means access to digital infrastructure, affordable tools, AI literacy,

local-language resources, trusted public platforms, community co-design, reskilling support, and governance models that protect workers and citizens without freezing innovation.

The cases identified in the report, including AI education in Piauí, robotics education in Romania, Spain's computational thinking and AI program, Singapore's SkillsFuture infrastructure, Kenya's AI-enabled financial ecosystem, and Alberta's community-designed health applications, show that inclusive AI is possible when institutions begin with the question: who is most likely to be left out, and what design choices would bring them in?

3

Institutional readiness is the real competitiveness frontier

The report rightly frames the current moment as an institutional gap. AI capability is progressing rapidly, but the systems needed to align, govern, absorb, and distribute its value are lagging. This gap is visible in education, regulation, labor markets, organizational design, and public trust.

The GFCC survey found overwhelming concern about education and workforce preparedness. Skills mismatch and education lag were the top barriers to human-AI complementarity. The same pattern is visible in public-domain research. The World Economic Forum identifies skills gaps as a major barrier to business transformation and expects 39% of workers' core skills to change by 2030. OECD research also warns that much AI-related training is still targeted disproportionately toward higher-skilled adults and AI professionals, rather than the wider workforce that needs practical AI literacy and transition support.

This is where competitiveness strategy must evolve. In the past, national, and regional competitiveness was often measured through productivity, innovation intensity, investment attraction, research output, and export performance. Those remain important. But in the AI economy, competitiveness will also depend on adaptive capacity: the ability of people, institutions, and regions to learn quickly, redesign work, build trust, deploy technology responsibly, and distribute productivity gains fairly.

A nation, region or university that deploys AI without strengthening human capability may produce short-term efficiencies but long-term fragility. By contrast, those that align AI adoption with skills, inclusion, governance, and innovation ecosystems will build more durable advantage.

Strategic recommendations



1. Establish AI-ready education and skills frameworks as a global priority.

The clearest priority emerging from the survey is the need for AI-ready education and skills frameworks. These should not be narrowly technical. They should define the full capability architecture required for human-AI complementarity: AI literacy, critical thinking, creativity, systems thinking, ethics, adaptability, collaboration, entrepreneurial mindset, and applied problem-solving. These frameworks should be adaptable across the entire continuum of education, including K-12, vocational education, universities, professional education, and workforce reskilling and upskilling.



2. Build lifelong learning as infrastructure, not as a series of disconnected programs.

The AI transition will not be managed through short-term training schemes alone. Countries and institutions need permanent lifelong learning infrastructure: portable learning accounts, modular credentials, employer co-investment, stackable qualifications, recognition of prior learning, and pathways that allow people to move from AI-exposed roles into AI-augmented roles. Singapore's SkillsFuture model is instructive not because it can be copied exactly, but because it shows the importance of building the infrastructure before crisis hits.



3. Redesign work, not just curricula.

The reports make an important point: the value of AI is realized when workflows are redesigned, not simply when tools are adopted. Organizations need to move from task automation to work transformation. That means identifying where AI can remove low-value cognitive burden, where human judgement must remain central, where new roles are needed, and how staff can be involved in shaping adoption. This is essential for productivity, trust, and wellbeing.



4. Protect the apprenticeship and graduate pathways into expertise.

One of the most under-appreciated risks is the potential erosion of

entry-level learning. If AI performs the routine work through which graduates and early-career professionals learn judgement, then the pathway to expertise and experience may be weakened. This is particularly relevant in law, finance, consulting, accounting, engineering, healthcare administration, software development, and research. Universities and employers should jointly design new AI-era apprenticeships and graduate route models in which early-career workers learn with AI, not around it or in competition with it.



5. Make inclusion a design principle, not a downstream mitigation strategy.

AI policy should begin with the most exposed groups and regions. This includes the Global South, rural communities, SMEs, low-wage workers, women in routine roles, older workers, and communities with limited digital infrastructure. Inclusive AI requires deliberate investment in connectivity, local-language models, community-based training, public digital infrastructure, accessible tools, and trusted intermediaries.



6. Develop a human-centered AI competitiveness scorecard.

The reports make an important point: the value of AI is realized when workflows are redesigned, not simply when tools are adopted. Organizations need to move from task automation to work transformation. That means identifying where AI can remove low-value cognitive burden, where human judgement must remain central, where new roles are needed, and how staff can be involved in shaping adoption. This is essential for productivity, trust, and wellbeing.



7. Move from voluntary principles to accountable governance.

Ethical AI frameworks are necessary, but insufficient if they remain voluntary and disconnected from practice. Governments, businesses, and universities should develop governance models with clear accountability, audit requirements, human-in-the-loop standards, worker consultation, transparency expectations, and mechanisms for redress. The aim should not be to slow innovation, but to build trust in innovation through effective and responsible adaptation.

Concluding reflection



The greatest danger is not that AI becomes too powerful. The greater danger is that human systems remain too passive, too fragmented, and too slow to adapt.

The reports show that leaders across the GFCC network understand both the opportunity and the risk. They are not rejecting AI. They are calling for a more deliberate model of human–AI co-evolution, one that strengthens human capability, redesigns institutions, protects inclusion, and measures success by lived outcomes as well as aggregate growth.

This is the moment for GFCC to move from diagnosis to leadership. The work ahead should not be framed as a technology agenda, but as a competitiveness and human prosperity agenda. AI will reshape productivity, work, education, and innovation. Whether it also expands opportunity, dignity, and inclusive prosperity will depend on the choices we make now.

The next phase should therefore be practical, collaborative, and measurable: build the alliance, establish the hubs, define the scorecard, mobilize universities, engage the Global South, and bring the voices of workers and communities into the center of the agenda.

The report should not close the conversation. It should launch a movement.

Prof. Aleks Subic

Chair, Human and AI Taskforce

Chair GFCC University Research and Leadership Forum

Vice Chancellor and President, Torrens University Australia



The Moment We're In: The State of AI in 2026



CHAPTER 1

AI capability is exponentially growing. In 2025, performance on SWE-bench Verified — a benchmark measuring whether AI can resolve real software engineering problems — jumped from 60% to near 100% of the human baseline in a single year.¹ Frontier models now meet or exceed human performance on PhD-level science questions, competition mathematics, and multimodal reasoning. In the same year, global corporate AI investment reached \$581.7 billion — a 130% increase from the prior year.²

These are not projections. These are data points from the past twelve months. The gap between what AI can do and what the institutions around it — education systems, governance frameworks, labor markets — are prepared for seems to be widening.

Capability: What Has Changed

The pace of change since late 2022 has been without precedent in the modern technology industry. Generative AI reached 53% of global population adoption within three years — faster than the personal computer or the internet. But adoption is not uniform: uptake correlates strongly with GDP per capita.¹ Organizational adoption has kept pace: 88% of companies now use AI in at least one business function, up from 55% two years ago.³

The speed of this adoption has no historical precedent. ChatGPT reached one million users in five days and 100 million in two months — the fastest growth curve in consumer technology history. TikTok took nine months and Instagram two and a half years to reach the same milestone. UBS analysts noted they could not recall a faster ramp in a consumer internet app in twenty years of following the internet space.⁴ By February 2026, ChatGPT had reached 900 million weekly active users — more than double its base of twelve months earlier.⁵

The capability frontier is also uneven in ways that matters for anyone deploying these tools. Stanford HAI researchers describe it as "jagged": the same model that wins gold at the International Mathematical Olympiad reads analogue clocks correctly only 50.1% of the time. Extraordinary capability in one domain does not predict reliable performance in another. For organizations, this means that a model's published test results are a poor guide to how it will perform on the specific work they care about — and that the gap between what AI achieves in a controlled test and what it delivers in practice remains significant.

One further signal of the moment's significance: the US-China performance gap, which defined AI competition for most of the past decade, has effectively closed. As of March 2026, the top six positions in community model rankings are held by Anthropic, xAI, Google, OpenAI, Alibaba, and DeepSeek – separated by margins of 3 percentage points or less.¹ The technology is no longer a single-country phenomenon.

The dissemination story adds a further dimension. China's domestic AI adoption – accelerated by DeepSeek's open-source release and the aggressive rollout of low-cost domestic models – is moving at a pace that capability benchmarks alone do not capture. And the geography of adoption is shifting faster than most assume: India and Brazil now rank second and third globally in ChatGPT usage, ahead of every European nation – Brazil ranks second on the use of OpenAI's APIs, reflecting a broader pattern in which the fastest growing AI markets are no longer in the Global North.

88%

of organizations globally now use AI in at least one business function – up from 55% two years ago. Yet 94% report not yet seeing significant enterprise-wide value from those investments.

McKinsey, The State of AI in 2025 (November 2025); Stanford HAI, AI Index 2026

\$581.7B

in global corporate AI investment in 2025 – a 130% increase year on year. Generative AI alone captured \$170.9 billion, growing over 200% from 2024.

Stanford HAI, AI Index Report 2026 (April 2026)

53%

global population adoption of generative AI within three years of launch – faster than the PC or the internet, though adoption correlates strongly with GDP per capita.

Stanford HAI, AI Index Report 2026 (April 2026)

The Economy and Jobs: What the Evidence Shows

The economic picture is real but uneven. Stanford HAI documents productivity gains of 14–26% in customer support and software development — measured in output per worker hour — and up to 72% in marketing output.¹ McKinsey estimates AI's long-term productivity growth potential across corporate use cases at \$4.4 trillion.⁶ These numbers are not evenly distributed: AI high performers — approximately 6% of organizations — already attribute more than 5% of earnings before interest and tax to AI. The majority are still in pilot mode.³

On jobs, the forward-looking assessment from the World Economic Forum's Future of Jobs Report 2025 projects that 170 million new roles will be created and 92 million displaced by 2030 — a net gain of 78 million jobs, but a structural churn of 22% of the formal workforce.⁷ The same report finds that 39% of existing skill sets will be outdated by 2030, and that 63% of employers already cite skills gaps as their primary barrier to transformation.⁷

The near-term data tells a more immediate story. In 2025, nearly 55,000 job cuts in the US were directly attributed to AI — out of 1.17 million total layoffs, the highest since the pandemic.⁸ In Q1 2026 alone, nearly 80,000 tech workers were laid off, with almost half of those cuts explicitly linked to AI-driven automation.⁹ Anthropic CEO Dario Amodei has warned that AI could eliminate up to 50% of white-collar jobs within five years — a view echoed in a February 2026 scenario paper by Citrini Research, which modelled a potential feedback loop between AI-driven white-collar displacement, falling consumer demand, and economic contraction.¹⁰ Economists have questioned the timeline, but not the direction of travel.

The IMF's analysis of 125 countries adds the distributional dimension that aggregate job counts obscure. Nearly 40% of global employment is exposed to AI — rising to 60% in advanced economies where cognitive-intensive roles predominate.¹¹ The IMF's assessment is that AI is likely to exacerbate labor income inequality if the complementarity between AI and high-income workers is strong, while capital returns will simultaneously increase wealth inequality. Emerging markets and developing economies face a structural disadvantage — they lack the digital infrastructure and institutional capacity to capture AI's productivity benefits while remaining exposed to its displacement effects.¹¹

The distributional consequences are already visible in labor markets. PwC's 2025 Global AI Jobs Barometer — drawing on close to a billion job advertisements across six continents — found that workers with AI skills command a 56% wage

premium over peers in equivalent roles without those skills, up from 25% the prior year.¹² Industries most exposed to AI are generating three times the revenue growth per employee of those least exposed. The gap between AI-skilled and non-AI-skilled workers is a current and widening reality.

The perception gap compounds the policy challenge. Stanford HAI data shows that 73% of US AI experts view AI's impact on the job market positively. Only 23% of the general public agrees.¹ The people closest to the technology and the people most affected by it are not having the same conversation.

Projections	Source
170 million jobs created, 92 million displaced by 2030 – net gain of 78 million	WEF, Future of Jobs Report 2025
39% of job skills will be outdated by 2030	WEF, Future of Jobs Report 2025
40% of global employment exposed to AI; 60% in advanced economies	IMF, Cazzaniga et al., SDN/2024/001
Productivity gains of 14–72% documented across sectors	Stanford HAI, AI Index Report 2026
73% of AI experts positive on job market impact; 23% of general public	Stanford HAI, AI Index Report 2026
63% of employers cite skills gaps as primary transformation barrier	WEF, Future of Jobs Report 2025
56% wage premium for AI-skilled workers; 3x revenue growth per employee in AI-exposed industries	PwC, Global AI Jobs Barometer 2025

The Institutional Gap

One thing is clear: AI capability is moving faster than the systems around it. That is not unusual for a technology this new — organizations and governments are adapting, learning, and course-correcting in real time. But the data suggests the gap is wide and the pace of adaptation, which is naturally uneven across geographies, is insufficient.

The Stanford HAI AI Index 2026 documents this at the level of the technology itself. On transparency, the Foundation Model Transparency Index — which rates AI companies on a 100-point scale based on how openly they disclose their models' training data, capabilities, and risks — saw average scores drop from 58 to 40 points in a single year. The most capable models now disclose the least about how they work. On governance, the US ranks last among surveyed countries in public trust in its own government to regulate AI, at just 31%. On education, four in five university students now use AI for their coursework, but only 6% of teachers say their institution's AI policies are clear. The technology has arrived in classrooms, offices, courtrooms, boardrooms, and beyond. The frameworks for managing it have not.

