Connecting the Dots

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

EDITION 3 · SEPTEMBER 2026



IN THIS EDITION

Australia's AI strategy

How Australia's AI governance compares internationally

Is AI making us stupid?

How a generation is losing the ability to think independently

Dr Sukh Rashpal Sandhu | Subscribe free on LinkedIn

EDITION 3 | SEPTEMBER 2026

Contents

Connecting the Dots: September 2026 3

Australia's AI Strategy: A Global Comparison How does Australia’s AI governance stack up internationally? 4

Analysis: Is AI making us stupid? How a generation is losing the ability to think independently 38

Follow us on LinkedIn 52

About Connecting the Dots 54

INTRODUCTION

Connecting the Dots: September 2026

By Dr Sukh Rashpal Sandhu

Edition 3 of Connecting the Dots: 2 pieces for September 2026.

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

This edition brings together 2 pieces first published in Connecting the Dots on LinkedIn, each carrying the date it first appeared. Republished pieces appear as they were first published, with their original images; facts, figures and links are as they stood on that date.

In this edition

- Australia's AI Strategy: A Global Comparison How does Australia’s AI governance stack up internationally?
- Analysis: Is AI making us stupid? How a generation is losing the ability to think independently

Thank you to every reader who subscribes, shares and responds.

Warm regards,
Dr Sukh Rashpal Sandhu
Author, Connecting the Dots

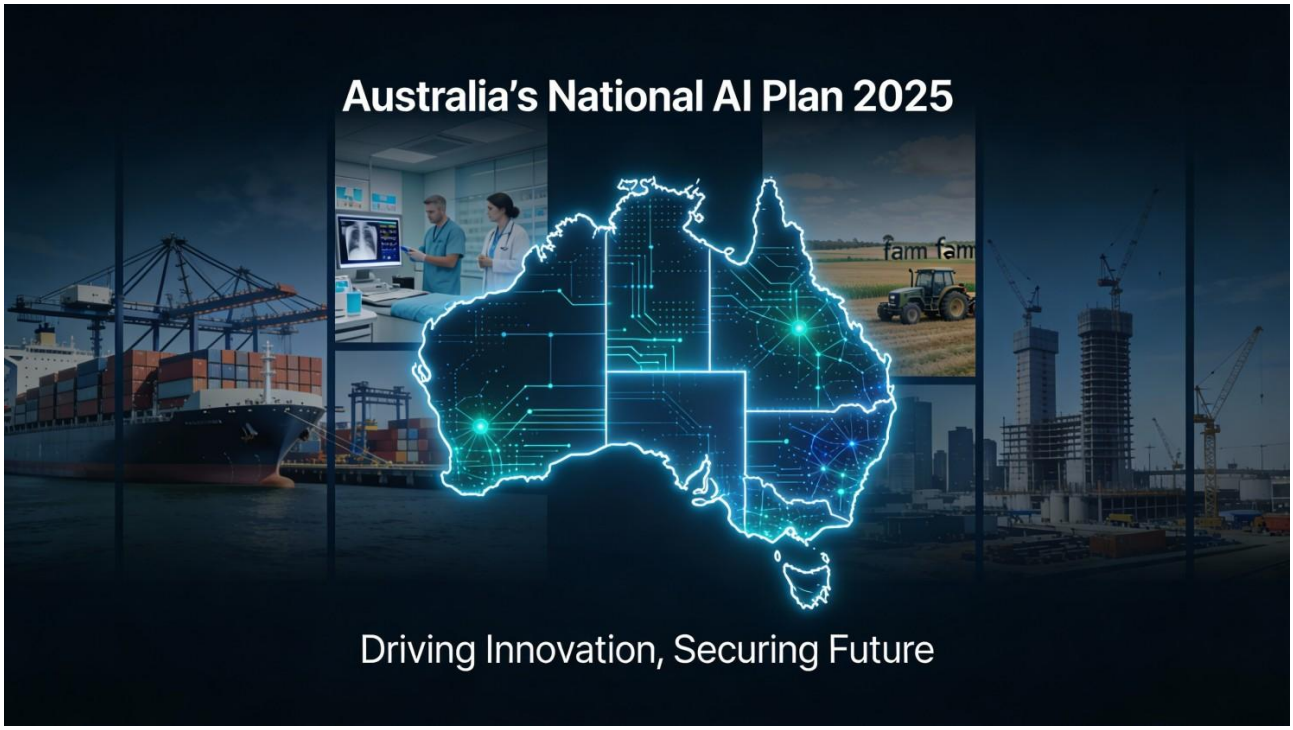
CONNECTING THE DOTS

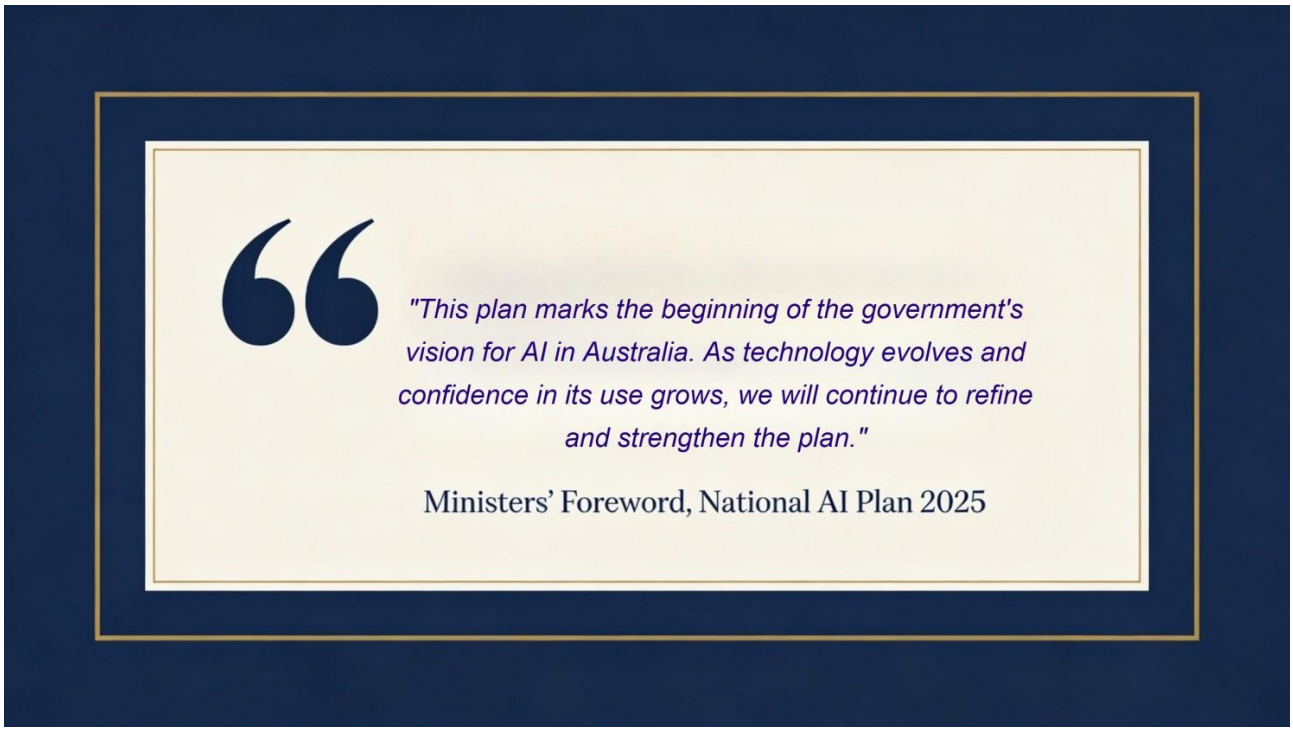
Australia's AI Strategy: A Global Comparison How does Australia’s AI governance stack up internationally?

By Dr Sukh Rashpal Sandhu | First published 22 December 2025

Over the past few weeks, I have been deeply engaged in projects, audits, interviews, and discussions, which left me little time to sit down and write an article on this important topic. Now that I have had the chance to reflect on my experiences and insights, I am excited to share my thoughts on the plan, its implications, and the opportunities it presents for Australia's industries.