On the workforce side, McKinsey's research identifies a consistent pattern: the redesign of workflows, not the adoption of AI tools, is what drives measurable business value. Organizations that treat AI as a capability to layer onto existing processes see marginal gains. Those that redesign how work is done see compounding advantages.⁴ The majority of organizations have not yet made that transition. Most are running pilots. Almost none have achieved what McKinsey defines as AI maturity — AI fully integrated into workflows, driving substantial and measurable business outcomes.³

Gallup's State of the Global Workplace 2026 — drawing on workers across more than 160 countries — finds that only 12% of employees in organizations that have deployed AI strongly agree it has transformed how work gets done. The gap is also deepening along existing fault lines. Frequent AI use stands at 44% among organizational leaders compared to only 23% among individual contributors.¹³

The training supply gap sharpens this picture further. The OECD's 2025 analysis of AI skills training across member countries found that most AI-related training is disproportionately focused on AI professionals — with higher prerequisites than average courses — suggesting that current provision is more targeted towards higher-skilled adults than the broader workforce that needs it most.¹⁴

This is the moment the GFCC's survey captures: a technology advancing faster than the human systems around it, with productivity gains real but unevenly distributed, and an institutional architecture — in education, governance, and labor markets — that was not built for what is now required of it.

The gap between AI capability and institutional readiness is not a temporary lag. It is the defining challenge of the transition.

Why This Survey, Why Now

Against this backdrop, the GFCC surveyed leaders across five sectors and around the world to ask a question that the technical and economic data alone cannot answer: how are the people responsible for navigating this transition experiencing it? What do they believe is at stake? What do they think needs to change?

The data in this chapter establishes the scale of the challenge. The chapters that follow present what leaders across 32 countries believe should be done about it.

Beyond the Survey

The qualitative feedback sessions highlighted the need to monitor how AI investments evolve vis-à-vis AI's impact on organizations. At the moment, ninety-four percent of organizations report not yet seeing significant enterprise-wide value from their AI investments³, a gap that matters beyond the balance sheet. If expectations continue to run ahead of realized productivity, income growth and tangible human benefit, the result may be disappointment, backlash or a market correction that sets the broader agenda back and forces adjustments across the AI economy, including the valuations of AI labs and corporations heavily invested in the underlying technologies that power it.

The discussions made the case for clear-eyed realism. AI progress should be judged not by investment flows, adoption rates, or model benchmarks, but by whether productivity gains diffuse broadly, whether job quality improves, and whether people experience greater prosperity.

Notes

1. Stanford Institute for Human-Centered AI, AI Index Report 2026, April 2026. hai.stanford.edu/ai-index/2026-ai-index-report
2. Stanford HAI, AI Index Report 2026. Global corporate AI investment data.
3. McKinsey & Company, The State of AI in 2025: Agents, Innovation, and Transformation, November 2025. mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
4. TechnologyChecker.io, ChatGPT Statistics 2026, May 2026. technologychecker.io/blog/chatgpt-statistics
5. UBS, as cited in Time Magazine, "How ChatGPT Managed to Grow Faster Than TikTok or Instagram", February 2023. time.com/6253615/chatgpt-fastest-growing
6. McKinsey & Company, The Economic Potential of Generative AI: The Next Productivity Frontier, June 2023; and AI Productivity Gains and the Performance Paradox, April 2026.
7. World Economic Forum, Future of Jobs Report 2025, January 2025. reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf
8. Challenger, Gray & Christmas, Job Cuts Report, October 2025.
9. Nikkei Asia / Tom's Hardware, Tech Industry Layoffs Q1 2026, April 2026. tomshardware.com/tech-industry/tech-industry-lays-off-nearly-80-000-employees-in-the-first-quarter-of-2026
10. Citrini Research, The 2028 Global Intelligence Crisis, February 2026. citriniresearch.com/p/2028gic
11. IMF Staff Discussion Note SDN/2024/001, Gen-AI: Artificial Intelligence and the Future of Work, Cazzaniga et al., January 2024. imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf
12. PwC, Global AI Jobs Barometer 2025, June 2025. pwc.com/gx/en/services/ai/ai-jobs-barometer.html
13. Gallup, State of the Global Workplace 2026. gallup.com/workplace; and Gallup Workforce Panel Survey, Q4 2025. gallup.com/workplace/701195/frequent-workplace-continued-rise.aspx
14. OECD, Bridging the AI Skills Gap: Is Training Keeping Up?, April 2025. doi.org/10.1787/66d0702e-en



CHAPTER 2

**How the World
Sees the
Human-AI Co-
evolution: What
Leaders Across
32 Countries
Told Us**

Between December 2025 and April 2026, the GFCC surveyed leaders from 32 countries – government officials, business executives, academics, and civil society representatives – asking them how they are experiencing the rise of intelligent systems. What emerged is a mosaic of cautious optimism, institutional anxiety, and a shared conviction that the window to act is narrowing.

This chapter presents the headline findings from that survey. It is not a comprehensive data appendix – that sits in the Annex. It is a window into how global leaders across sectors and geographies are thinking about human-AI co-evolution in 2026: what they fear, what they hope for, and where they believe the leverage lies.

A Snapshot of Who Responded

The respondent pool skewed senior: 77% hold C-level, senior leadership, or senior management roles. Industry and business contributed the largest share (42% respondents), followed by academia and research (25%), civil society (11%), government and policy (7%), and other sectors (15%). Expertise concentrated on technology and AI, innovation and entrepreneurship, and economic development.

Geographically, responses came from six continents. Brazil represented the largest single-country group; a sensitivity analysis excluding this cohort confirmed no significant impact on the overall findings. The United States, the United Kingdom, India, and the UAE rounded out the top five by response volume.

Sentiment: Transitional, Not Catastrophic

When asked to characterize the current relationship between humans and AI in their context, the largest group of respondents (42%) chose "transitional": AI is changing the nature of work and skills. A further 27 described the relationship as "predominantly supportive", with AI enhancing human capability. Only 22 respondents selected "disruptive" – AI is displacing human roles and creating uncertainty – and 11 said AI adoption in their context remains underdeveloped.

Global leaders are not panicking about AI. They are navigating it – and finding the terrain uncertain.

Confidence in the prospect of a positive co-evolution tells a similar story. Asked to rate their confidence on a scale of 1 to 5, the center of gravity settled between 3 and 4, with an average score of 3.5 across all respondents: cautious optimism rather than either alarm or euphoria, and no strong clustering at either extreme.

This leadership cohort mostly believes that positive co-evolution is achievable – but they do not believe it will happen by default. The question of whether it happens is a question of policy, investment, and institutional will.

Where Leaders See the Opportunity

Respondents were asked to identify the top opportunities for AI to augment human capability and competitiveness. The responses converged strongly into five themes, together accounting for the overwhelming majority of all mentions:

1	Productivity & Task Automation ~62% of the respondents	AI removes repetitive, low-value cognitive work – freeing humans for creative and strategic tasks. The most widely cited opportunity by a significant margin.
2	Decision-Making & Data Analysis ~44% of the respondents	AI synthesizes vast data at speed to surface insights and supports faster, better-informed decisions at every level of an organization.

3	Education & Personalized Learning ~35% of the respondents	AI enables hyper-personalized education at scale, acting as tutor and skills accelerator for learners of all ages and backgrounds.
4	Healthcare & Scientific Discovery ~27% of the respondents	AI accelerates diagnostics, drug discovery, and research capacity, while expanding the reach of healthcare into underserved communities.
5	Democratization of Expertise ~18% of the respondents	AI levels the playing field – giving individuals, small businesses, and emerging economies access to capabilities previously reserved for large organizations.

The first and fifth opportunities form a natural pair that runs through the entire survey: AI saves time for those who already have it, and AI can extend opportunity to those who have been so far excluded from it. Whether the second effect materializes depends on deliberate design choices – in technology, policy, and institutional architecture.

Where Leaders See the Risk: skills mismatch and education

Asked to select the main barriers to achieving human-AI complementarity, respondents painted a coherent picture of institutional and systemic failure rather than technological malfunction. The top concern, selected by 58% of respondents, was skills mismatch and education lag – the gap between the competencies AI-driven economies demand and the systems currently equipped to deliver them.

Ethical or bias concerns and lack of regulatory clarity or trust were each mentioned by 45% of the respondents – reflecting a shared anxiety that the governance architecture for AI is inadequate to the pace of change. Workforce displacement and job loss (mentioned by 42% of respondents) and the digital divide and unequal access (mentioned by 40% of respondents) completed the top five. Cultural resistance to change, while real, ranked last.



Top barrier:
Skills mismatch (mentioned by 51% of respondents)

Education and workforce systems are not producing the competencies that the AI economy demands – and the gap is widening faster than institutions can close it.



Close second:
Ethics & regulatory gaps (mentioned by 45% of respondents each)

Governance frameworks for AI are seen as inadequate to the pace of change – creating uncertainty that slows adoption and erodes public trust simultaneously.

A minority view worth noting – expressed by roughly 5-10% of respondents – flagged risks that went beyond the mainstream framing. These included structural unemployment, AI-amplified misinformation, and the possibility that AI would widen rather than close existing inequalities. These voices are worth tracking: in previous technological transitions, minority concerns at the early stage often became majority realities later.

Three Tensions the Data Reveals

Across all sections of the survey, three recurring tensions shaped the findings. They are not contradictions — they are the fault lines where action is most needed.

1

Optimism about AI's potential, pessimism about institutional readiness.

Respondents broadly believe in the upside of human-AI co-evolution. They are far less confident that the systems around them — education, governance, workforce institutions — are equipped to deliver it.

2

Recognition that AI is not inherently inclusive, combined with limited examples of deliberate inclusion by design.

Most respondents acknowledge that AI amplifies existing conditions. Yet examples of AI that actively reverses inequality — rather than merely not worsening it — remain rare and fragile.

3

Agreement on what needs to happen, disagreement on who is responsible for making it happen.

Cross-sector collaboration dominated the response to every question about shared responsibility. But several respondents noted that calling for collaboration without addressing misaligned incentives is wishful thinking.

The survey data reveals something more than a simple optimism-pessimism divide. Cross-analyzing confidence scores against other responses shows that high- and low-confidence respondents are not just differently hopeful — they see a fundamentally different AI. Those with low confidence (scores 1–2) flag workforce displacement as their top concern at nearly three times the rate of high-confidence respondents (58% vs 20%), while placing notably more weight on technical AI expertise as the critical human capability. High-confidence respondents, by contrast, lean heavily into adaptability and lifelong learning (61% vs 29%) and entrepreneurial mindset (47% vs 25%).

The dividing line, in other words, is not temperament. It is a disagreement about whether humans can adapt fast enough — and whether the distinctive value of human capability will be recognized and rewarded in an AI-driven economy. That disagreement runs underneath every section of this survey, and it is worth keeping in mind as the chapters that follow make the case for investing in the human edge.

Reading the Regional Divide

While the global aggregates are informative, the survey's most revealing moments came in the divergences. Respondents from high-income economies with strong digital infrastructure were more likely to describe the human-AI relationship as supportive or transitional. Those from the Global South, or from contexts with weak institutional capacity, were more likely to flag underdevelopment or disruption — and to raise concerns about whether global frameworks designed in wealthy nations would be applicable to their realities.

Several respondents from emerging economies specifically challenged the assumption embedded in the three-way responsibility model — government, business, universities — noting that in contexts where governments are under-resourced, universities are underfunded, and businesses are primarily focused on short-term survival, that framework may not be the right starting point.

The survey data reflects this divergence. Respondents from high-income economies were most likely to describe AI as transitional (36%) or supportive (25%) — consistent with contexts where adoption infrastructure exists. Respondents from lower-income economies were more divided: 27% described AI

as disruptive and 27% as supportive, suggesting a more polarized experience where AI is simultaneously enabling access to services previously out of reach and displacing roles with no transition pathway.

Notably, despite — or perhaps because of — that polarization, respondents from lower-income economies registered the highest average confidence in positive co-evolution (3.53 out of 5), compared to 3.30 in high-income countries. A different model — one that begins with community, digital infrastructure, and foundational access — may be needed for contexts where the institutional architecture the mainstream model assumes is still being built.

The collaboration model assumes functional institutions. In contexts where that assumption fails, the model needs rethinking from the ground up.

This is not a reason to abandon a global framework. It is a reason to ensure that framework has enough flex to accommodate the full range of institutional realities among GFCC's diverse membership.

What the Numbers Don't Capture

Quantitative survey data can map the terrain, but it cannot fully convey the texture of how people are experiencing this transition. The open responses in this survey contained something that bar charts cannot — a sense of urgency, of complexity, and in some cases of genuine anxiety about the speed of change relative to the speed of institutional adaptation.

The following chapter, *Voices of the Transition*, gives those individual perspectives more room. It does not resolve the tensions in the data — it makes them more vivid. That is precisely the point. This report is not a search for a single answer. It is an attempt to hold together the full range of leaders across 32 countries and four sectors, are actually experiencing. The complexity is the message.

Key Takeaways



- Most leaders see AI as transitional rather than purely disruptive – cautious optimism dominates, but it is contingent on institutional action.
- Productivity and task automation ranked first among perceived opportunities; skills mismatch ranked first among perceived risks – a clear signal that the human side of the transition is the priority.
- Ethical, governance, and equity concerns are not fringe anxieties; they are mainstream concerns held by more than half of respondents.
- A minority view – expressed by roughly 5-10% of respondents – flagged the risk that AI will widen rather than narrow inequality unless explicitly designed to do otherwise. In previous technological transitions, minority concerns at the early stage became majority realities later.
- The frameworks for measuring success – how we know whether human-AI co-evolution is actually working – are largely absent. Building them is urgent.



Voices of the Transition: Perspectives from the Global Community

CHAPTER 3

The previous chapter told us what global leaders think – in aggregate terms, across sectors and geographies. This chapter tries to convey something different: the texture of conviction, anxiety, and lived experience underneath the numbers. It does so through four personas, developed by clustering survey responses around recognizable patterns of belief and experience present across the respondent pool. Each persona is built from real responses and represents a worldview that appeared consistently enough to deserve its own voice.¹

Every phrase attributed to the personas draws directly from submitted answers. They are composites by design – each assembles the most coherent and representative expression of a perspective that appeared across multiple respondents, geographies, and sectors. The goal is giving form to the diversity of views, so the tensions within the data become human rather than statistical.