The National AI Plan 2025 establishes a whole-of-government framework designed to ensure technology works for people by fostering a fairer, stronger Australia. While the plan marks a significant shift in national innovation, it is notably light on enforceable AI-specific obligations and definitive implementation timelines. Australia's approach currently lacks a formal risk classification system, leaving businesses to guess which governance measures apply to their systems. Furthermore, while the plan contemplates future mandatory guardrails for high-risk AI, it provides no

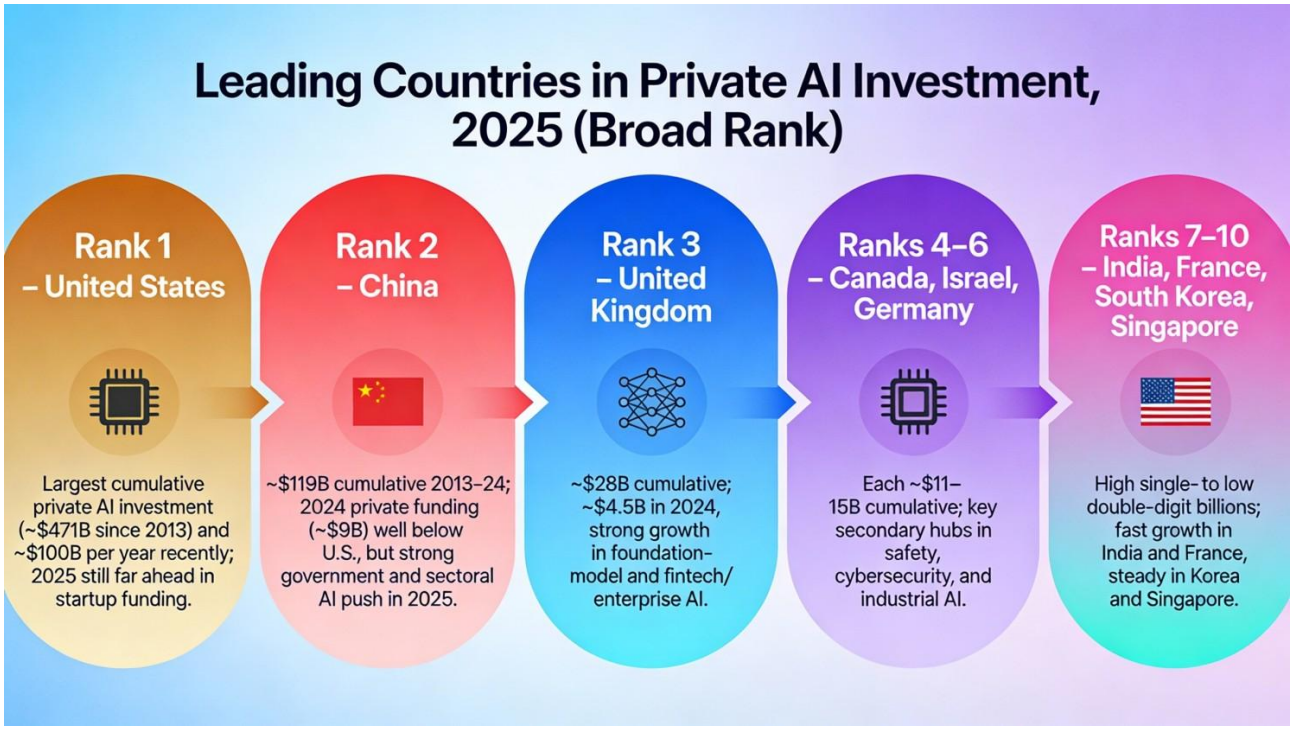
specific timeline or penalties for non-compliance, contrasting sharply with the phased, binding implementation of the EU AI Act.

A defining moment for the Australian industries



Australia's National AI Plan 2025 is more than just another policy document. It marks a significant shift in how our nation will innovate, compete, and grow in the coming decades. For those of us who have helped build Australia's economic and industrial foundations, this is both a pivotal and exciting moment.

The plan enters a global landscape of competing AI governance frameworks.



Each reflects unique political, economic, and philosophical priorities. For example, the European Union's AI Act is comprehensive but rigid. The United States' approach shifts with executive orders. The UK promotes pro-innovation principles, while Singapore offers a voluntary governance model. Meanwhile, China adopts a modular regulatory system.

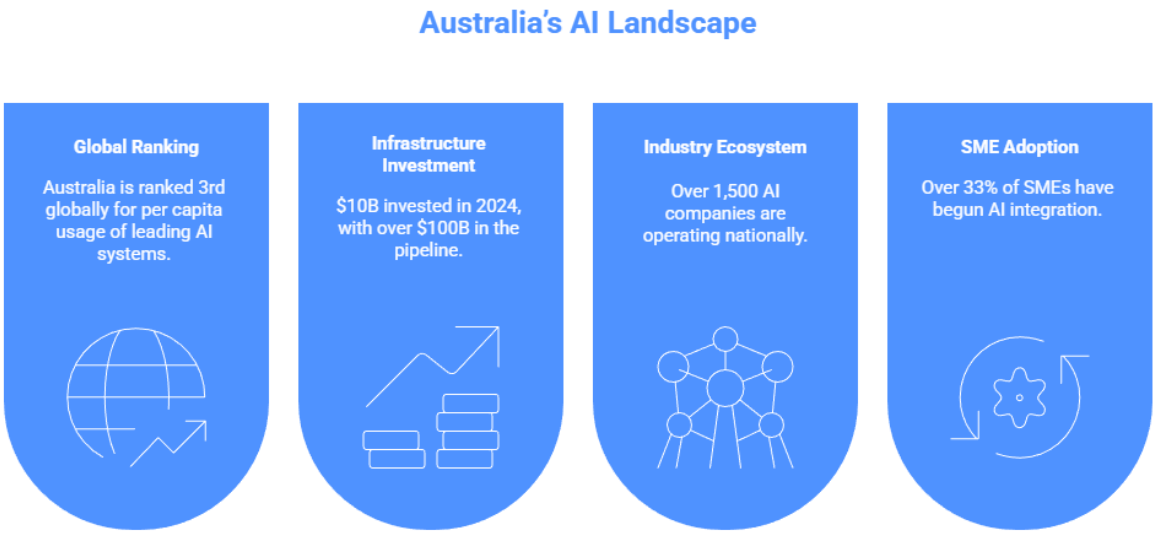


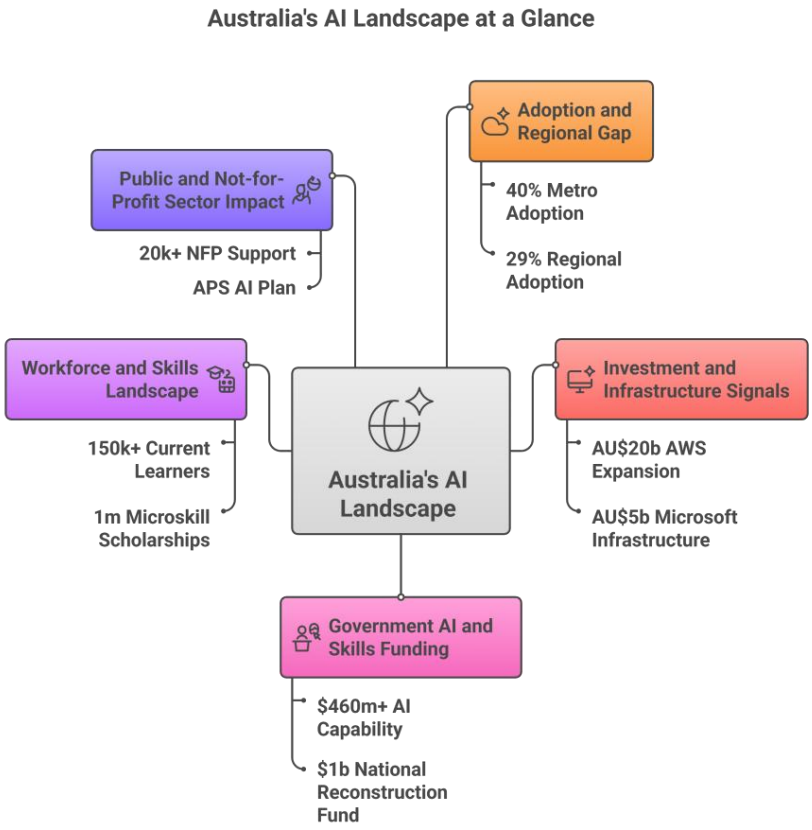
Australia takes a pragmatic middle path. It avoids the strictness of the EU and the lack of regulation seen under the Trump administration. However, this balance comes with risks. It may not offer enough protection for citizens or clarity for businesses seeking compliance.

Having worked in Australia's professional and industrial sectors for over 30 years, I have seen many technological changes reshape industries. But AI is different. It does not just change how tasks are done, it redefines what tasks are possible, what knowledge is accessible, and what new opportunities emerge.

The government's plan comes at the right time. The scale of AI's impact is already visible in our infrastructure and economic data.