The numbers tell us what people think. The voices tell us why it matters to them.

Four personas emerged clearly from the qualitative response: **The Optimist**, who sees AI as a historic opportunity to democratize capability and accelerate human progress; **The Skeptic**, who acknowledges the potential but remains unconvinced that institutions will move fast enough to prevent serious harm; **The Systems Thinker**, who focuses less on AI itself and more on the governance, incentive structures, and cross-sector coordination required to make co-evolution work; and **The Inclusion Advocate**, who measures every development against a single question – who benefits, and who gets left behind.

They do not map neatly onto sectors, geographies, or confidence scores – The Optimist appears in Tokyo and Lagos; The Skeptic in London and São Paulo. What divides them is not where they sit but how they weigh up the evidence, which risks they find most credible, and what kind of action they believe will make the difference.

The Optimist

Sees the transition as fundamentally generative. Believes human creativity and AI capability are complementary, not competing. Frustrated by institutional slowness but not alarmed by it. Confident that if the right investments are made, the window is still open.

"Phase one: eliminate repetitive cognitive work that nobody should be doing. Phase two: democratize analytical capacity at scale. Phase three: create services that only exist economically because of AI."



Profile

**Business executive,
tech sector**

Industry / Business

Confidence in co-evolution

Rated 4 or 5 (out of 5) when asked about confidence in positive co-evolution between humans and AI

AI relationship

Predominantly supportive

"AI removes the grudge work — the repetitive, low-cognition tasks that consume human time and produce no meaning. What's left is the good stuff: creativity, judgement, connection. The opportunity is not to automate humans out of work. It is to automate humans into better work. The people who understand that are pulling ahead. The institutions that don't are falling behind, and they will keep falling behind until they decide to change."

CORE BELIEF

The friction between human potential and AI capability is real, but it is temporary as it is not a fundamental one — institutions lag, but they adapt. Curiosity and adaptability — not technical expertise — are the differentiators that determine who thrives.

THE TENSION THEY CARRY

Optimism is conditional. Several respondents in this persona noted that their confidence rests on a bet that institutions will act and adjust in time. If they don't — if education systems keep producing graduates for jobs that no longer exist, if businesses keep treating AI as a cost-cutting tool rather than a capability multiplier — the upside scenario closes. The Optimist is not naive. They are impatient.

The Skeptic

Not opposed to AI — often familiar with the technology, their skepticism comes from proximity, not distance. Skeptical of the dominant narrative that frames AI as an unambiguous opportunity. Believes the disruption is structural, the timeline is shorter than institutions admit, and the costs will fall disproportionately on those least able to absorb them.

"Every previous revolution looked like progress in the aggregate. The aggregate masked a lot of suffering."



Profile

**Senior researcher /
civil society leader**

Academia / Civil Society

Confidence in co-evolution

Rated 1 or 2 (out of 5) when asked about confidence in positive co-evolution between humans and AI

AI relationship

Disruptive

"We keep talking about reskilling as if it is a solved problem. It is not. The evidence from previous transitions is that retraining programs mostly don't work at scale, that the people who need them most are the least likely to access them, and that the window between displacement and redeployment is long enough to cause permanent damage — to careers, to communities, to trust in institutions. I am not saying AI is bad. I am saying we are running the experiment on real people without a safety net."

CORE BELIEF

Aggregate economic gains from AI are likely real. Distribution of those gains is not guaranteed — and historical precedent suggests it will be deeply unequal without structural intervention. The default trajectory is not prosperity for all; it is prosperity for some and dislocation for many.

THE TENSION THEY CARRY

The Skeptic often sees things others miss — the disruption of traditional apprenticeship models, the energy cost of AI infrastructure, the geopolitical concentration of AI capability. But their focus on what could go wrong sometimes leaves them without a positive program. Several respondents in this persona struggled to articulate what success would look like, not because they lack imagination but because they distrust the framing of the question.

The Systems Thinker

Sees the human-AI transition as a second-order governance and institutional design problem. The systems around us were already broken before AI arrived — misaligned incentives, coordination failures, institutions designed for a world that no longer exists. AI does not create these problems. It compounds them. Most current policy responses, in this view, are solving the wrong problem.

"The constraint is not the technology. It is our willingness to redesign the systems around it."



Profile

Economist / policy advisor / academic

Government / Academia

Confidence in co-evolution

Rated 3 (out of 5) when asked about confidence in positive co-evolution between humans and AI.

AI relationship

Transitional

"The three-sector model — government sets policy, business invests in people, universities reform curricula — is right in principle and largely broken in practice. Governments want jobs and votes. Businesses want efficiency. Universities want funding. None of those incentives naturally produces the outcome we are describing. You don't solve a misaligned-incentive problem by asking everyone to collaborate more. You solve it by redesigning the incentives. Until someone is willing to say that, we are going to produce a lot of excellent reports and very little change."

CORE BELIEF

The barriers to human-AI co-evolution are not primarily technical or even educational — they are institutional and political. Skills alliances and reskilling programs are necessary, but they will not be sufficient as long as the underlying incentive architecture remains unchanged: governments optimizing for short-term electoral cycles, businesses optimizing for quarterly returns, universities optimizing for research rankings. Good intentions coordinated badly still produce bad outcomes.

THE TENSION THEY CARRY

The Systems Thinker is often the most analytically rigorous voice in the room — and sometimes the most paralyzing. Several respondents in this persona raised

challenges so systemic that the implied conclusion was close to despair. But the most compelling version of this perspective does not stop at diagnosis. It couples structural critique with a concrete theory of change: start with the countries and institutions small enough to move, demonstrate what works, and build outward from there. The system cannot be redesigned all at once. It can be redesigned at the edges, repeatedly, until the edges become the center.

The Inclusion Advocate

Focuses primarily on who is left out of the AI transition and why. Sees AI not as inherently beneficial or harmful but as an accelerant of existing conditions — one that will widen inequality unless deliberately designed not to. Brings urgency that other personas often lack, grounded in direct experience of the communities most at risk.

"AI is not creating a new divide. It is turbocharging the ones we already have."



Profile

Civil society leader / non-profit director

Non-profit / Civil Society
/ Government

Confidence in co-evolution

Rated 2, 3 or 4 (out of 5) when asked about confidence in positive co-evolution between humans and AI. However, confidence depends entirely on whether inclusion is being actively designed in — not assumed.

AI relationship

Exclusionary or limited inclusivity

"The groups most at risk are the ones who were already excluded before AI arrived. The Global South. Women in clerical roles. Workers in routine jobs with no pathway to reskilling. Young graduates entering professions where AI already does the entry-level work — which means they never get the learning-by-doing experience that built the expertise of the generation above them. We are not just risking displacement. We are risking the permanent break of the apprenticeship model that creates expertise. That is a civilizational problem, not just an economic one."


CORE BELIEF

The trajectory of AI is not fixed. Inclusive outcomes are possible — Singapore built national AI literacy programs that reach workers across every age group and career stage; Romania's Nextlab.tech brought robotics and AI education to rural and economically disadvantaged schools through low-cost hardware and human mentorship; Alberta's Amii used AI-powered diagnostics to extend healthcare to remote and Indigenous communities; and Brazil's Piauí program integrated AI ethics and digital literacy into public schools in one of the country's least resourced states. What these cases share is not scale or budget — it is deliberate design. But inclusive outcomes are not automatic. They require policy choices, funding commitments, and institutional designs that actively redistribute access rather than allowing market forces to concentrate the benefit at the top.

THE TENSION THEY CARRY

The Inclusion Advocate carries the weight of urgency most acutely — and sometimes most exhaustingly. Several respondents in this persona flagged that the examples of successful inclusion they could cite were small-scale, fragile, or dependent on exceptional local leadership. Scaling what works in Singapore to what is needed in sub-Saharan Africa or rural South America is not just a matter of replication. It requires a level of sustained global political will that has not yet materialized.

The four personas can be mapped across two dimensions – their level of confidence in positive co-evolution, and whether their focus is primarily on individual adaptation or systemic change.



What These Four Voices Reveal

The four personas agree on more than they disagree. All four see education and skills as the central battleground — the place where the human-AI transition will be won or lost. All four believe the window for action is narrowing. All four, in different registers, express frustration with the pace of institutional response relative to the pace of technological change.

What they reveal, taken together, is that every one of them is making a different bet on human agency. The Optimist bets on individuals adapting faster than institutions. The Skeptic bets that without structural support, adaptation will not be enough. The Systems Thinker bets on redesigning the institutional environment. The Inclusion Advocate bets on deliberate investment in the communities, demographics, and regions farthest from the frontier. They disagree about which bet is the right one — but none of them has abandoned the idea that human choices still determine the outcome.

What follows in the next four chapters is an attempt to honor all four voices: to make the case for human capability investment that the Optimist calls for, to take seriously the structural risks the Skeptic names, to apply the institutional lens the Systems Thinker demands, and to keep the inclusion imperative the Inclusion Advocate insists on at the center.

Note

I. To estimate the prevalence of each persona in the full respondent pool, a scoring model was applied to closed-question data. Each respondent received points toward each persona based on their answers to Q6 (AI relationship), Q7 (confidence in co-evolution), Q9 (main barriers), Q10 (critical human capabilities), and Q14 (AI inclusivity). Respondents were assigned to the persona with their highest aggregate score. Distribution: **The Optimist** 45 (40%), **The Systems Thinker** 26 (23%), **The Inclusion Advocate** 22 (19%), **The Skeptic** 20 (18%). Of these, 54 (48%) were clearly assigned with a score margin of 2 or more points over the next-closest persona; 37 were tight calls (margin of 1); and 22 were tied between two personas. The Optimist plurality is the most robust result across alternative scoring specifications. The three-way split among the remaining personas should be treated as indicative. The high proportion of respondents sitting at persona boundaries (52%) reflects a genuine feature of the data: these worldviews are not mutually exclusive, and the survey population includes many respondents whose perspectives draw on more than one persona simultaneously.



The Human Edge: What Humans Bring to the Age of AI

CHAPTER 4

There is a question underneath every section of this survey, and it is not about AI. It is about humans. Specifically: what do humans bring to an economy increasingly run by intelligent machines – and will those contributions be recognized and rewarded, or quietly devalued as machines become more capable?

The respondents to this survey gave a clear answer. Creativity and innovation topped the list of critical human capabilities, selected by 65% of the respondents. Critical thinking and ethics came second, being mentioned by 62% of the sample. Adaptability and lifelong learning – the capacity to learn continuously – ranked third. Technical AI expertise – the ability to build and engineer AI systems – came last, chosen by just 18% of the respondents.

In a moment of rapid AI capability growth, leaders in this survey pointed consistently to the capabilities that are hardest to automate: judgement, imagination, the ability to adapt, the willingness to keep learning. They were not dismissing technical expertise. They were saying it is not the differentiator.

The Capability Rankings

The survey offered respondents eight human capabilities and asked them to select up to five they considered most critical for competitiveness in an AI-driven economy. The table below shows how each capability ranked across the full respondent pool, with a confidence split: the percentage of high-confidence respondents – those who scored 4 or 5 on the co-evolution confidence question – who selected each capability, versus the percentage of low-confidence respondents (scores 1-2) who did the same. The gap between those two groups reveals something important: the capabilities that most divide optimists from pessimists are not the ones you might expect. The table is ranked by that gap, largest to smallest.

	Capability	What this means in an AI economy	Optimist–Pessimist gap
1	Adaptability & lifelong learning 53% of respondents	The capacity to learn new things continuously and reorient when conditions change — the meta-skill that makes all other skills renewable.	+32pp hi 61% · lo 29%
2	Entrepreneurial mindset 36% of respondents	The disposition to identify opportunities, tolerate uncertainty, and build something new — increasingly important as AI lowers barriers to creation.	+22pp hi 47% · lo 25%
3	Emotional intelligence & collaboration 42% of respondents	Reading people, managing relationships, building trust, and coordinating across differences — the social infrastructure of organizations.	+12pp hi 37% · lo 25%
4	Digital & AI literacy 50% of respondents	Understanding how AI systems work, when to trust them, when to question them, and how to work alongside them effectively.	+9pp hi 51% · lo 42%
5	Critical thinking & ethics 62% of respondents	Evaluating evidence, identifying flawed reasoning, and applying moral judgement to decisions with real-world consequences.	+9pp hi 59% · lo 50%
6	Complex problem-solving 43% of respondents	Navigating multi-variable, ambiguous problems where context, stakeholder dynamics, and incomplete information all matter.	+7pp hi 49% · lo 42%
7	Creativity & innovation 65% of respondents	Generating ideas, making unexpected connections, and producing original work — tasks where human variability is a feature, not a flaw.	+3pp hi 61% · lo 58%
8	Technical AI / engineering expertise 19% of respondents	The ability to build, fine-tune, and maintain AI systems. Valued by fewer respondents overall — and significantly more by pessimists than optimists.	-21pp hi 12% · lo 33%

Two findings on this table deserve particular attention. First, adaptability and lifelong learning show the largest gap between high- and low-confidence respondents: 61% of optimists selected it, against only 29% of pessimists — a 32-percentage-point difference. The belief that humans can keep relearning is possibly optimism about education and the mechanism by which optimists expect the positive scenario to materialize.

Second, technical AI expertise runs in the opposite direction, chosen by 33% of low-confidence respondents but only 12% of high-confidence ones. One possible interpretation is that pessimists are reaching for technical depth as a response to uncertainty. But a more accurate reading is possibly that respondents who flagged technical expertise were thinking about something closer to AI literacy in practice — the ability to understand how systems work, identify when outputs are unreliable or biased, and exercise meaningful oversight rather than passive acceptance. That framing is consistent with the same group's elevated concern about ethical and regulatory gaps: if you're worried about AI governance failures, human judgement becomes imperative to understand the technology well enough to critique it.

A system not built for what it needs to produce

The survey also asked whether current systems are producing them. The answer was unambiguous. Eighty-nine per cent of respondents described education and workforce systems as either "poorly" or "somewhat inadequately" prepared for the AI transition. Only 10 respondents said systems were adequate. Not a single respondent said "very well."

89% say systems are inadequately prepared — an opinion shared equally by optimists and pessimists, government and industry, the Global South and the Global North.

What makes this finding striking is its consistency across sector, confidence level, and view of the AI relationship. Government officials: 88% negative. Industry executives: 89%. Civil society leaders: 92%. Even the most optimistic respondents — those with confidence scores of 4 or 5 — returned an 86% negative rating. This is not a finding on which the entire leadership cohort agrees.

The implication is direct. The capabilities the table above identifies as critical are not being cultivated at the scale the transition requires. The gap is not a perception problem. It is a systems problem.

What Needs to Change – and the Urgency Behind It

Respondents were asked what key reforms or investments are needed to align human capital systems with the AI economy. The respondents converged on five priorities that diagnose what is missing.

Reform priority		What should be included
1	Education system overhaul 72 mentions	Reform curricula from K–12 through university, integrating AI literacy, critical thinking, systems thinking, and problem-solving as foundational skills rather than electives.
2	Curriculum and skills reform 55 mentions	Shift from static degrees and rote learning to stackable credentials, skills-based hiring, and adaptive learning aligned to AI-era competencies.
3	Infrastructure and digital access 45 mentions	Large-scale reskilling programs, portable learning accounts, and continuous upskilling pathways for workers at all career stages.
4	Infrastructure and digital access 28 mentions	Invest in connectivity, compute access, and universal access to technology and AI tools as a foundation for participation
5	Governance, regulation, and ethics 22 mentions	Develop regulatory frameworks, ethical guidelines, and labor protections to manage AI-driven workforce disruption

Education reform dominated by a wide margin — nearly two-thirds of all mentions. But the most analytically interesting responses pushed beyond the obvious. Several respondents argued that the framing of "reform" itself is insufficient: the scale of disruption may require reimagining human capital systems from scratch rather than updating existing ones. One respondent argued that the transition requires moving from a human-centered world of work to one where the fundamental unit of value has changed — a shift that conventional education reform does not address. Another respondent raised the possibility that if AI approaches the end of work as we know it, preparing people for a "transition" may be the wrong frame entirely.

Two specific reform ideas surfaced repeatedly in the open responses that the aggregate counts do not fully capture. The first was portable benefits and social protection — decoupling healthcare, retirement, and income support from specific employers to support a workforce in permanent transition. The second was systems thinking as a core literacy — teaching people to reason about second and third-order effects, non-linear causality, and emergent behavior, not just AI tools. This is a fundamental skill reserved that many education systems have simply failed to prioritize — and one that becomes more valuable as the problems people face grow more interconnected and harder to predict. One respondent made the case most fully: most of the problems people will face are systems problems, and almost no education system currently teaches people to think that way.

What Leaders Actually Say About Human Advantage

The open responses in this survey gave the capability rankings texture and argument. Several stood out for the precision and conviction with which they made the case for the human edge.

"Curiosity is now the single highest-leverage human attribute. Every question you ask returns real, actionable value. But curiosity isn't something most educational systems nurture — they optimize for compliance and correct answers. People who are genuinely curious will compound their capabilities exponentially in an AI-augmented environment. Those who aren't will plateau regardless of what tools they have access to."

"AI accelerates the decay of static content knowledge. What lasts is problem-solving, critical thinking, and the ability to learn fast, work with machines, and adapt on the fly. The shift from knowing to learning is the central challenge of the transition — and most education systems are built for the former."

A recurring theme in the open responses was the distinction between knowing and doing — between holding information and being able to act on it in ways that create value. Several respondents articulated this as the collapse of the barrier between understanding something and building something, made possible by AI tools that handle technical execution:

AI enables non-experts to analyze data through natural language, making creative vision the primary driver of competitiveness.

The entrepreneurial dimension of this argument surfaced repeatedly. If AI lowers the barrier between idea and product, between insight and implementation, then the humans who will benefit most are not necessarily those with the deepest technical skills — they are those with the sharpest vision of what to build and the agency to pursue it. One respondent framed this directly: *the barrier between understanding something and creating something is collapsing. The people who thrive will be those who do the most with what they know.*

The Uncomfortable Truth

The case for the human edge is compelling. But it comes with a condition that this survey's respondents named clearly and that this report will not soften: the capabilities that matter most are not equally available. And beneath that observation sits a harder question that the data cannot fully answer but that every policymaker and institution leader must confront: how can these capabilities actually be developed at scale? Curiosity, adaptability, and systems thinking are not innate traits. They emerge — or fail to emerge — in response to the conditions people grow up and work in. If those conditions are not deliberately built, the human edge remains real but narrow — a competitive advantage for the few rather than the foundation of broad prosperity.



The Human Edge Needs Infrastructure

Creativity, critical thinking, and adaptability are present in human beings across every geography and background. But they are not equally cultivated. The conditions that develop these capabilities — high-quality education, autonomous and meaningful work, psychological safety, time to think, exposure to diverse ideas — are distributed with profound inequality.

This matters for the AI transition because AI will reward the human edge handsomely — but only for those who have it or have access to the infrastructure to properly develop it. For those who don't, the AI transition might offer even fewer compensations than previous technological transitions did. The current wave is moving up the skill ladder faster, compressing the space between displacement and the new frontier.

The human edge is not a given. It is something societies choose — or fail — to develop.

Research evidence bears this out. A field experiment in Turkey found that students using GPT-4 performed up to 48% better on tasks while they had access – but 17% worse once access was removed.¹¹ Task performance improved; genuine learning did not. It is a small but striking illustration of the gap between capability and dependency that sits at the heart of this challenge.

The disruption of traditional apprenticeship models – raised by several respondents and discussed in detail in Chapter 5 – is perhaps the most concrete expression of this problem. If AI eliminates entry-level roles in law, finance, accounting, and other knowledge professions, the pathway by which young people build the deep expertise that eventually becomes a human edge disappears. The senior lawyer's judgement is valuable precisely because it was built on years of doing the work that AI now does instead. Remove the learning-by-doing foundation and the expert edifice above it becomes harder to build.

One respondent articulated this with particular clarity, asking how new graduates will acquire the experiential expertise their predecessors developed when AI already performs the entry-level work. Another framed it as a generational problem: young graduates entering professions where AI co-pilots make senior employees more productive have no clear path to becoming those senior employees.

Three Investments the Survey Calls For

From the capability rankings, the cross-analysis, and the open responses, three investment priorities emerged with clarity and consistency across sectors and geographies.

1 Redesign education around durable human capabilities

The most consistent reform signal in the survey was an overhaul of education systems from K-12 through university — not to teach AI, but to teach the capabilities that complement it. Respondents called for critical thinking, systems thinking, problem-solving, creativity, and communication to become foundational across all subjects and disciplines, not electives or extras. Several pushed further: one respondent argued that curiosity should be cultivated as a core value, not just a trait. It captures the spirit of what many respondents were reaching for. An education system that rewards compliance and correct answers is actively counter-productive in an AI economy.

2 Build lifelong learning as infrastructure, not a program

Adaptability and lifelong learning ranked third overall — and showed the largest gap between high- and low-confidence respondents. The respondents who believe in positive co-evolution believe in humans' capacity to keep learning. This is not an argument for more training programs. Several respondents were explicit that reskilling programs mostly don't work at scale, that the people who need them most are least likely to access them, and that the window between displacement and redeployment is long enough to cause permanent damage. What is needed is infrastructure: portable learning accounts, stackable credentials, modular pathways, and the financial support that allows people to retrain without falling off an economic cliff.

3 Redesign work to cultivate human capability, not just deploy it

Several respondents pushed beyond education reform to argue that the organization of work itself needs to change. Human capabilities like creativity, judgement, and collaboration are not fixed quantities that education delivers and employers consume. They are cultivated — or eroded — by the conditions under which people work. Autonomous work, meaningful tasks, psychological safety, and time to think are the environmental conditions in which human capability compounds. Workplaces that strip these out in pursuit of short-term efficiency — reducing humans to monitors of AI outputs — will hollow out the human edge they depend on. One respondent put it simply: the less you focus on the technology and the more you focus on human purpose, the more impactful your contribution will be.

Beyond the Survey

Feedback sessions reinforced the report's central argument about the human edge and deepened the conversation around what it actually takes to develop the capabilities that matter most. There was broad agreement that creativity, critical thinking, ethics, adaptability, collaboration and judgment are the key differentiators in an AI-driven economy. The harder question that emerged from the discussions is how to cultivate these capabilities deliberately in environments where digital platforms are ubiquitous and AI is increasingly present.

Educators, institutions and societies are still learning what works and what does not when it comes to the use of AI in education and its effects on learning. In some settings, AI may encourage passivity, dependency and shallow engagement rather than deeper thinking. In light of this, the design of educational environments and learning experiences takes on critical importance.

The experience of digital technology in education also offers a useful caution: access alone does not produce capability. How technology is introduced, what it replaces and what it leaves room for matter just as much as whether it is present at all. The challenge, therefore, is not simply to put AI into classrooms and workplaces. It is to design environments that use AI to deepen human capability rather than substitute for it.

That design challenge extends to assessment. If assignments and tests can be completed with a simple AI prompt, the issue is not cheating. It is that the assessment model itself has been exposed as measuring the wrong things. Curriculum reform and micro-credentials will only be credible if what is being assessed reflects genuine capability: curiosity, critical thinking, judgment, applied problem-solving and the responsible use of AI tools.

The question is not only what humans can do with AI, but also what kind of humans our AI-enabled systems are helping to create.

A Signal Worth Watching

The cross-analysis in this survey produced one interestingly unique finding: the belief in human adaptability is the load-bearing pillar of confidence in positive AI co-evolution.

Respondents who selected adaptability and lifelong learning as a critical capability were 32 percentage points more likely to hold high confidence in the co-evolution scenario. No other capability showed a gap of that magnitude. The implication for policy is direct: if you want societies to navigate this transition with confidence, the single most powerful investment is in building the conditions that make lifelong learning genuinely accessible: portable credentials, funded learning accounts, employers who reward curiosity, and education systems that develop adaptability rather than just measure compliance.

That reality does not yet exist at the scale this transition requires. It is what Chapter 5 turns to next.

Chapter 4

Key Takeaways



- Creativity, critical thinking, and adaptability ranked far above technical AI expertise – the survey's clearest signal that uniquely human capabilities, not technical depth, are the perceived differentiator.
- Adaptability and lifelong learning shows the largest confidence gap of any capability (+32 percentage points) between respondents who are highly confident in positive co-evolution and those with low confidence. Among the first group, 61% selected this capability, compared to only 29% of those with low confidence. Believing humans can keep relearning is the mechanism by which optimists expect positive co-evolution to materialize.
- Pessimists lean toward technical AI expertise as the solution – suggesting they see the future as a competition with machines rather than a collaboration alongside them.

- The human edge is real but unequally distributed. The conditions that cultivate it – quality education, meaningful work, time to think – are not available to everyone. Underneath all the discussion of advanced capabilities and reskilling programs sits a more fundamental requirement that is easy to overlook: without a baseline of quality foundational education – literacy, numeracy, basic digital access – none of the skills development agenda has any ground to stand on.
- The disruption of traditional apprenticeship models is an underappreciated structural risk: if AI eliminates entry-level cognitive work, the pathway by which people build the expertise that becomes their human edge may close before a replacement pathway opens. Several respondents hinted at cross-generational cooperation as a potential response – pairing experienced workers with younger entrants who bring AI fluency but lack the accumulated judgement. The opportunity is real but the mechanisms to enable it at scale do not yet exist. Building those mechanisms is one of the biggest design challenges of the transition.
- Three investments follow from the evidence: redesign education around durable capabilities; build lifelong learning as infrastructure, not programs; and redesign work itself to cultivate human potential rather than simply consume it.

Note

II. OECD (2026), OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education, OECD Publishing, Paris. <https://doi.org/10.1787/062a7394-en>



CHAPTER 5

The Inclusion Gap: AI's Unequal Footprint

AI is not destined to widen inequality. But without deliberate action, it will. That is the verdict of 70% of this survey's respondents who described the current AI transformation in their context as either limited in inclusivity or actively exclusionary. Only 9% said it was highly inclusive. The gap between where we are and where we need to be is the central challenge of the transition.

70%

of respondents describe AI transformation in their context as having limited inclusivity or actively exclusionary impacts – only 9% say it is highly inclusive.

$r=0.38$

The strongest correlation in the survey: confidence in positive AI co-evolution tracks closely with whether someone believes AI is inclusive in their context. Pessimism and exclusion perception travel together.

100%

of respondents from the non-profit and civil society sector – those closest to the communities most at risk – described AI transformation as limited or exclusionary. Not one saw it as moderately or highly inclusive.

The Inequality Amplifier

The concept that recurred most consistently across the open responses was not displacement or disruption. It was amplification. Respondents across geographies and sectors described AI as potentially accelerating existing inequalities – turbocharging the gaps in digital access, education quality, institutional capacity, and economic power that were already present before the first language model was trained.

This framing has a specific and important implication. If AI were creating genuinely new inequalities, the response would be to mitigate those specific harms. But if AI is primarily an accelerant of existing conditions, the response must be structural: address the underlying inequalities that AI will amplify, not just the AI-specific symptoms.