Australia’s AI Landscape at a Glance



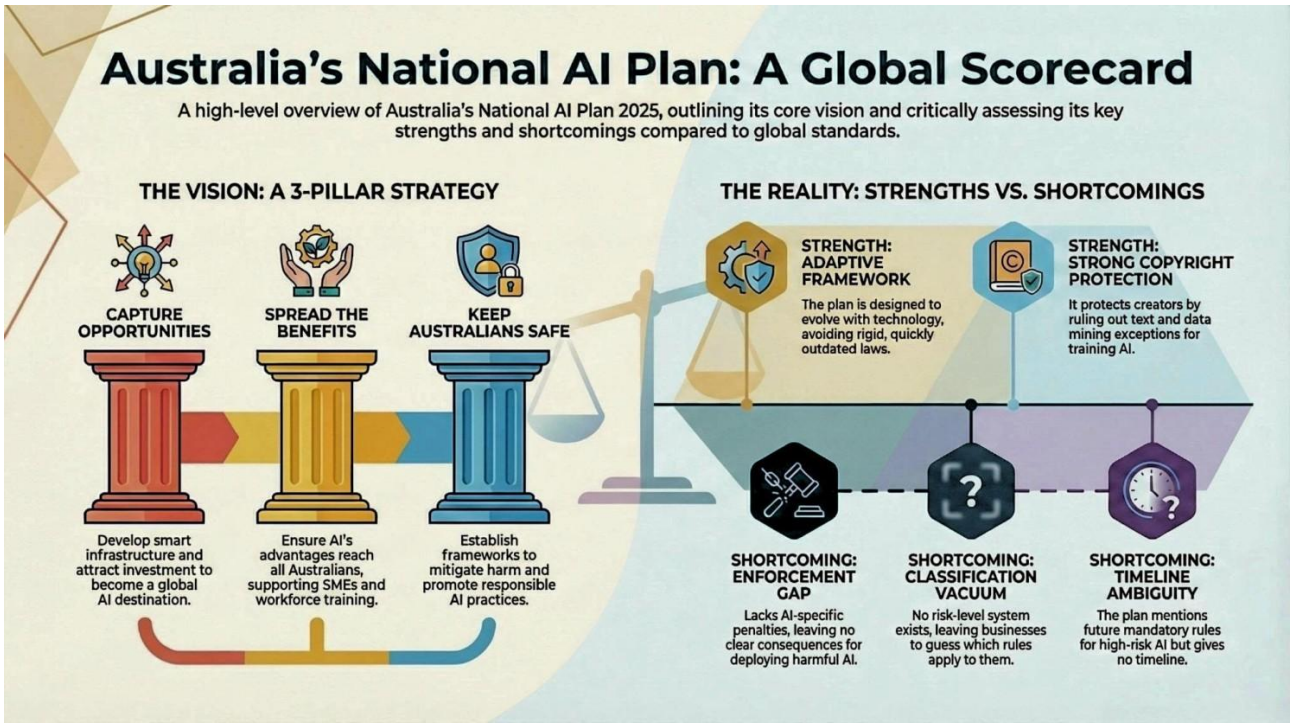


These figures tell an encouraging story, but they also reveal challenges we must address. The gap between metropolitan and regional AI adoption stands at 40% versus 29%. Many organisations across agriculture, manufacturing, and services remain unaware of AI's potential benefits. Furthermore, legitimate concerns about safety, privacy, and workforce disruption demand thoughtful, ethical responses.

This article provides a comprehensive analysis of the National AI Plan 2025, examining its implications for the full breadth of Australian industries, from the resources and energy sectors to healthcare, finance, and logistics.

> *My purpose is not merely to summarise the plan, but to translate its vision into actionable insights that business leaders, entrepreneurs, and industry stakeholders can apply immediately to ensure Australia remains a global leader in the intelligent economy.*

The Vision: AI for Australia



The plan's vision crystallises into a powerful statement:

AI for Australia:

1. **capturing opportunities,**
2. **spreading benefits,**
3. **keeping us safe.**

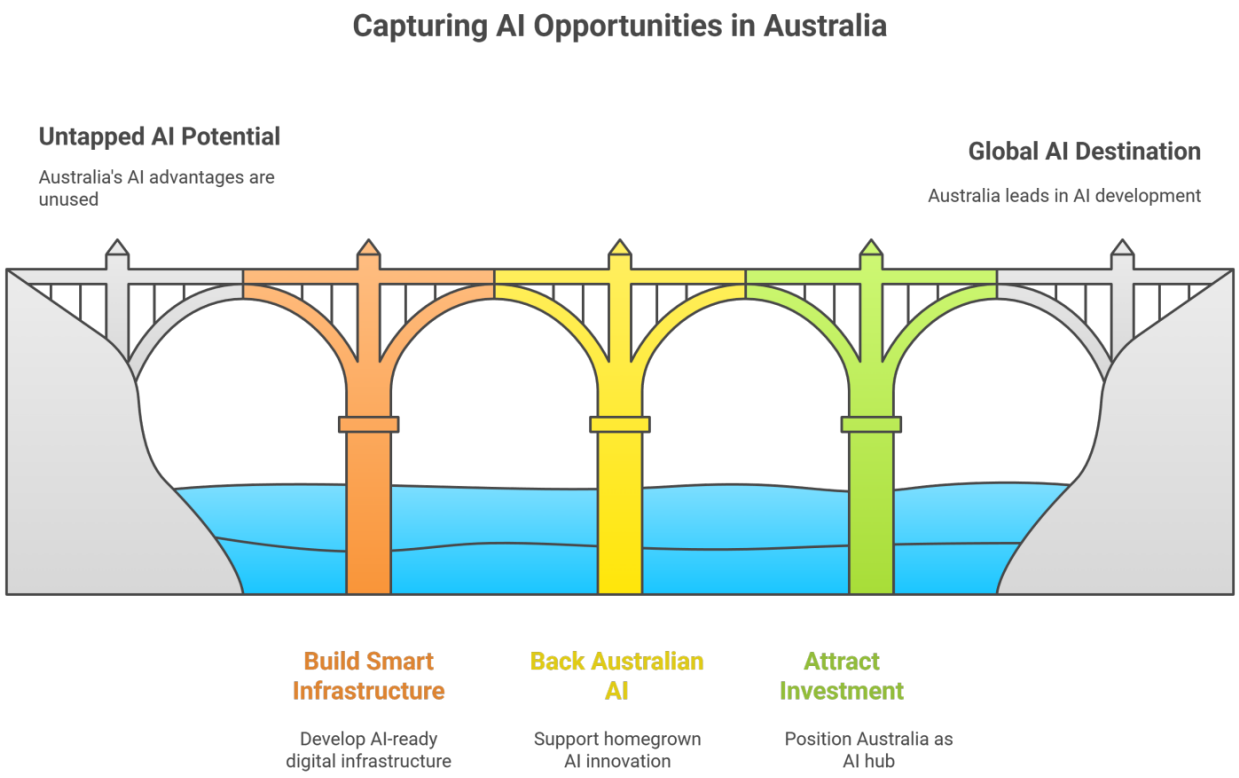
This tripartite framing reflects sophisticated policy thinking. Rather than treating AI as merely an economic opportunity to exploit or a risk to manage, the government recognises that sustainable success requires balancing all three dimensions simultaneously.

Let us examine each in turn.

Read more about the National AI Plan by following the link below.

[![National AI Plan # National AI Plan The National AI Plan is the Australian Government's plan to grow the AI industry in Australia. The plan sets out the steps the government will take to support Australia to build an AI-enabled economy that is more competitive, productive and resilient.... industry.gov.au](https://www.industry.gov.au/publications/national-ai-plan)

Goal One: Capture the Opportunities

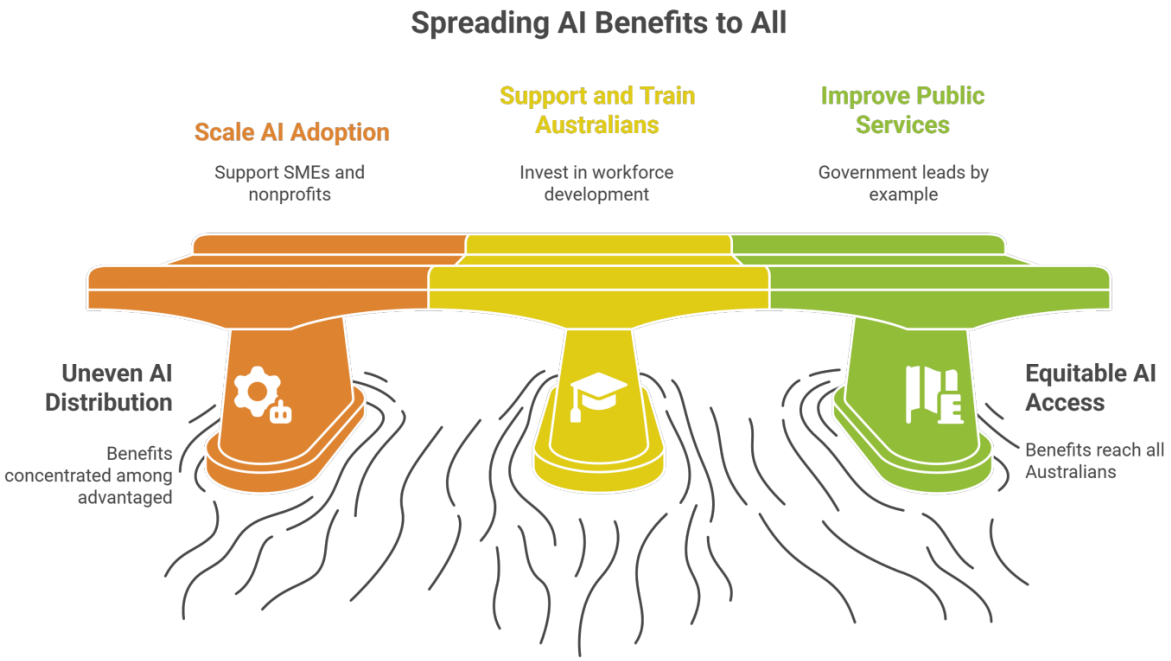


Australia possesses genuine competitive advantages in the global AI landscape. Our abundant renewable energy resources make us attractive for the energy-intensive data centres that power AI systems. Our geographic position and stable political environment appeal to international technology investors. Our world-class research institutions continue producing innovations that translate into commercial applications.