The cross-analysis in this survey adds a further dimension. When we compared respondents' confidence in positive AI co-evolution with their perception of how inclusive the current AI transformation is, the two tracked together more strongly than any other pair of variables in the dataset – a correlation of 0.38. In plain terms: respondents who see the AI transformation as broadly inclusive tend to be significantly more confident that co-evolution will go well; those who see it as exclusionary or unequal tend to be significantly less confident. This is not a trivial finding. It suggests that the inclusion question and the optimism question are the same debate. How someone experiences or perceives the distribution of AI's benefits shapes their entire outlook on whether the transition can succeed. A significant implication is that building confidence in the co-evolution scenario requires making visible and tangible progress on inclusion.

Who Is Most at Risk

Respondents identified five groups and geographies as bearing disproportionate risk from AI transformation. The survey coded these from open responses across respondents who answered the question (81%). They are presented here in order of the frequency with which they were named and with the specific mechanisms of exclusion that respondents identified.

Beyond the Survey

The survey's findings on cross-sector collaboration and shared responsibility resonated strongly with feedback session participants – and sparked a broader insight that extends beyond the AI transition: the places best positioned to seize the opportunities ahead are those where stakeholders across government, business, universities, and civil society are actively shaping the future together.

Where mechanisms exist for meaningful public-private partnerships and cross-sector collaboration – where institutions are trusted, accountability is real, and different actors can align around shared goals – communities and countries are better equipped to capture the benefits of AI while managing its risks. The reverse is also true. In places with weak institutions, low enforcement capacity, or limited public trust, AI risks amplifying fraud, manipulation, cybercrime, and disinformation rather than opportunity. This is not an argument against AI access in emerging economies – it is an argument for making governance capacity, cyber resilience, trusted digital infrastructure, and cross-border cooperation part of the inclusion agenda alongside access itself.

This is not unique to AI. Strong collaborative governance is the foundation of competitiveness in any setting, particularly in times of rapid change.



1. The Global South and developing economies

Structural

Africa, South Asia, Latin America, and rural regions globally face a combination of barriers — insufficient digital infrastructure, weak institutional capacity, limited capital, and language exclusion from AI systems trained predominantly in English and Mandarin — that make participation in the AI economy structurally difficult rather than merely challenging. Respondents were clear that this is not primarily a technology access problem. It is a political economy problem: the investment, policy capacity, and institutional weight needed to participate on equal terms are absent in precisely the places where they are most needed.

"The ones already lacking meaningful connectivity and meaningful digital skills — most of them in Global South countries. The gap does not close by itself."



2. Regions with weak digital infrastructure

Geographic / cross-cutting

Limited connectivity, inadequate compute access, and low digital literacy exclude entire geographies from the AI economy — including rural areas within wealthy nations. This was one of the survey's underappreciated findings: the digital divide is not a developed-versus-developing world binary. Rural Ireland, remote Canada, and the North and Northeast regions of Brazil face structural exclusion that mirrors, in different proportions, the exclusion facing the Global South. Without infrastructure, participation is not difficult — it is impossible. One respondent raised a further dimension: countries whose tax revenue declines as intellectual activity offshores to AI company data centers face a fiscal erosion that compounds their exclusion.

"Knowledge workers and countries who don't own AI infrastructure. Countries whose tax revenue declines due to the offshoring of intellectual activities to AI company data centres."



3. Low-skilled and routine-task workers

Occupational

Workers in data entry, customer service, basic administration, logistics, and manufacturing face the most direct and immediate displacement risk, with the fewest pathways to transition. What makes this group particularly vulnerable is the combination of high automation exposure and low employer investment in their futures: companies facing

"Workers in routine jobs face higher risk, and populations in developing economies that lack institutions, capital, or infrastructure to compete in an AI-driven economy."

cost pressure tend to cut these roles rather than redesign them. Respondents noted a further dynamic: even where employers want to reskill these workers, the existing training infrastructure is inadequate to the scale of transition required.



4. Older workers and the apprenticeship model

Generational

Mid-career and older workers face a combination of retraining barriers, hiring bias, outdated credentials, and digital fluency gaps that make the transition exceptionally difficult. But this group carries a second risk that several respondents raised with particular force: if AI eliminates entry-level roles in law, finance, accounting, and other knowledge professions, the pathway by which young graduates build the expertise that eventually makes them irreplaceable disappears. The senior partner's judgement is valuable because it was built on years of doing the work AI now does instead. Remove that learning-by-doing foundation and the expert architecture above it becomes impossible to construct. This is not a future risk — it is happening now in several professional services sectors.

"The entry-level workforce faces risk as AI co-pilots make senior employees more productive, the need for junior associates to do initial drafting or data cleaning is vanishing — potentially breaking the traditional apprenticeship model."



5. Women, minorities, and marginalized groups

Structural / intersectional

Women are disproportionately represented in clerical, administrative, and customer service roles — the occupational categories with the highest automation exposure. Minorities, migrants, and linguistic minorities face compounding disadvantages: AI systems trained in majority-language, and majority-culture data perform poorly for them, while the institutions that might advocate for their interests are often under-resourced or absent. Several respondents noted the intersectional nature of this risk: the most excluded groups tend to face multiple simultaneous barriers rather than a single addressable problem, which makes conventional targeted interventions — designed to address one dimension of exclusion — systematically insufficient.

"Women in administrative roles, informal and gig workers, people with disabilities, linguistic minorities, and migrants — these groups face direct displacement with compounding disadvantages."

Beyond the Survey

Feedback sessions reframed the disruption of traditional apprenticeship pathways — identified in Chapter 5 as a significant risk — as an opportunity worth seizing. Organizations that deliberately pair the complementary strengths of different generations may find a genuine competitive edge in the transition.

Senior professionals bring judgment, tacit knowledge, context, and institutional memory that AI cannot replicate. Younger workers bring AI fluency, comfort with experimentation, and ease with new tools. These are not competing assets — they are complementary ones. Organizations that create structured opportunities for cross-generational collaboration and reverse mentoring can redesign work for the AI era while preserving the pathways through which people build genuine expertise.

The opportunity is real, but the mechanisms to act on it at scale are not yet in place. Developing, experimenting with, and scaling cross-generational learning programs — across sectors and geographies — is one of the more actionable design challenges of this transition.

Two Underappreciated Risks



Beyond the five groups above, two specific risks emerged from the open responses that are less frequently discussed in mainstream AI policy discourse but that several respondents – including some of the survey's most rigorous contributors – identified as significant.

01 The fiscal erosion risk

As AI systems perform knowledge work that was previously performed by human professionals in specific locations, the tax revenues that fund local governments, schools, and social services in those locations erode. This is not primarily about job losses – it is about the geographic concentration of AI value creation. A law firm in a major capital that uses AI to replace the analytical work previously done by junior associates located in smaller towns generates tax revenue in the jurisdictions where its AI provider is domiciled, not where its displaced workers live. Several respondents raised this as an underappreciated dimension of AI-driven inequality that current governance frameworks do not adequately address. As the Citrini Research^{III} scenario modelled in early 2026, the feedback loop is potentially self-reinforcing: AI displaces workers, reducing household income; reduced income contracts consumer demand; contracted demand weakens the broader economy; a weakening economy reduces tax revenues further, cutting the very public services – education, retraining, social protection – that displaced workers depend on to transition. This is not a prediction but a potential trajectory.

02 The narrative exclusion risk

AI systems are trained predominantly on English and Mandarin language data, and their outputs reflect the cultural assumptions, reference points, and knowledge structures. For communities whose languages and knowledge systems are underrepresented – including most of the Global South, indigenous communities, and smaller national languages – this means that the AI economy is not just economically exclusive. It is epistemically exclusive: the tools produce outputs that are less accurate, less useful, and less culturally appropriate for those communities. One respondent identified this directly: populations where English or Chinese are not the primary language face a compounding disadvantage that goes beyond access to the technology itself.

The Evidence That Change Is Possible

What the survey shows — and what the next chapter examines — is that respondents did not stop at diagnosis. They named examples: programs, platforms, and initiatives where AI is reaching people who would otherwise be excluded. Chapter 6 looks at what those examples have in common, and what it would take to replicate them at scale.

Beyond the Survey

Feedback sessions pointed to a dimension of the AI transition where better data and clearer frameworks are needed: AI has a physical footprint, and it is significant. Data centers, energy grids, water supplies, land, cooling infrastructure, semiconductors, and supply chains are the material foundation of the AI economy — and their presence in communities carries real costs and real benefits that are not yet well understood.

The public narrative around AI infrastructure has skewed negative, and in some cases the concerns are legitimate. Energy consumption, water use, and land requirements are substantial and growing. Communities hosting data centers do not always see commensurate local benefit. But the picture is not uniformly negative. AI infrastructure can bring investment, jobs, grid modernization, and long-term economic activity to regions that need it — and there are cases where communities can benefit meaningfully from hosting it.

What is missing is the rigorous impact modelling needed to tell these cases apart. As AI adoption grows, communities will increasingly ask who benefits from AI infrastructure, who bears the costs, and whether public engagement in those decisions is meaningful. Responsible AI governance must include honest answers to those questions — not just advocacy in either direction, but the data and analytical frameworks needed to understand the trade-offs clearly and make better decisions.



- 70% of respondents describe AI transformation in their context as limited or exclusionary – the strongest and most consistent verdict in the entire survey.
- AI is not creating new inequalities so much as amplifying existing ones. The policy response must therefore be structural – addressing the underlying conditions that AI accelerates – not only AI-specific.
- The correlation between inclusivity perception and confidence in co-evolution ($r = 0.38$) is the survey's strongest relationship: people who see an exclusionary AI do not just hold a different opinion – they see a different future.
- Five groups carry disproportionate risk: the Global South; regions with weak digital infrastructure; low-skilled and routine-task workers; older workers and new entrants facing a disruption in the apprenticeship pathway; and women, minorities, and marginalized groups facing intersectional disadvantage.
- Two underappreciated risks deserve more policy attention: fiscal erosion as AI concentrates value geographically, and narrative exclusion as AI systems trained on majority-language data underperform for minority-language communities.

Note

III. Citrini Research, The 2028 Global Intelligence Crisis, February 2026. citriniresearch.com/p/2028gic



What Good Looks Like: A Gallery of Human-AI Integration That Is Working

CHAPTER 6

The previous chapter documented the scale of the inclusion gap. This chapter documents evidence that it can be closed. The six cases presented here were identified through survey testimony from respondents and validated through research. They are not presented as finalized benchmarks – the contexts are too different for direct replication. They are presented as proof of concept: demonstrations that when the design choices matter and that the AI transformation can reach people who would otherwise be left behind.

Each case is examined through five dimensions: impact, replicability, scalability, resource intensity, and time to results. Impact captures the evidence of measurable human benefit – not just adoption rates, but whether people's capabilities, opportunities, or quality of life changed. Replicability asks whether the design choices that made it work are transferable to other contexts, or whether success depended on conditions that are unique to one place. Scalability examines whether the model can grow beyond its original setting without losing what made it effective. Resource intensity assesses what it costs in money, time, and institutional capacity to implement. And time to results matters because policymakers and institutions need to know whether they are making a short-term or long-term bet.

Together, the six cases span four sectors, six countries, and three income levels. What they share is more instructive than what distinguishes them. Below, we separate the cases into two different categories: (i) cases focused on preparing people to live and work in the AI age and (ii) cases in which AI is being used to improve people's lives.

Preparing People for the AI Age

The **four cases in this section** address the same fundamental challenge from different angles and at different scales: how do you ensure that the capabilities people need to participate in an AI-driven economy are not reserved for those who already have the most? Each case represents a different level of intervention — a state government in one of Brazil's least resourced regions legislating AI into the mandatory curriculum, a platform designed to reach rural students who would otherwise be excluded, a national co-governance model delivering AI education to 5.5 million children, and a decade of national infrastructure investment that made Singapore's AI skilling agenda possible. What they share is a deliberate design choice to start with the question of who gets left out — and to build the answer into the architecture of the program from the beginning.

List of Cases

Case 1: Brazil — Piauí AI Program

Case 2: Romania — Nextlab.tech

Case 3: Spain — Código Escuela 4.0

Case 4: Singapore — SkillsFuture and the AI Layer

Case 1: Brazil – Piauí AI Program

One of the first mandatory AI curricula in the Americas – UNESCO Prize 2025

Education

Context

Piauí is one of Brazil's least economically developed states, located in the Northeast – a region historically underserved by both federal investment and private sector activity. In 2023 the state government, in partnership with three federal universities – UFRGS, IFFar, and Unipampa – launched a mandatory AI curriculum across its entire public school network. The program now reaches over 120,000 students and more than 800 teachers. It is one of the first jurisdictions in the Americas to make AI a compulsory subject. In 2025 it was awarded the UNESCO ICT in Education Prize, recognizing its contribution to the responsible use of AI in education – one of four initiatives honored globally that year.

"The program integrates AI ethics, digital literacy, and hybrid learning modules into public schools. It reaches students who often have limited access to high-end private schools or digital tools – something most AI education initiatives completely miss."

What makes it work

Five design choices distinguish this program from generic digital education initiatives:

- **Mandatory, not optional:** making AI a compulsory subject – rather than an elective or enrichment program – ensures reach extends to the students least likely to self-select into technology courses. Piauí's decision to legislate AI into the curriculum rather than pilot it as a project is the single most important design choice.
- **Three-year structured curriculum:** a sequenced learning path covering Data Curation, Machine Learning, and Natural Language Processing gives the program intellectual depth and prevents the surface-level digital literacy that characterizes many national AI education initiatives.
- **Plugged and unplugged delivery:** activities are designed to work with and without connectivity, making the program viable in the low-resource, low-bandwidth environments that characterize much of rural Piauí. This is the mechanism that makes universal reach possible.
- **Sustained teacher investment:** over 680 educators trained annually – not a one-off launch event. The program treats teacher capacity as a continuous investment, not a prerequisite to be satisfied and forgotten.
- **University partnership for credibility and content:** collaboration with three Brazilian federal universities – the Federal University of Rio Grande do Sul (UFRGS), the Federal Institute Farroupilha (IFFar), and the Federal University of Pampa (Unipampa) – provides both the academic rigor and the institutional relationships that allow a state government with limited internal capacity to build something substantive.