The plan identifies three strategic actions to capture these opportunities:

1. Build Smart Infrastructure: Developing the computing power, data centres, and digital connectivity that AI requires. Major investments from Firmus ($4.5 to $73.3 billion), Amazon ($20 billion), and Microsoft ($5 billion) demonstrate international confidence in Australia's infrastructure potential.
2. Back Australian AI Capability: Supporting homegrown AI innovation through $460 million in existing funding, $1 billion from the National Reconstruction Fund, and initiatives like the GovAI sovereign platform. Case studies like Harrison.ai, which received $32 million to develop medical imaging AI capable of detecting up to 124 clinical findings on chest X-rays, illustrate the practical impact of this investment.
3. Attract Investment: Positioning Australia as a global destination for AI development through the R&D Tax Incentive, streamlined regulatory frameworks, and strategic partnerships with leading technology companies.

Goal Two: Spread the Benefits

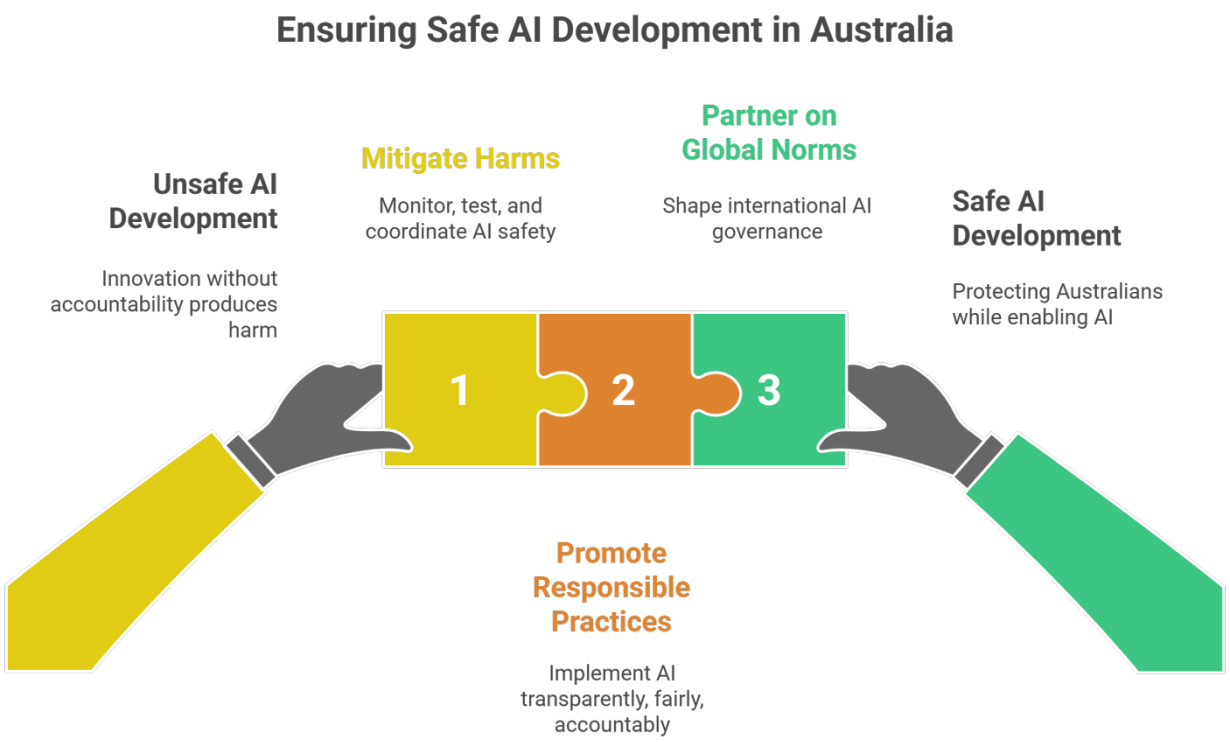


Economic opportunity means little if its benefits concentrate among the already advantaged. The plan explicitly commits to ensuring AI's transformative potential reaches all Australians, across geographic regions, business sizes, and demographic groups.

Three actions advance this goal:

1. Scale AI Adoption: The National AI Centre (NAIC) and AI Adopt Program ($17 million) provide guidance, training, and support to help organisations, particularly SMEs and not-for-profits, navigate AI implementation. Infoxchange's partnership with NAIC, which has helped over 20,000 not-for-profit organisations build AI capability, demonstrates this approach's effectiveness.
2. Support and Train Australians: Recognising that AI reshapes tasks rather than simply eliminating jobs, the plan invests heavily in workforce development. More than 150,000 learners are engaged in AI-related and digital skills programs across the VET and TAFE system, and Fee-Free TAFE is being used to support AI-related units and microcredentials. The $47 million Next Generation Graduates Program funds AI postgraduate scholarships.
3. Improve Public Services: Government leads by example through the AI Plan for the Australian Public Service, GovAI platform deployment, and practical applications like AI-enhanced lung cancer detection and National Library transcription projects.

Goal Three: Keep Australians Safe



Innovation without accountability produces harm. The plan establishes robust frameworks to protect Australians while enabling beneficial AI development.

Three actions deliver this protection:

1. Mitigate Harms: The Australian AI Safety Institute will monitor emerging capabilities, test AI systems for safety concerns, and coordinate information sharing across government. This addresses both upstream risks (how AI is built) and downstream harms (how AI affects people).
2. Promote Responsible Practices: NAIC guidance helps organisations implement AI transparently, fairly, and accountably. The Policy for Responsible Use of AI in Government establishes standards that the public sector will model for broader industry adoption.
3. Partner on Global Norms: Australia's participation in the Bletchley Declaration, Seoul Declaration, Paris Statement, UN Global Digital Compact, and Global Partnership on AI ensures our voice shapes international AI governance. Bilateral agreements with Singapore, the UK, South Korea, India, and the United States extend this influence.

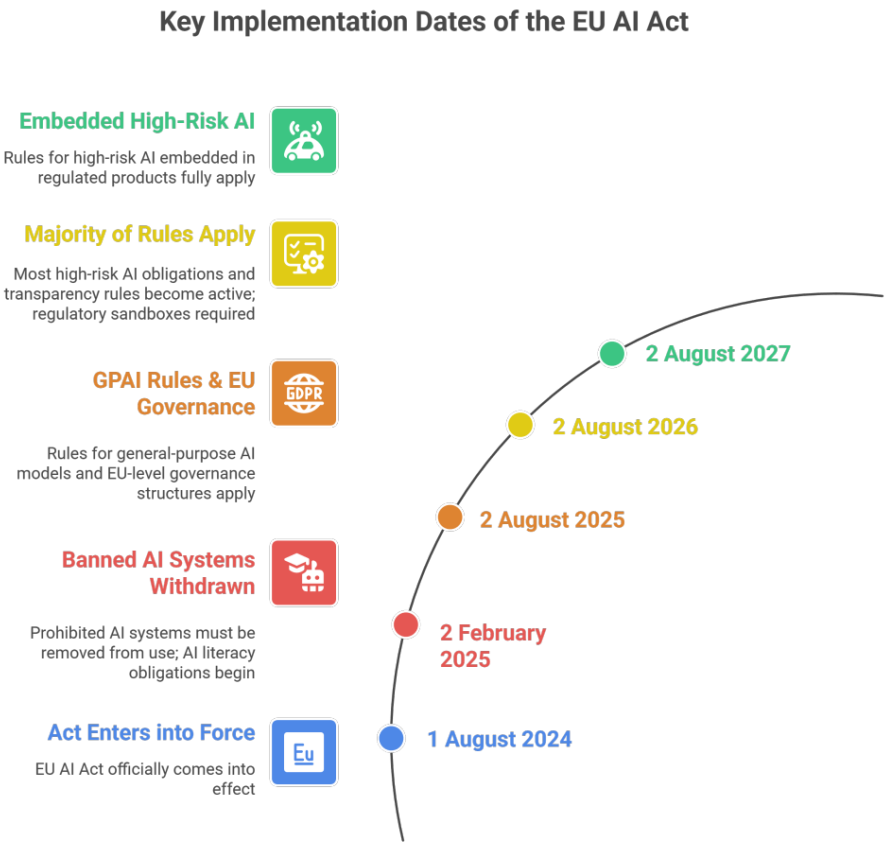
The Global AI Regulatory Landscape: A Fragmented Picture

Before examining Australia's position, we must understand the international context. The global AI governance landscape in late 2025 presents a study in philosophical contrasts and considerable instability.

The European Union

[![Artificial intelligence act | Think Tank | European Parliament
Artificial intelligence act | Think Tank | European Parliament Artificial intelligence act
europarl.europa.eu](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)698792)

The EU AI Act, which entered into force on 1 August 2024, represents the world's first comprehensive binding legal framework for artificial intelligence. Spanning 180 recitals and 113 articles, it establishes a risk-based classification system with four tiers: unacceptable risk (prohibited), high risk (heavily regulated), limited risk (transparency obligations), and minimal risk (largely unregulated).



The Act's extraterritorial reach means it affects any AI system whose outputs are used within the EU, regardless of where the provider is based. Non-compliance carries maximum penalties of EUR 35 million or 7% of worldwide annual turnover, figures that command serious attention from global technology companies.

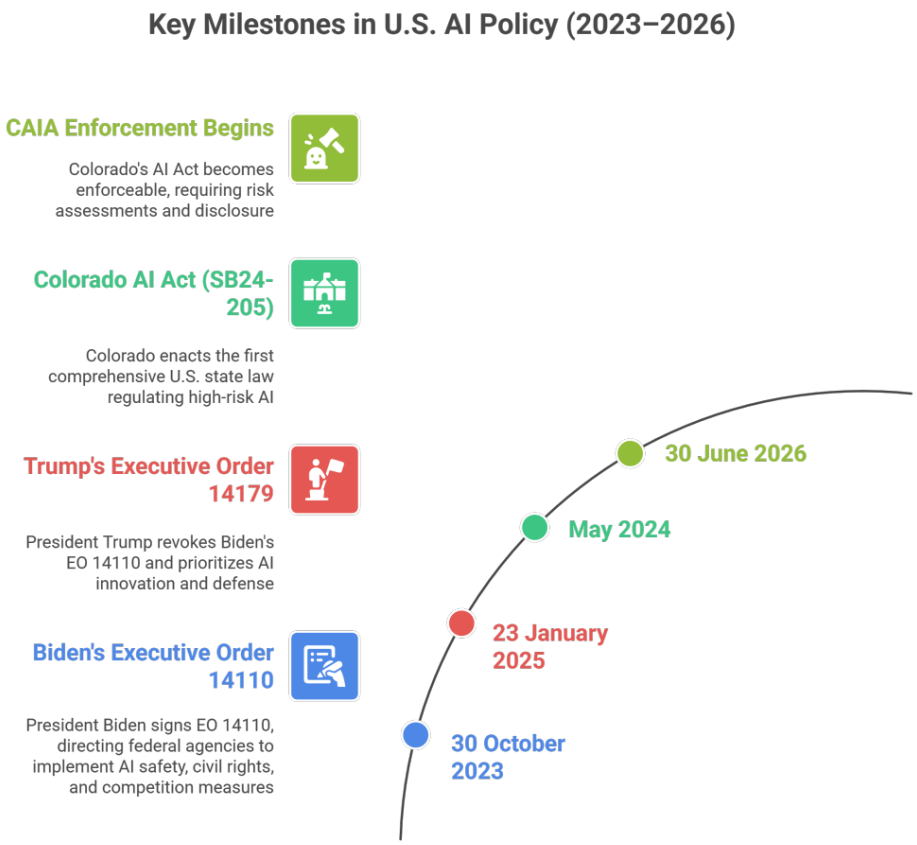


Prohibited practices include social scoring systems, real-time biometric identification in public spaces (with limited exceptions), and AI that deploys subliminal techniques to distort behaviour. High-risk AI systems face extensive requirements covering data governance, technical documentation, human oversight, and cybersecurity.

The United States: Regulatory Whiplash

[![Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence
Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence
Search, browse and learn about the Federal Register. Federal Register 2.0 is the unofficial daily publication for rules, proposed rules, and notices of Federal agencies and organizations, as well as executive orders and other presidential documents.
Federal Register](https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence)

The American approach presents a cautionary tale of governance instability. President Biden's Executive Order 14110, signed in October 2023, was described as "the most comprehensive piece of governance by the United States regarding AI." Its 110 pages directed over 50 federal agencies to undertake more than 100 specific actions, emphasising safety testing, transparency, equity, and civil rights protections.



Within hours of taking office on 20 January 2025, President Trump rescinded the Biden order entirely. His replacement Executive Order 14179, titled "Removing Barriers to American Leadership in Artificial Intelligence," explicitly frames AI development as a matter of national competitiveness requiring deregulation. The order criticises "engineered social agendas" and "ideological bias" in AI systems, terminology signalling fundamental disagreement with the Biden administration's equity focus.

A subsequent December 2025 executive order goes further, targeting "excessive" state AI laws and calling for a national standard that would pre-empt state-level regulation. The administration has explicitly criticised Colorado's algorithmic discrimination law as potentially requiring "false results" to avoid differential treatment of protected groups.

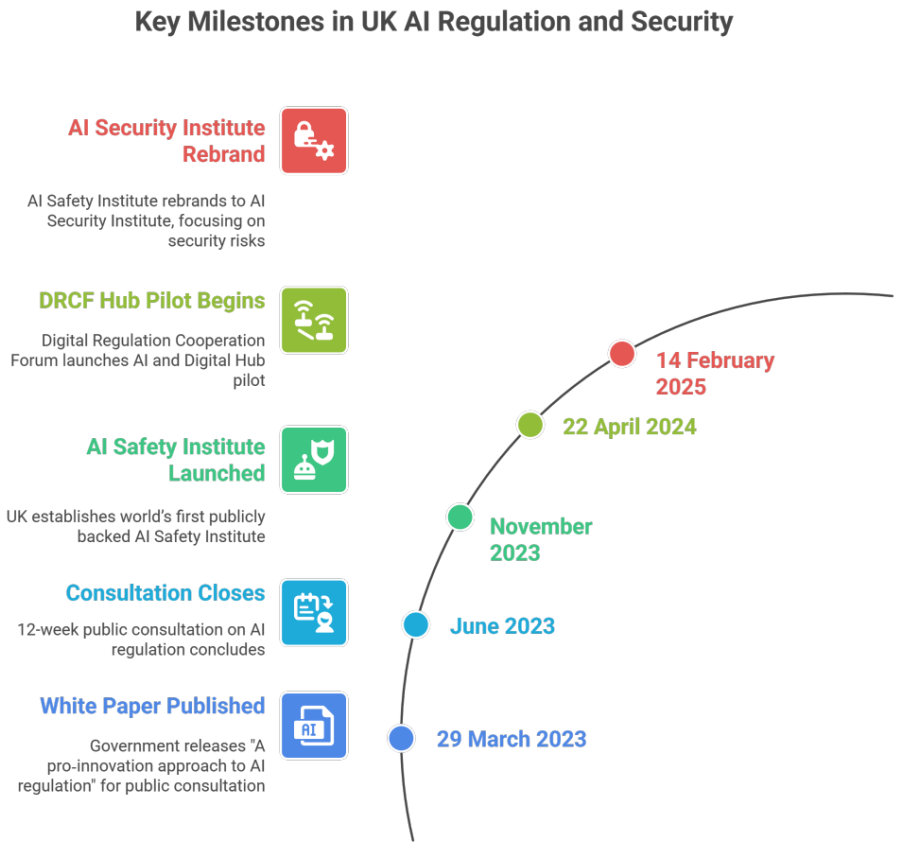


This regulatory volatility creates profound uncertainty. Policies established over years can be dismantled overnight. International partners and businesses cannot rely on American AI governance remaining stable across electoral cycles.

The United Kingdom: Pro-Innovation Principles

[![A pro-innovation approach to AI regulation # A pro-innovation approach to AI regulation AI helps piece together the first complete image of a black hole AI can enable scientific discovery. A computer vision model was used to piece together the first ever image of a black hole 55 million light years away, combining images from 8... GOV.UK](https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper)

The UK has deliberately chosen a different path from its former EU partners. Rather than comprehensive AI-specific legislation, the 2023 White Paper "A Pro-Innovation Approach to AI Regulation" establishes five non-statutory, cross-sectoral principles: safety, transparency, fairness, accountability, and contestability.



These principles are implemented by existing sectoral regulators, including the Information Commissioner's Office, Ofcom, the Competition and Markets Authority, and the Financial Conduct Authority, each interpreting and applying them within their domains. The Digital Regulation Cooperation Forum coordinates cross-regulator AI efforts through initiatives like the AI and Digital Hub pilot.

The UK government has indicated that comprehensive AI legislation could emerge in 2026, but current signals emphasise continued "pro-growth and pro-innovation" regulation. The February 2025 rebranding of the AI Safety Institute as the AI Security Institute signals a pivot toward national security concerns over broader ethical considerations.