Design insight: *Making AI education mandatory rather than optional is the design choice that determines reach. Everything else – curriculum depth, teacher training, low-resource delivery – follows from that prior commitment.*

Impact	120,000+ students and 800+ teachers across Piauí's entire public school network; UNESCO ICT in Education Prize 2025 – one of four initiatives recognized globally ^{IV}
Replicability	High – the mandatory-subject model, plugged/unplugged delivery design, and university partnership structure are directly replicable in other Brazilian states and across Latin America; UNESCO recognition provides a ready-made validation framework for adoption elsewhere
Scalability	Medium-high – mandatory status and state-level governance create the institutional architecture for scale; independent outcome evaluation not yet publicly available
Resource intensity	Medium – requires curriculum development, university partnerships, and sustained teacher training investment; plugged/unplugged design keeps hardware costs low relative to reach
Time to results	Medium-term – curriculum integration yields measurable AI literacy outcomes over academic years; UNESCO recognition confirms program is beyond pilot stage

Case 2: Romania – Nextlab.tech

AI-guided robotics education reaching rural and disadvantaged students at national scale

Education

Context

Nextlab.tech is a Bucharest-based platform that delivers AI-guided robotics and STEM education to students aged 8–16, with a deliberate focus on rural and economically disadvantaged communities. Founded in 2017 and operating across more than 6,000 schools and universities in five countries, the platform uses adaptive learning technology to personalize content to each student's level. Free robotics kits — sponsored by corporate partners including BCR, Romania's largest bank — are distributed to qualifying students, with priority given to those from disadvantaged backgrounds. The annual national competition is listed in Romania's Ministry of Education calendar and open to all students regardless of location or family income. What makes Nextlab distinctive is that it has solved the access problem not by lowering quality but by adapting delivery: the same knowledge base reaches a child in rural Moldova and one in Bucharest.

"We have children from disadvantaged communities who did not initially believe they would be able to build a working robot. There is little difference between a child in a remote rural area and one from a wealthy city — they share the same knowledge base, the same content, and the same narratives."

What makes it work

The design choices that make this model work include:

- Corporate co-funding of hardware: free kits distributed to disadvantaged students remove the cost barrier without requiring government procurement, while keeping the program's independence from political budget cycles.
- Ministry of Education accreditation: official inclusion in the national competition calendar gives the program institutional legitimacy and reach through existing school networks.
- Adaptive platform design: AI-personalized learning paths mean a child with no prior exposure to robotics can start where they are, not where the curriculum assumes they should be.
- International expansion model: the platform is now operating in five countries, demonstrating that the model travels beyond its country of origin.

Design insight: Access gaps in AI and STEM education can be closed without lowering standards – if the delivery model is designed to adapt to the learner rather than requiring the learner to adapt to the delivery.

Impact	Thousands of students in rural and disadvantaged communities building technical capabilities previously inaccessible to them; multiple countries now implementing the model
Replicability	High – the adaptive platform model and corporate kit-sponsorship structure are replicable in other middle-income countries with functional school systems
Scalability	High – already operating in 6,000+ institutions across five countries; Ministry accreditation model is scalable through national curriculum integration
Resource intensity	Low – low-cost semiconductor kits (entry-level hardware kit starts at around \$30 USD, making it among the most affordable robotics education hardware available globally) and an enterprise platform license model reduce per-student costs significantly
Time to results	Medium-term – skill development outcomes visible within an academic year; national scale achieved over several years

Case 3: Spain – Código Escuela 4.0

National program integrating computational thinking, AI, and robotics into all public school levels

Education / Government

Context

Launched under Spain's Digital Agenda 2026, Código Escuela 4.0 is a national program co-designed between the Ministry of Education and Spain's 17 autonomous communities. It integrates computational thinking, programming, AI concepts, and robotics into curricula from early childhood through secondary education in all publicly funded schools. The program targets 5.5 million students and includes both hardware provision – robotics kits, programmable boards, sensors – and sustained teacher training through regional implementation bodies. Unlike top-down mandates that fail at implementation, the program explicitly distributes execution responsibility to regional governments, which ensures regional governments have a genuine stake in the outcome.

"Spain's Código Escuela 4.0 stands out for its federal-regional co-governance model – the Ministry sets the framework and funds it, but the autonomous communities implement and adapt it. That design choice is what makes it credible at scale rather than a central mandate that local systems quietly ignore."

What makes it work

Three features make Código Escuela 4.0 a particularly instructive model for national-scale AI education reform:

- Federal-regional co-governance: the Ministry funds and sets standards; autonomous communities implement and adapt. This structure avoids both the inflexibility of centralized mandates and the fragmentation of purely local initiatives.
- Universal public school reach: by targeting all publicly funded schools from kindergarten through secondary, the program avoids the trap of AI education becoming a premium offering available only to well-resourced schools.
- Teacher training as a parallel track: hardware deployment is accompanied by mandatory teacher development, addressing the most common failure mode of technology-in-schools programs – equipment that arrives without the pedagogical support to use it.

Design insight: *National scale in education reform requires distributing both responsibility and authority across levels of government – not just funding from the top and delivery at the bottom.*

Impact	5.5 million students across all publicly funded schools in Spain; one of the largest national AI education deployments in the EU
Replicability	High within federal/devolved governance systems – the co-design model between national and regional governments is directly applicable to other countries with similar constitutional structures
Scalability	High – already operating at national scale; the model is inherently designed for system-wide deployment rather than pilot logic
Resource intensity	High – significant government investment required (€14M-€16M per region for hardware/training); EU Digital Education Action Plan funding provides a replication pathway for other EU member states
Time to results	Medium-term – curriculum integration yields measurable digital competency outcomes over academic years

Case 4: Singapore – SkillsFuture and the AI Layer

A decade of lifelong learning infrastructure, with AI programs built on top of it
Workforce / Government

Context

SkillsFuture^v was launched by the Singapore government in 2015 – a decade before AI disruption became the dominant workforce concern – as a national movement to embed lifelong learning into the culture and infrastructure of work. Every Singaporean adult receives a SkillsFuture Credit to spend on approved training courses. All autonomous universities, polytechnics, and institutes of technical education have made lifelong learning part of their institutional mission. By 2026, more than one in two eligible Singaporeans had used their SkillsFuture credit at least once since 2015, and more than 600,000 individuals took up training with SkillsFuture support in 2025 alone. The AI-specific layer was built on top of this existing infrastructure from 2023 onward. AI for Everyone and AI for Industry programs provide basic and advanced AI training across age groups and sectors. The SkillsFuture Level-Up Program, targeting mid-career workers, has reached more than 60,000 Singaporeans aged 40 and above since its 2024 launch. A target has been set to train 100,000 workers to be 'AI bilingual' – able to apply AI tools to real workplace problems – by 2029. In February 2026, Singapore's Budget announced that those enrolling in selected AI courses would receive six months of free access to premium AI tools, removing the cost barrier between training and application.

"Singapore's AI programs work because SkillsFuture was already there. You cannot separate the AI layer from the infrastructure underneath it. What Singapore built first – the credits, the pathways, the employer partnerships, the cultural expectation of continuous learning – is what makes the AI-specific programs land."

What makes it work

The Singapore case has one primary lesson and several secondary ones:

- Infrastructure before content: the SkillsFuture platform, credit system, employer networks, and cultural normalization of adult learning existed for nearly a decade before AI became the focus. The AI programs succeeded because the infrastructure was already trusted and used.
- Mid-career workers as a priority, not an afterthought: the Level-Up Program explicitly targets workers aged 40 and above – the group most at risk from the disruption of the apprenticeship dynamic – rather than focusing only on new entrants.

- Connecting training to tools: providing six months of free AI tool access alongside course enrolment closes the gap between learning and doing that defeats most reskilling programs.
- Non-PMET workers remain a challenge: Singapore's own data shows that non-professional, managerial, executive, and technical workers — those in lower-skilled roles — take up reskilling at lower rates despite higher need. The infrastructure exists; the uptake gap among the most vulnerable remains an open problem.

Design insight: *The question for other countries is not how to copy Singapore's AI programs. It is whether they have the underlying infrastructure — portable credits, sustained employer investment in workforce development, and a culture of lifelong learning — that would make such programs land.*

Impact	600,000+ trained with SkillsFuture support in 2025; 60,000+ mid-career workers through Level-Up since 2024; more than half of eligible Singaporeans have used their credit at least once
Replicability	Medium — the specific SkillsFuture credit and infrastructure model requires sustained political commitment and public investment over many years; elements are replicable, but the full system took a decade to build
Scalability	High — Singapore's model is already informing workforce policy in several Asian and European countries; the AI-bilingual target of 100,000 by 2029 demonstrates intentional scaling
Resource intensity	High — significant sustained public investment required; not viable without strong government commitment and employer co-funding

Time to results

Long-term – cultural normalization of lifelong learning is a generational project; AI-specific outcomes measurable within 3–5 years

Using AI to Improve Lives

The **two cases in this section** operate on a different logic from the previous ones. They are not primarily about preparing people for the AI economy — they are about using AI to extend the reach of essential services to populations that existing systems, at human-delivered scale and cost, could never serve. In Kenya, AI layered onto a trusted mobile payment platform reaches 30 million users with financial and agricultural services that were previously economically unviable at their transaction size. In Alberta, AI-assisted diagnostics and community co-designed safety tools reach rural and Indigenous communities that face significant barriers to specialist services. What these cases share is a design principle: AI works best not when it replaces existing human systems, but when it extends what those systems cannot reach.

List of Cases

Case 5: Kenya — M-Pesa and the AI-Enabled Financial Ecosystem

Case 6: Canada (Alberta) — Amii

Case 5: Kenya – M-Pesa and the AI-Enabled Financial Ecosystem

AI layered onto trusted mobile payment infrastructure to reach the previously unbanked

Financial inclusion / Private sector

Context

M-Pesa began in 2007 as a mobile money transfer service and has grown into the financial backbone of Kenya, handling more than 61 million daily transactions. What makes the M-Pesa case relevant to this report is not its origin, but what has been built on top of it. Safaricom has systematically layered AI capabilities across the platform: the Zuri generative AI chatbot handles complex queries for over 30 million M-Pesa users; AI-driven credit scoring and fraud detection enable micro-loans for populations with no formal credit history; and FarmerAI – launched in 2025 – gives smallholder farmers access to verified agricultural advisory information in real time. These services reach populations that formal financial and agricultural advisory services could never serve economically at human-delivered scale. The AI layer does not replace the human ecosystem of M-Pesa agents, local knowledge, and community trust – it extends what that ecosystem can offer.

"The most underappreciated inclusion gains may come not from new AI programs but from layering AI onto existing trusted infrastructure that already reaches excluded populations. M-Pesa is the clearest global example of what that looks like at scale."

What makes it work

The M-Pesa model demonstrates three principles with particular clarity:

- Trusted infrastructure as the foundation: the AI services work because M-Pesa is already trusted, already embedded in daily life, and already reaching people that formal institutions cannot. AI without that foundation would face the adoption barriers that defeat most financial inclusion initiatives.
- Unit economics reversal: micro-loan assessment, fraud detection, and agricultural advisory services were not economically viable at human-delivered scale for populations transacting at M-Pesa ticket sizes. AI reduces the marginal cost of each interaction toward zero, making inclusion economically rational rather than philanthropically dependent.
- Language and cultural context: FarmerAI delivers advisory content in local languages to smallholder farmers – directly addressing the narrative exclusion risk identified in Chapter 5. The service is designed for its users, not adapted for them as an afterthought.

Impact	30+ million M-Pesa users with access to AI-powered services; smallholder farmers across Kenya accessing verified agricultural advisory information for the first time
Replicability	Medium – the M-Pesa model itself is specific to contexts with a dominant mobile money platform; the principle of layering AI onto existing trusted infrastructure is broadly replicable
Scalability	High – already operating at national scale in Kenya; model is expanding across East Africa; the FarmerAI approach is being studied by agricultural development organizations globally
Resource intensity	Low – AI services built on existing infrastructure avoid the cost of building new distribution channels; commercially self-sustaining rather than grant-dependent
Time to results	Immediate – AI services deployed on existing platform reach existing users from day one; no adoption infrastructure required

Case 6: Canada (Alberta) – Amii

AI-powered diagnostics and community-designed solutions for rural and Indigenous communities

Healthcare / Research institute

Context

The Alberta Machine Intelligence Institute (Amii) is one of Canada's three named institutes in the Pan-Canadian AI Strategy, based at the University of Alberta in Edmonton. Its rural healthcare work demonstrates two distinct models of human-AI integration. The first involves portable AI-assisted ultrasound technology: machine learning models trained on thousands of images guide non-specialist operators to position the scanner correctly and flag anomalies for review, extending diagnostic capability to communities that currently face significant travel distances for appointments. The second is a community co-designed online hate detection tool built in partnership with Indigenous communities – using natural language processing trained on culturally sensitive language to identify both overt and covert racist speech in a way that generic moderation tools miss entirely. Both cases embody a principle that runs through Amii's work as an institution: AI for good and for all means designing with communities, not for them.

"Amii's work demonstrates what human-centered AI looks like in practice – not AI replacing clinicians in remote communities, but AI enabling a clinician with limited training to do what previously required a specialist. The human is still the point. The AI extends their reach."

What makes it work

Two features of the Amii cases are particularly relevant for this report:

- Community co-design as a non-negotiable: the Indigenous online safety tool was built with Indigenous experts leading the cultural and linguistic training, not as a retrofit of an existing tool. The result is a system that can detect coded language and veiled meanings that generic moderation tools miss – because only the community the tool is designed to protect could identify those patterns.
- Human-in-the-loop as the design principle: the portable ultrasound AI guides and augments the operator; it does not replace the clinician's judgement. The community hate detection tool flags and supports moderation; it does not automate the decision. In both cases, AI extends human capacity without removing human accountability.