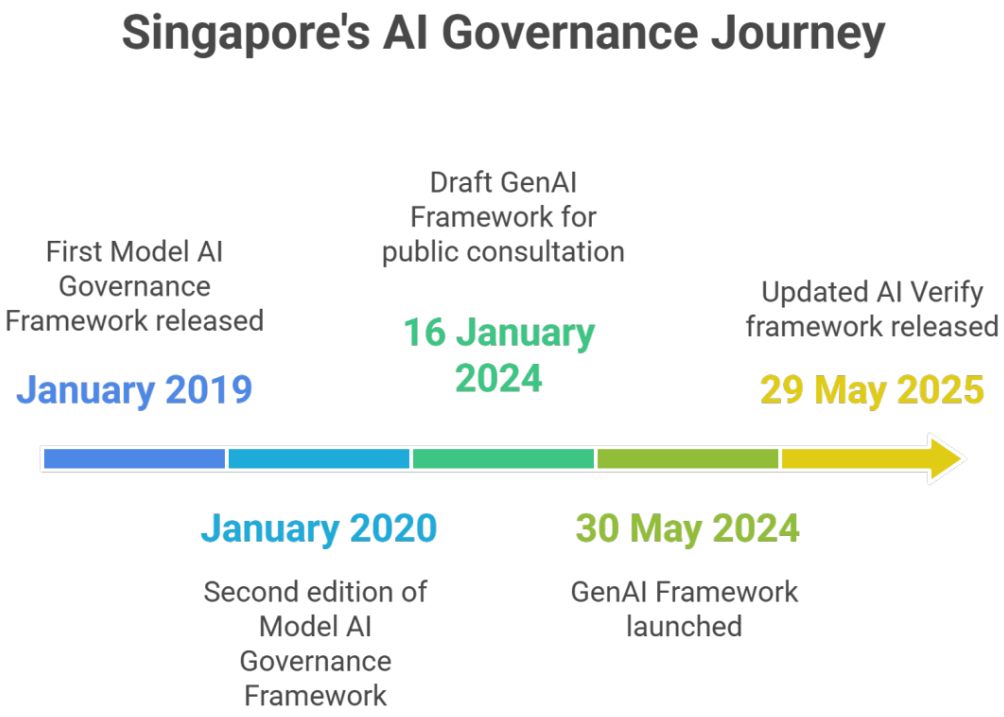


Critics argue this approach risks regulatory fragmentation, inconsistent interpretation across sectors, and inadequate protection for citizens. Supporters contend it enables flexibility and avoids premature constraints on rapidly evolving technology.

Singapore: Voluntary Governance Frameworks

Untitled MODEL ARTIFICIAL INTELLIGENCE GOVERNANCE FRAMEWORK 1 ARTIFICIAL INTELLIGENCE GOVERNANCE FRAMEWORK MODEL SECOND EDITION SUMMARY OF UPDATES... pdpc.gov.sg

Singapore pioneered AI governance frameworks with its 2019 Model AI Governance Framework, updated in 2020 and expanded in May 2024 to address generative AI. The Infocomm Media Development Authority (IMDA) and AI Verify Foundation have developed comprehensive guidance across nine dimensions: accountability, data governance, trusted development, safety, content provenance, disclosure, third-party concerns, incident reporting, and AI for public good.

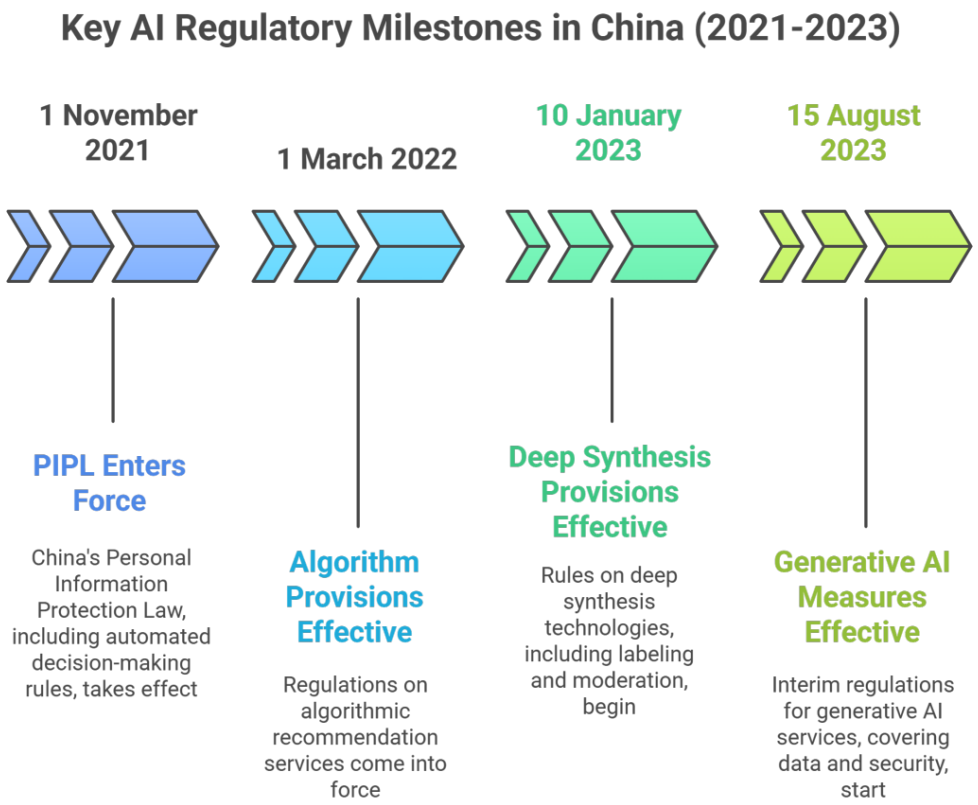


Singapore's approach remains fundamentally voluntary. The Model AI Governance Framework provides best practice guidance rather than binding legal requirements. However, the framework explicitly contemplates that existing laws, including personal data protection, cybersecurity, and sector-specific regulations, already address many AI-related concerns. Singapore has signalled willingness to introduce binding measures if voluntary approaches prove insufficient.



The AI Verify testing toolkit, released as open-source software, allows organisations to assess their AI systems against governance principles. This practical tooling distinguishes Singapore's approach from purely theoretical frameworks.

China: Modular Targeted Regulation



China presents the most distinctive regulatory model and potentially the most instructive for understanding alternatives to comprehensive horizontal legislation. Rather than a single overarching AI law (though one is reportedly in preparation), Chinese authorities have issued targeted regulations addressing specific AI applications.

The 2022 Algorithm Recommendation Provisions govern AI systems that influence public opinion or drive social engagement. The 2023 Deep Synthesis Provisions address deepfakes and synthetically generated content. The August 2023 Generative AI Measures became among the world's first regulations specifically targeting generative AI, requiring algorithm registration, security assessments, and content that is "true and accurate."

China's algorithm registry, administered by the Cyberspace Administration of China (CAC), now contains over 3,700 registered algorithms from more than 2,300 companies. This provides regulators unprecedented visibility into the AI ecosystem, a level of transparency that Western jurisdictions have not achieved.



The Chinese model demonstrates that agile, technology-specific regulation can address emerging risks without the multi-year legislative processes that characterise EU-style comprehensive frameworks.

Comparative Analysis: Key Dimensions

AI Regulation Comparison

Characteristic	AI-Specific Statute	Fines/Penalties	Enforcement	Registration
European Union	Yes	Up to EUR 35 million or 7% turnover	AI Office, AI Board, Scientific Panel	High-risk AI systems
United States	No	USD 20,000 per violation (Colorado)	Attorney General (Colorado)	No
United Kingdom	No	GBP 17.5 million or 4% turnover	ICO, FCA	No
Singapore	No	10% of turnover, capped at SGD 1 million	PDPC	No
China	Yes	Up to RMB 50 million or 5% turnover	Regulators can order suspension or fines	Algorithms with public-opinion functions
Australia	No	Up to AUD 50 million, 3x benefit, or 30% turnover	OAIC, Australian Consumer Law	No

The following analysis examines how Australia's approach compares across critical governance dimensions:

Legal Force and Enforceability

EU AI Act: Binding regulation with direct effect across 27 member states. Maximum penalties of EUR 35 million or 7% of global turnover. Third-party conformity assessments required for high-risk systems.

United States (Trump): Executive orders without legislative force, subject to reversal by future administrations. No federal AI-specific legislation. Emphasis on preempting state laws.

United Kingdom: Non-statutory principles implemented through existing regulatory frameworks. No AI-specific penalties beyond those available under existing laws.

Singapore: Voluntary governance frameworks. Existing laws (PDPA, cybersecurity) provide enforcement mechanisms. AI Verify toolkit enables self-assessment.

China: Binding regulations with administrative enforcement. Penalties up to RMB 15 million or 5% of turnover. Algorithm registration mandatory.

Australia: Policy framework without legislative force. Relies on existing regulatory frameworks (Privacy Act, consumer law). AI Safety Institute established for monitoring. Future mandatory guardrails contemplated.

Critical Assessment: Australia's approach currently lacks enforceable AI-specific obligations. While the plan contemplates future mandatory measures for high-risk AI, it provides no timeline or specifics. This creates uncertainty for businesses seeking compliance clarity and limited recourse for citizens harmed by AI systems.

Risk Classification Frameworks

EU AI Act: Four-tier classification (unacceptable, high, limited, minimal risk) with detailed criteria and sector-specific annexes. Additional category for general-purpose AI models with systemic risk thresholds.

United States (Biden, rescinded): Distinguished "rights-impacting" and "safety-impacting" AI. Focused on foundation models posing national security risks.

United Kingdom: No formal risk classification. Regulators apply principles contextually within their sectors.

Singapore: No mandatory classification. Guidance distinguishes traditional AI from generative AI with different governance considerations.

China: The Generative AI Measures contemplate "categorical and hierarchical regulation" but detailed classification criteria remain under development.

Australia: No formal risk classification system. The AI Safety Institute will "monitor emerging capabilities" and "test AI systems for safety concerns," but without defined thresholds or categories.

Critical Assessment: Australia's lack of defined risk categories represents a significant gap. Organisations cannot determine with certainty what governance measures apply to their AI systems. The EU's classification approach, while complex, provides actionable clarity that Australia's framework currently lacks.

Balancing Innovation and Protection

EU AI Act: Criticised by some as innovation-inhibiting through compliance burdens. Includes regulatory sandboxes and SME exemptions to mitigate impact. The December 2024 Digital Package on Simplification proposes amendments to "ensure the rules remain clear, simple, and innovation-friendly."

United States (Trump): Explicitly prioritises innovation over regulation. Frames safety requirements as "barriers" to be removed. No mandatory consumer protections.

United Kingdom: "Pro-innovation" approach explicitly stated. Prime Minister confirmed in March 2025 that regulation would be "pro-growth and pro-innovation." Comprehensive legislation delayed until 2026 or later.

Singapore: Explicitly seeks to "address generative AI concerns while facilitating innovation." GenAI Sandbox provides SMEs access to enterprise AI solutions for experimentation.

China: Balances innovation promotion with content control imperatives. State actively supports AI development through funding, open-source initiatives, and infrastructure investment alongside regulatory requirements.

Australia: Three-goal framework explicitly includes both "capture opportunities" and "keep Australians safe." Substantial investment in infrastructure ($10+ billion data centre investment) alongside AI Safety Institute establishment.

Critical Assessment: Australia's rhetorical balance between innovation and safety does not yet translate into concrete mechanisms that achieve both. The investment figures demonstrate commitment to infrastructure; the safety mechanisms remain largely aspirational.

Copyright and Training Data

EU AI Act: GPAI providers must comply with Copyright Directive and publish training data summaries. The text and data mining exception allows training on publicly available content unless rights holders opt out.

United Kingdom: December 2024 consultation proposed text and data mining exception with opt-out rights for creators, similar to the EU approach. Consultation closed February 2025; outcome awaited.

United States: No federal AI-specific copyright provisions. Pending litigation (New York Times v. OpenAI, etc.) may establish precedents. Fair use doctrine provides uncertain protection for training.

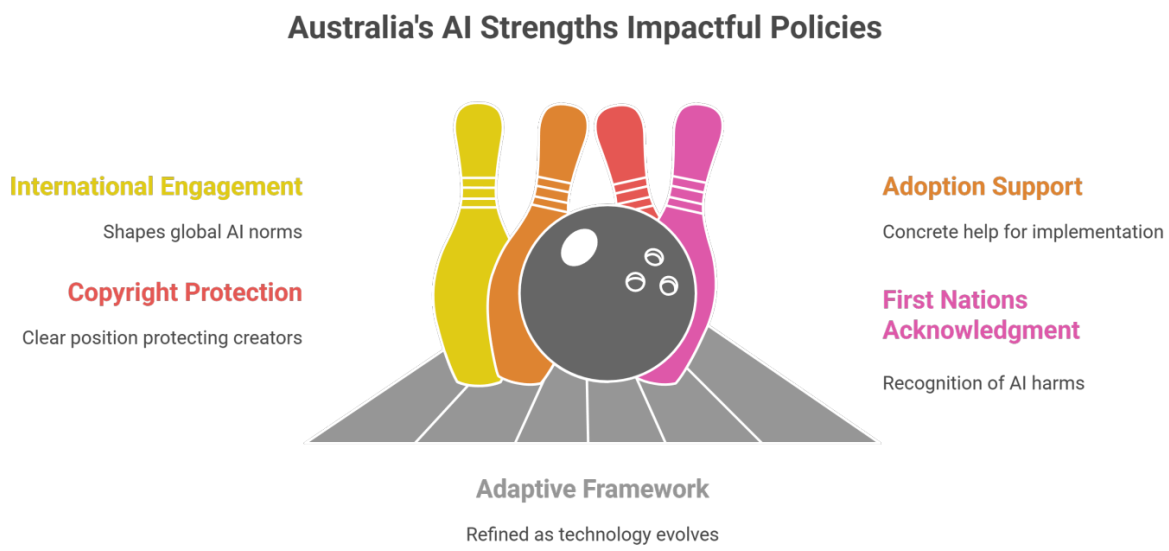
Singapore: GenAI Framework recommends that policymakers "foster open dialogue among all relevant stakeholders" on copyright issues. No specific requirements established.

China: Courts have begun addressing AI-generated content copyright. January 2024 ruling found AI-generated images can qualify for copyright protection if meeting originality standards.

Australia: The National AI Plan explicitly rules out text and data mining exceptions, protecting content creators. AI developers must secure appropriate permissions for training data.

Critical Assessment: Australia's decision to protect copyright holders represents a significant departure from the EU and UK direction. This is good news for content creators but may disadvantage Australian AI developers relative to competitors in jurisdictions with broader training exemptions.

What Australia Gets Right



Critical analysis requires acknowledging genuine strengths alongside weaknesses. Here is where Australia got it right:

1. Adaptive Framework Design

Unlike the EU AI Act's fixed legislative structure requiring formal amendment processes, Australia's plan explicitly describes itself as "a beginning" that will be refined as technology evolves. This adaptive design reflects wisdom about the pace of AI development. The EU is already proposing simplification amendments less than eighteen months after the Act entered into force, suggesting that more flexible approaches may prove more durable.

2. International Engagement Strategy

Australia's participation in the Bletchley Declaration, Seoul Declaration, Paris Statement, and Global Partnership on AI provides genuine influence in shaping international AI norms. Bilateral agreements with Singapore, UK, South Korea, India, and the United States extend this reach. As a middle power, Australia cannot unilaterally set global standards, but effective multilateral engagement can amplify our voice beyond our economic weight.

3. Practical Focus on Adoption Support

The National AI Centre and AI Adopt Program provide concrete support for organisations navigating AI implementation. The partnership with Infoxchange, which has helped over 20,000 not-for-profit organisations build AI capability, demonstrates practical impact. Many international frameworks focus on regulation without equivalent attention to enabling adoption. Australia's dual focus represents sound policy design.

4. Copyright Protection Decision

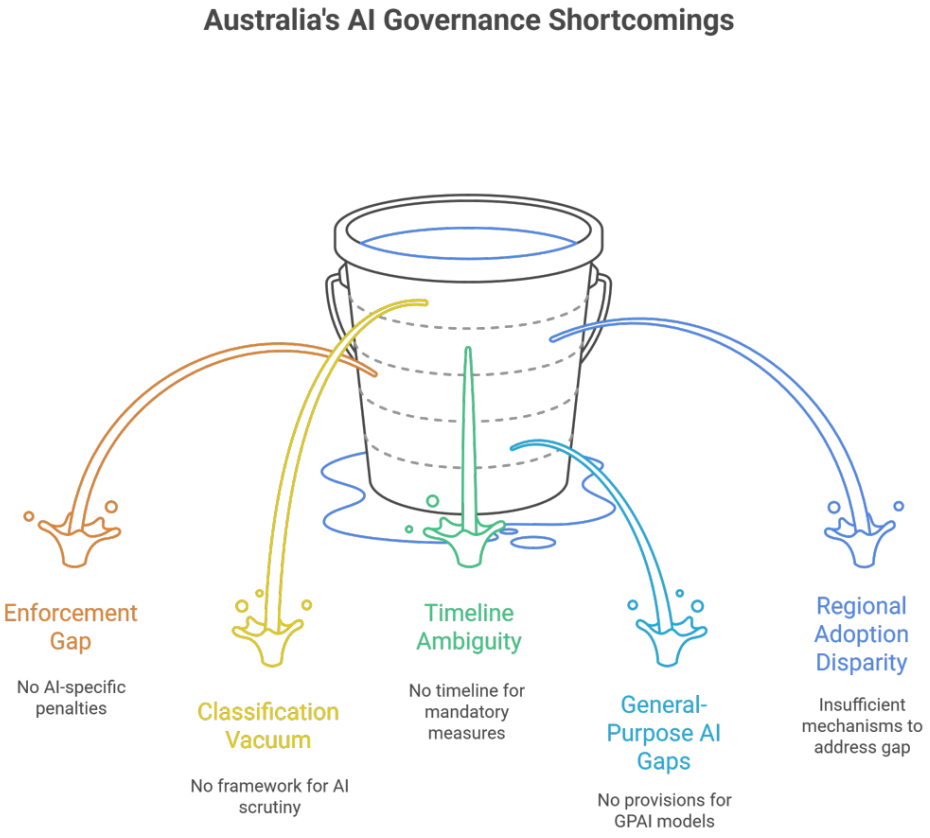
By ruling out text and data mining exceptions, Australia has taken a clear position protecting creative industries. This contrasts with the UK's consultation and the uncertain

US litigation landscape. Certainty, even if some stakeholders disagree with the direction, has value for both creators and developers who can now plan accordingly.