Design insight: *AI designed with excluded communities – rather than for them – produces outcomes that purpose-built tools cannot achieve, because the knowledge required to build them only exists inside those communities.*

Impact	Rural and Indigenous communities in Alberta accessing diagnostic and safety services not previously available at their location; model being studied as a framework for AI deployment in remote and underserved communities across Canada
Replicability	Medium – the community co-design model is highly replicable in principle but requires sustained relationship-building that cannot be shortcut; the portable ultrasound approach is already being adapted in other low-resource contexts globally
Scalability	Medium – the portable diagnostics model scales with hardware access; the community co-design model scales more slowly because genuine co-design cannot be industrialised
Resource intensity	High – requires sustained research institute investment, community relationship-building, and clinical validation; not a low-cost model
Time to results	Medium-term – diagnostic tools deployable within months of community partnership; long-term outcomes require sustained evaluation

What the Six Cases Have in Common

Six cases across six countries and four sectors produce a set of design lessons synthesized by table below:

Case	Sector	Impact	Replicability	Scalability	Resource intensity	Time to results
Brazil – Piauí	Education	High	High	Medium	High	Medium-term
Romania – Nextlab	Education	Medium	High	High	Low	Medium-term
Spain – Escuela	Education	High	High	High	High	Medium-term
Singapore – SF	Workforce Training	Very High	High	High	High	Long-term
Kenya – M-Pesa	Financial	Very High	Medium	High	Low	Immediate
Canada – Amii	Healthcare	Medium	Medium	Medium	High	Medium-term

Three patterns can be identified regardless of sector, geography, or income level.

- **Deliberate targeting from the outset.** Every case in this chapter required an explicit decision to serve excluded groups – rural students in Piauí and Romania, non-English speakers in Kenya, Indigenous communities in Alberta, lower-wage workers in Singapore.
- **Human oversight as the design principle, not the safeguard.** In every case, AI augments human capacity rather than replacing it. The farmer using FarmerAI still makes the planting decision. The operator using the portable ultrasound tool still refers the patient. The student building a Nextlab robot is still doing the building.
- **Infrastructure over programs.** The cases are built on permanent institutional architecture – SkillsFuture's credit system, M-Pesa's trusted payment network, Spain's federal-regional co-governance structure, Romania's Ministry-accredited competition calendar.

Inclusive AI transformation does not happen by accident. It happens by design.

The cases in this chapter are not finished stories. Singapore's non-PMET workers still take up reskilling at lower rates than their more educated peers. Romania's model has not yet been independently evaluated for learning outcomes at national scale. Kenya's FarmerAI is months old. Alberta's community co-design work is painstaking and cannot be easily replicated at speed. These are signals of what becomes possible with sustained commitment – not evidence that the problem is solved.

Chapter 6

Key Takeaways



- Inclusive AI outcomes are achievable across geographies, sectors, and income levels – the six cases span three continents, four sectors, and very different resource environments.
- All six cases required deliberate design choices to reach excluded groups. None happened because AI naturally distributed its benefits broadly.
- Human oversight is the consistent design principle – AI augments and extends human capacity in every case; it does not replace human judgement or agency.
- Infrastructure outlasts programs. The cases with the greatest scale and durability are built on permanent institutional architecture, not project-funded initiatives.
- The Singapore case carries a specific lesson for other countries: the question is not how to copy the AI programs but whether the underlying infrastructure – trusted, used, funded – exists to make them land. For most countries, that infrastructure needs to be built first.

- The unit economics of exclusion can be reversed. AI reduces the marginal cost of reaching excluded populations to a point where inclusion becomes commercially rational – not just socially desirable.
- None of these cases are finished. Each faces its own open challenges – uptake gaps, evaluation gaps, scale gaps. They are proof of concept, not proof of completion.

Notes

IV. The 2025 UNESCO King Hamad Bin Isa Al-Khalifa Prize for the Use of ICT in Education recognized four initiatives globally on the theme of ethical and responsible AI in education: the Piauí AI Programme (Brazil), Experience AI by the Raspberry Pi Foundation and Google DeepMind (UK), AI4InclusiveEducation – La Scientothèque (Belgium), and Mahara-Tech Information Technology Institute (Egypt). unesco.org/en/prizes/ict-education/laureates

V. In 2021, the GFCC launched the report “Future Skills” highlighting trends and national presenting skills innovative initiatives focused on workforce development. SkillsFuture was one example of a government-funded program aimed at providing solutions to professionals, workers, employers, and the Singaporean society at large: https://www.thegfcc.org/_files/ugd/f344ed_61eef8e479f64175acb8d38510e12792.pdf



The Work Begins Now: Turning Insight into Action

CHAPTER 7

This report captures a moment in time and opens a conversation about the future of humanity with AI. The survey behind it reached leaders across 32 countries and five sectors. It points towards a global perception that AI is transitional rather than transformational and that the systems around it are failing to keep pace. In addition, it leads to the conclusion that the distribution of its benefits is already unequal and becoming more so. It also produced a set of directions for action that the respondents identified as meaningful.

This chapter distills what the survey said about what should happen next — which initiatives, which priorities, which measures of success — and frames them as raw material for a common action agenda. But before turning to priorities and action, it addresses a question the survey answered with unusual clarity: who is actually responsible for making the transition work, and what does each actor need to do differently? Eighty-six per cent of respondents said they want to stay involved. The question we try to answer here is what they are being invited into.

Who Is Responsible for What

Cross-sector collaboration was the dominant answer when respondents were asked how governments, businesses, and universities should share responsibility for preparing people for the AI economy. Sixty per cent of respondents called for joint action across all three institutions. But the survey went further than that general call. Respondents were specific about the distinct role each actor should play, and notably honest about the structural obstacles to each playing it well.

The five themes that emerged, in order of frequency:

1

Cross-sector collaboration

(60% of respondents)

No single institution can carry this alone; joint councils, co-funded programs, and aligned strategies are necessary.

2

AI literacy as a shared public good

(46% of respondents)

AI fluency should be treated like numeracy: a universal baseline funded publicly and reinforced by employers and universities alike.

3**Business: workforce training and upskilling****(40% of respondents)**

Companies must invest in their own people rather than waiting for the education system to produce ready-made talent.

4**Universities: curriculum reform and research****(34% of respondents)**

Redesign curricula across all disciplines, produce independent research on AI's social impacts, and act as neutral conveners between government and industry

5**Government: safety nets, regulation, and funding****(28% of respondents)**

Set policy, fund digital infrastructure and reskilling, protect displaced workers, and ensure AI benefits are distributed equitably – including responsibility for foundational education (K-12).

The consensus model is intuitive and appeared consistently across sectors and geographies: government sets the floor and the policy framework; business invests in its own people rather than waiting for the education system to catch up; education institutions at all levels reform curricula with universities producing independent research on AI's social impacts. The survey's most candid voices noted that calling for collaboration without addressing the structural conditions that make it difficult is wishful thinking. Each actor operates under different pressures, different timelines, and different accountability systems. Genuine collaboration requires more than shared intent – it requires designed incentive structures and governance arrangements that make joint action rational and sustainable for each party.





What Needs to Happen: Priorities Identified

When asked which areas should be prioritized in a Human and AI Global Action Plan, respondents could select up to five from eight options. The results show a clear hierarchy – and some meaningful divergences across sectors.

What Needs to Happen: Priorities Identified

When asked which areas should be prioritized in a Human and AI Global Action Plan, respondents could select up to five from eight options. The results show a clear hierarchy – and some meaningful divergences across sectors.

1	AI-ready education and skills frameworks Top choice for Government (75%), Academia (79%), and Industry (66%)	69%	
2	Innovation and entrepreneurship for inclusive growth Strongest for Non-profit (67%); reflects the inclusion-through-opportunity argument	55%	
3	Reskilling and lifelong learning ecosystems Non-profit sector leads at 67%; consistent with civil society's focus on workers most at risk	54%	
4	Human-centered AI ethics and governance Highest for Government (62%) and Industry (57%); reflects urgency around regulatory gaps	53%	

5	AI for global development and sustainability	41%	
	<i>Notably low for Non-profit (25%), suggesting they prioritize workforce transitions over global frameworks</i>		
6	Human-AI productivity partnerships in industry	44%	
	<i>Highest for Government (62%); relatively lower for civil society</i>		
7	Public trust, transparency, and communication	35%	
	<i>Consistently lower across all sectors – not dismissed, but not a top-tier priority</i>		
8	Social safety nets and workforce transition support	32%	
	<i>Lowest overall – a significant gap given the scale of displacement risk identified in Chapter 5</i>		

Three findings in this ranking deserve particular attention. First, AI-ready education and skills frameworks dominated by a wide margin – the only option chosen by more than two-thirds of respondents, and the top choice in every sector. This is consistent with the survey's most persistent signal: the human side of the transition, not the technology side, is where the agenda must start.

Second, the most counterintuitive finding is about social safety nets. At 32% overall, they ranked last — a significant gap given that Chapter 5 identified low-skilled and routine-task workers as facing the most immediate displacement risk. But the sector breakdown reveals something more specific: non-profit and civil society respondents — those closest to vulnerable populations — selected safety nets at just 8%, the lowest of any sector. Government respondents, by contrast, selected it at 50%. One interpretation is that civil society organizations have concluded that safety nets are a government responsibility and are focusing their attention on what they can directly influence: reskilling and innovation. Another is that the urgency of the capability gap has crowded out the protection agenda even among those who know displacement best. Either way, the gap between identified risk and stated priority is large enough to warrant explicit attention.

Third, the non-profit sector stands out across the table. It leads on reskilling and lifelong learning ecosystems (67%) and innovation and entrepreneurship for inclusive growth (67%), but scores lowest on human-centered AI ethics and governance (25%) and AI for global development and sustainability (25%). Civil society in this survey is focused on near-term, people-centered interventions — not on frameworks and governance architecture. That is a useful signal for how to sequence the action agenda: capability and opportunity first, governance architecture in parallel but not as a prerequisite.

Five Key Initiatives the Survey Calls For

Five initiative directions emerged with clarity and consistency. They are presented here as the directions respondents identified as most meaningful, with enough specificity to begin design conversations.



1: A Global Human–AI Skills Alliance

What it involves

Develop shared AI competency frameworks, stackable credentials, and reskilling standards that travel across borders and sectors. The alliance would not produce another skills taxonomy – it would create the institutional architecture for recognizing and building AI-era capabilities at scale, including apprenticeship pathways into AI-augmented roles for workers displaced from routine positions.

"The ones already lacking meaningful connectivity and meaningful digital skills – most of them in Global South countries. The gap does not close by itself."



2: A Network of Human–AI Innovation Hubs

What it involves

Applied hubs where countries pilot AI-assisted workflows in healthcare, education, manufacturing, public services, and SME contexts – with a shared learning infrastructure that turns local experiments into global playbooks. The hubs would generate ethical guidelines rooted in practice rather than theory, and cost-effective adoption models designed for lower-resourced settings. Critically, hubs in the Global South would not be satellites of northern models: they would generate original knowledge about what works in lower-resource, lower-institution contexts.

"A network of applied hubs where countries pilot AI-assisted workflows with shared learning. Outputs: proven implementation playbooks, ethical guidelines rooted in practice, cost-effective adoption models for lower-resourced regions."



3: A Global Human-AI Competitiveness Index

What it involves

A shared index tracking how countries are integrating AI to enhance — not replace — human competitiveness. The index would measure augmentation rates, displacement rates, skills gaps, reskilling access, and equity of outcomes — not just AI adoption rates or productivity aggregates. It would fill the data void that respondents consistently identified: there is currently no authoritative global picture of what AI is actually doing to work and prosperity at population scale. The index would make that picture visible and make it politically difficult to claim success while distribution gaps are widening.

"Develop a global index to measure how effectively countries are integrating AI to enhance human competitiveness — tracking augmentation, displacement, skills gaps, and readiness."



4: An AI Literacy Initiative for the Global South

What it involves

Targeted programs ensuring Global South participation — AI literacy in local languages, digital infrastructure investment advocacy, open tools accessible to low-resource settings, and policy templates adaptable to weaker institutional environments. The Piauí model — mandatory curriculum, plugged and unplugged delivery, university

"AI literacy programs focused on the Global South — not just technical training, but teaching people when AI works, when it doesn't, and how to think with it. The literacy gap is what determines who benefits from AI and who gets displaced by it."

partnership for quality assurance — provides a ready-made blueprint. The initiative would also address the narrative exclusion risk: supporting the development of AI tools and training datasets in underrepresented languages, not only adapting English-language tools for non-English speakers.



5: A Governance and Accountability Standards Program

What it involves

International frameworks, governance templates, and accountability standards ensuring AI is transparent, human-centered, and equitable — with consequences, not just voluntary commitments. Respondents were explicit that the world has enough ethics statements and not enough enforcement. The program would develop model regulatory frameworks adaptable to different institutional contexts, worker voice mechanisms for AI deployment decisions, and audit standards for bias monitoring and human-in-the-loop requirements. It would also advocate for the one structural reform most consistently identified as missing: making governance frameworks binding rather than aspirational.

"The biggest opportunities are pushing for standards that actually have consequences rather than the usual voluntary ethics theatre, funding real experiments to figure out what transition support works at scale."

Across the survey's open responses, a consistent picture emerged not just about the key initiatives to pursue to advance the human edge in the age of AI, but also about the potential roles the GFCC can play in that context. Three potential roles came through clearly: a neutral convener across sectors and geographies, able to build bridges that no single actor can build alone; a trusted curator of what works, connecting doers with needers through validated implementation playbooks rather than more reports; and a champion of governance with real consequences, willing to advocate for enforcement rather than produce more frameworks.

How to Measure Whether This Is Working

How can we best measure success in ensuring that AI drives human prosperity rather than displacing it? This question produced the most analytically rich responses in the survey. Respondents pointed towards measurements that go well beyond conventional productivity metrics. The table below synthesizes the most frequently named indicators, organized into five categories that together constitute a scorecard for human-AI co-evolution.

Category	What respondents said should be measured
Economic outcomes	• Net job creation vs displacement – tracked at sector and skill level, not only in aggregate • Wage growth at median and lower-quartile levels – not averages that can mask concentration at the top • Productivity gains per worker – with evidence of how those gains are distributed between capital and labor • Emergence of new job categories in AI-augmented roles
Skills and capability growth	• AI literacy rates across income levels and geographies – not only in high-skill populations • Participation in reskilling programs among underrepresented groups • Speed and cost of workforce transitions from AI-exposed to AI-augmented roles • Share of workers successfully moving into higher-skill or higher-wage roles with AI support

Category	What respondents said should be measured
Workforce quality and wellbeing	• Worker autonomy and decision-making authority in AI-enabled environments • Reduction in administrative burden and low-value cognitive work • Job mobility – whether workers can move between roles and sectors • Worker-reported meaningfulness of work – not only employment rates
Inclusive access and equity	• Reduction in the digital divide – AI tool access across income levels and geographies • Distribution of productivity gains – Gini coefficient trends and SME adoption rates • Representation of underrepresented languages in AI training data and outputs • Poverty reduction attributable to AI-enabled services in previously underserved communities
Governance and accountability	• Percentage of AI deployments with documented human-in-the-loop requirements • Bias monitoring coverage – proportion of AI systems subject to regular audit • Worker voice inclusion – proportion of AI deployment decisions with formal worker consultation • Responsible use outcomes – not just adoption rates or capability benchmarks

Translating what respondents identified into a rigorous measurement framework is the next challenge – and it is not a simple one. Some of the dimensions people flagged, such as productivity growth and wage trends, have well-established metrics. Others, such as worker autonomy, meaningful work, and the lived experience of the transition, are harder to operationalize and not yet consistently tracked at global scale. The gap between what matters and what gets measured is itself part of the problem this report is trying to name.

At its core, the measurement question is a human one. Prosperity is not a GDP figure – it is whether people are thriving. If AI drives economic growth while workers feel more anxious, less secure, and further from a good life, then by any

meaningful definition the transition is producing growth without widespread prosperity – regardless of what the aggregate numbers show.

A more complete framework might combine three layers: the human and distributional dimensions respondents identified as most important; traditional economic indicators – productivity growth, wage trends, job creation – which remain relevant even if they don't tell the whole story; and human experience metrics that capture whether people are actually thriving.

Gallup, a GFCC member and one of the world's leading authorities on measuring human experience at scale, offers tools directly relevant to this third layer, more specifically in their Gallup Economics framework.^{VI} Gross Domestic Thriving – which measures whether people believe their lives are going well today and improving tomorrow – is perhaps the most important single signal of whether prosperity is real or merely statistical. But Gross Domestic Payroll to Population, which cuts through misleading unemployment figures to measure genuine full-time employment, and Gross Domestic Ambition, which tracks whether workers are genuinely engaged rather than simply present, are also worth serious consideration.

Gallup's framework and other existing tools can be potentially applied, adapted, or extended to track the human dimensions of AI-driven prosperity. For that to be possible, one methodological challenge will require particular attention: isolating the effect of AI from other forces shaping human prosperity metrics over time. If Gross Domestic Thriving declines, or if fewer people hold genuine full-time employment, how much of that change is attributable to AI displacement versus economic cycles, demographic shifts, or other structural forces? Answering that question rigorously is difficult – but asking it honestly is the starting point for measurement that is actually useful rather than merely reassuring.

But the most searching challenge to any measurement framework came from a survey respondent, who framed the stakes in historical terms:

Think about industrialization: GDP per capita climbed dramatically while huge numbers of people's actual lives got worse for decades before they got better. So, by the standard scorecard, success. By experience, a lost generation. Which is why the most rigorous answer is probably: measure a lot of things at once, including things that are hard to measure.

The challenge to measure the things that are hard to measure — instead of the comfort of metrics that look good in aggregate while masking who is paying the cost — is perhaps one of the most important commitments to ensure a better transition.

A small number of respondents went further still — questioning not the measurement framework but the frame itself. One respondent put it directly: the survey's questions assume that AI will follow the pattern of previous technological transitions. It will not. This is civilization-changing technology. The question may not be how to prepare people for an AI economy, but how to restructure society and distribute wealth in a world where work, as we know it, may no longer exist.

These voices are not the majority — they are rather outliers — but they are worth listening to. Not as a reason to abandon the agenda of action, but as a reminder that the goal is not AI adoption. The goal is human flourishing. AI is one of the means. Whether it becomes a good one depends on choices that have not yet been made — and on whether the frameworks we build are honest enough to tell us when they haven't been.

The Invitation

Eighty-six per cent of respondents said they want to receive updates and potentially get involved. That number is unusual for a survey of this kind. It is not just a courtesy response — it reflects an appetite for continued engagement from a leadership cohort that has thought carefully about these questions and wants to be part of what comes next.

The analysis is done. The directions are clear. What the GFCC does with them — which initiatives it commits to, which partnerships it builds, which governance positions it takes — is the question this report leaves open on purpose.

Beyond the Survey

Feedback sessions brought data quality to the center of the competitiveness conversation — not as a technical detail, but as a strategic issue that will determine who captures more value from AI.

AI systems are only as useful as the data that feeds them. Poor, incomplete, biased, outdated, or badly structured data can produce confident but unreliable outputs — and the consequences compound as organizations make decisions on that basis. Conversely, high-quality, trusted, proprietary, domain-specific, and local-language datasets may become core sources of economic growth and competitive advantage for the organizations and countries that invest in building them. Future human-AI competitiveness will depend not only on access to AI tools, but on the capacity to build, govern, verify, and use reliable data.

This has a direct implication for the measurement framework above. Tracking whether AI is driving human prosperity requires data that is itself trustworthy, representative, and comparable across geographies and over time. Data quality, provenance, interoperability, and ownership are not technical details — they are foundational to whether the indicators we care about can be measured honestly at all.



- AI-ready education and skills frameworks is the clear top priority – chosen by 69% of respondents and leading in every sector. The human side of the transition, not the technology side, is what respondents want the GFCC to lead on.
- Social safety nets ranked last despite being among the most pressing needs identified in Chapter 5 – a gap that likely reflects the seniority bias of the respondent pool and deserves explicit attention in the Action Plan.
- Five initiative directions emerged from the survey: a Global Skills Alliance, a Human-AI Innovation Hub network, a Global Competitiveness Index, a Global South AI Literacy Initiative, and a Governance and Accountability Standards Program.
- The measurement framework respondents described goes far beyond GDP or job counts – it encompasses workforce wellbeing, equity of outcomes, governance quality, and the distribution of productivity gains, not just their aggregate size.
- The GFCC is being asked to be three things: a neutral cross-sector convener, a trusted curator of what works, and a champion of governance with real consequences rather than voluntary commitments.
- 86% of respondents want to stay involved. That appetite for continued engagement is itself a resource – and an obligation to follow through.
- The most important next step is not the Action Plan. It is reaching the people this survey did not reach – the workers, communities, and populations most affected by the transition – and making their voices part of what the GFCC does next.

Note

VI. Gallup, Gallup Economics: Leading Indicators of Economic Growth, Annual Report 2026, published in partnership with Semafor World Economy, Washington D.C., April 2026.



**I – Survey
Methodology
Overview**

II – GFCC Members

III – GFCC Fellows

ANNEXES

Annex I – Survey Methodology Overview



The survey "Human and AI: Co-evolving to Drive Inclusive Prosperity" was designed and administered by GFCC as part of its Human and AI Global Task Force initiative. It ran from December 2, 2025 to April 4, 2026. **All responses were treated confidentially and synthesized anonymously.**

Data Collection and Respondent Profile

Responses were collected through four channels:

57% (64) The GFCC website	20% (23) Email blast	20% (23) Members and Fellows link	3% (3) Social media
-------------------------------------	--------------------------------	---	-------------------------------

The website link was distributed primarily through personal and professional networks.

The respondent pool skewed senior – 77% hold C-level, senior leadership, or senior management roles. By sector: Industry/Business (42%), Academia/Research (26%), Other (15%), Non-profit/Civil Society (11%), and Government/Policy (7%). Expertise concentrated in technology and AI, innovation and entrepreneurship, and economic development, with limited representation from workforce and labor policy specialists.

Geographically, responses came from 32 countries across 6 continents.

Survey Instrument and Analysis

The survey comprised 21 questions across five sections, combining single-select, multiple-select, 5-point Likert scale, and open-text formats. All closed questions achieved 100% response rates; open-text response rates ranged from 46% to 100%.

Quantitative analysis used descriptive statistics and cross-tabulations across key variables including confidence in AI co-evolution, AI relationship view, system

preparedness, and AI inclusivity perception. The strongest relationship identified was between confidence in AI co-evolution and AI inclusivity perception (Pearson $r = 0.38$), suggesting that how respondents perceive the distribution of AI's benefits shapes their overall outlook on the transition. All quantitative analysis was conducted in Python using pandas.

Qualitative analysis involved full reading and inductive thematic coding of all open-text responses across eight questions. Themes were derived from the responses rather than applied from a pre-existing framework. Direct quotations used in the report are verbatim from survey responses.

The four personas in Chapter 3 – Optimist, Skeptic, Systems Thinker, and Inclusion Advocate – were derived inductively from qualitative analysis of open-text responses, then validated quantitatively. Four recurring patterns of framing emerged consistently across geographies and sectors, distinguished by which risks respondents weighed most heavily, how they understood the role of institutions, and what kind of action they believed would make the decisive difference. Each persona was written as a composite of real response language.

To estimate the prevalence of each persona in the full respondent pool, a scoring model was applied to closed-question responses on AI relationship view, confidence in co-evolution, main barriers, critical human capabilities, and AI inclusivity perception. Results were: The Optimist (40%), The Systems Thinker (23%), The Inclusion Advocate (19%), and The Skeptic (18%). Of 113 respondents, 48% scored clearly for one persona; 52% sat at boundaries between two, reflecting genuine overlap in worldviews. These proportions should be treated as indicative rather than precise – the personas are analytical constructs, not discrete population segments.

Annex II – GFCC Members

Albania

Albanian Council on Competitiveness and Innovation

Australia

Adelaide University
Monash University
Western Sydney University

Brazil

Brazilian National Confederation of Industries (CNI)

Bulgaria

Glavbolgarstroy Holding AD

Canada

Sandra Duxbury

Greece

Council on Competitiveness of Greece—CompeteGR
Delphi Economic Forum
Piraeus Bank

Hungary

Bernadett Petri

Japan

Japan Science and Technology Agency
Tsukuba Global Innovation Promotion Agency (TGI)

Kazakhstan

Center for Research & Consulting LLC
Kazakhstan Growth Forum

Lebanon

American University of Beirut

Malaysia

Malaysian Industry-Government Group for High Technology (MIGHT)
Universiti Teknologi Petronas

New Zealand

Auckland University of Technology
University of Auckland

Oman

National Competitiveness Office

Portugal

Universidade Católica Portuguesa

Romania

Bucharest University of Economic Studies

Saudi Arabia

Ministry of Investment for Saudi Arabia (MISA)

UK

Centre for Competitiveness
Aston University
Harper Adams University
Queen's University Belfast

Ukraine

Kyiv National Economic University named after Vadym Hetman

USA

Academy of Art University
Arizona State University
Council on Competitiveness
Gallup, Inc.
Georgetown University
Lockheed Martin Corporation
University of Arizona
University of California San Diego
University of Illinois
University of North Carolina at Chapel Hill
University of Pittsburgh
University of South Carolina
University of Vermont
Vanderbilt University

Zimbabwe

National Competitiveness Commission

Annex III – GFCC Fellows

Distinguished Fellows

Amr Al-Dabbagh
Prince Saud bin Khalid Al-Faisal
Ahmad Tajuddin Ali
Hiroshi Amano
John Anderson
Spyros Artavanis-Tsakonas
C. Derek Campbell
Nathalie Cely
Athina Chatzipetrou
Lee Yee Cheong
Luciano Coutinho
Petros G. Doukas
Alexander Downer
Luis Fernandes
Hippolyte Fofack
Claudio Vilar Furtado
Nika Gilauri
Yanos Gramatidis
Michinari Hamaguchi
Zakri Abdul Hamid
Karen Holbrook
Jerry Hultin
Greg Hunt
Carole Hyatt
João A. H. da Jornada
Charles Kiefel

Stephen Kingon
Ashwani Kumar
Tae-Shin Kwon
Joan MacNaughton
Christos Megalou
Peter Meyers
Jan Mládek
Liam Nellis
Hiro Nishiguchi
Lucas Papademos
Emmanuel Pohl
Frank-Jürgen Richter
Laura Sandys
Lori Schmidt
Jack Sim
Arthur Sinodinos
Rogerio Studart
Yusoff Sulaiman
Manuel Trajtenberg
Gregory F. Treverton
Pete Worden
Kandeh Yumkella

Senior Fellows

Edward D. Agne
Chris Allen
Rehan Chaudhri
Fred T. Davies
Margareta Drzeniek
Mark Esposito
Marcelo Augusto de Felippes
Dionisio Garcia
Banning Garrett
Deborah Gordon
Kwanza Hall
Greg Horowitz
Elisa Jagerson
Nkem Khumbah
Susan McCalmont
James Metson
Mark Minevich
Michael Nelson
Peter Smith
Ted Zoller

JOIN OUR GLOBAL COMMUNITY

Contact us to set up a call:

 info@thegfcc.org

 www.thegfcc.org