5. First Nations Acknowledgment

The explicit recognition of AI harms to First Nations communities, including perpetuation of harmful stereotypes and misuse of cultural intellectual property, is notably absent from most international frameworks. The reference to the Kakadu AI project as a model for respectful Traditional Owner governance over cultural knowledge provides concrete guidance beyond rhetorical acknowledgment.

What Australia Gets Wrong



Honest assessment requires equal frankness about weaknesses. According to me, these are the areas that need improvement:

1. Enforcement Gap

The plan contains no AI-specific enforcement mechanisms, penalties, or compliance requirements. Unlike the EU's EUR 35 million fines or China's algorithm registration requirements, Australia relies entirely on existing regulatory frameworks.

This creates a fundamental asymmetry: the plan articulates principles without providing tools to ensure adherence. The AI Safety Institute will "monitor" and "test" AI systems, but what happens when problems are identified? The plan is silent on consequences for

deploying harmful AI. Without enforcement teeth, the framework may amount to well-intentioned guidance that bad actors can safely ignore.

2. Classification Vacuum

Australia provides no framework for determining which AI systems require greater scrutiny. The EU's four-tier risk classification, while complex, tells developers exactly what category their system falls into and what requirements apply. Australia's plan leaves organisations guessing.

This ambiguity may seem business-friendly in the short term, but it creates long-term uncertainty. Organisations investing in AI systems cannot determine what governance measures will ultimately be required. Some will over-invest in compliance; others will under-invest and face retrospective requirements.

3. Timeline Ambiguity

The plan contemplates future mandatory measures, including "guardrails for high-risk AI," but provides no timeline for their introduction. Contrast this with the EU AI Act's explicit phased implementation: prohibited practices from February 2025, GPAI requirements from August 2025, high-risk system requirements from August 2026 and 2027.

Businesses require planning certainty. Will binding requirements emerge in 2026? 2028? Never? The plan's silence on timing undermines the ability of responsible organisations to prepare.

4. General-Purpose AI Gaps

The EU AI Act includes specific provisions for general-purpose AI (GPAI) models, including transparency requirements, training data documentation, and additional evaluations for high-capability models with systemic risk potential. Singapore's GenAI Framework explicitly addresses generative AI governance across nine dimensions.

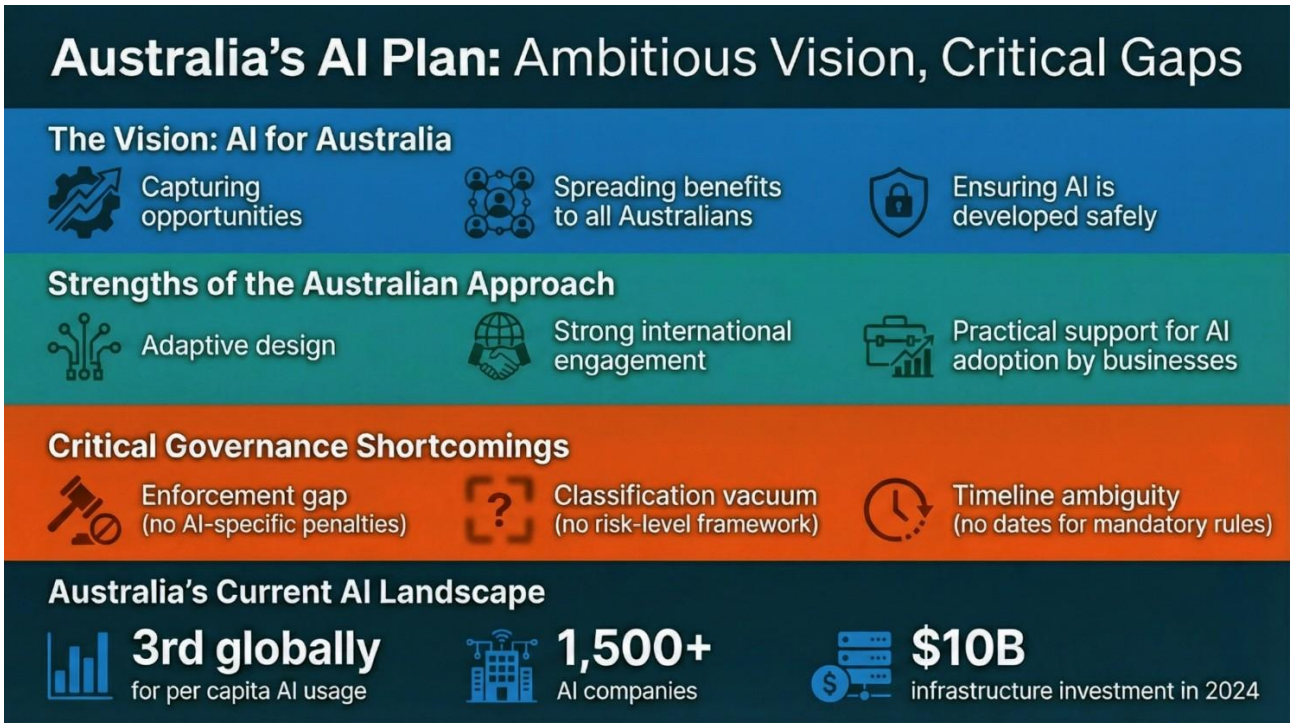
Australia's plan does not differentiate governance requirements based on AI model architecture or capability. As foundation models become increasingly central to AI deployment, this gap will become more significant.

5. Regional Adoption Disparity

The plan acknowledges the 40% metropolitan versus 29% regional AI adoption gap but provides insufficient mechanisms to address it. Digital connectivity improvements through NBN expansion and satellite coverage address infrastructure, but infrastructure alone does not drive adoption.

Without targeted programs addressing regional skills shortages, awareness gaps, and support structures, the plan risks exacerbating rather than reducing regional digital divides.

Comparative Verdict: Where Does Australia Stand?



Placing Australia on the global spectrum requires acknowledging that different approaches suit different national circumstances:

More Protective Than

United States (Trump administration): Australia maintains focus on safety, worker consultation, and benefit distribution that the current US approach explicitly rejects.

United Kingdom: Australia's AI Safety Institute and explicit commitment to future guardrails exceeds the UK's entirely voluntary current framework.

Less Protective Than

European Union: The EU's binding requirements, risk classification, and substantial penalties provide significantly stronger citizen protection than Australia's framework.

China: Despite authoritarian governance concerns, China's mandatory algorithm registration, content requirements, and active enforcement exceed Australia's voluntary approach.

Comparable To

Singapore: Both nations employ voluntary frameworks with practical tooling (Singapore's AI Verify, Australia's NAIC guidance) while contemplating future binding measures. Singapore's more developed governance dimensions and testing toolkit give it a slight edge in implementation maturity.

Lessons Australia Should Learn

Each international framework offers insights Australia could usefully adopt:

From the European Union

1. Clear risk classification: Develop defined categories with explicit criteria so organisations know what governance applies to their systems.
2. Phased implementation timelines: Announce binding requirements with sufficient lead time for preparation.
3. Proportionate penalties: Ensure consequences for non-compliance are sufficient to influence behaviour.

From Singapore

1. Practical testing toolkits: Develop Australian equivalents to AI Verify that enable organisations to assess their AI systems against governance principles.
2. GenAI-specific guidance: Expand framework to address generative AI's distinct governance challenges.

From China (With Appropriate Democratic Safeguards)

1. Algorithm transparency: Consider registration or disclosure requirements for high-impact AI systems to provide regulatory visibility.
2. Modular regulation: Address emerging AI applications through targeted measures rather than waiting for comprehensive legislation.

From the United States (As Cautionary Tale)

1. Bipartisan foundations: Build cross-party consensus for AI governance to prevent policy oscillation across electoral cycles.
2. Legislative codification: Embed critical protections in legislation rather than ministerial discretion to ensure durability.

Implications for Australian Industries

The National AI Plan creates distinct implications across Australia's major industry sectors. Understanding these nuances enables strategic responses tailored to each sector's unique circumstances, challenges, and opportunities.

The sector analyses that follow interpret the plan's directions and examples for specific industries, rather than restating explicit commitments word-for-word from the policy document.

Vocational Education and Training Sector

For Registered Training Organisations (RTOs), the plan presents both immediate opportunities and longer-term transformation imperatives.

Immediate Opportunities:

New AI-focused units of competency and microcredentials are entering training packages, creating immediate pathways for RTOs to expand course offerings.

Fee-Free TAFE funding for AI skills development stimulates learner demand that private RTOs can complement through specialised, niche or advanced training solutions.

Emerging partnerships between RTOs and AI-adopting industries create new work-integrated learning models and employer-sponsored upskilling programs.

The National AI Centre's openly accessible resources provide high-quality, free content that RTOs can contextualise for delivery within existing courses.

Transformation Imperatives:

RTOs must cultivate internal AI capability to model the competencies they deliver. Trainers and assessors require structured professional development in AI literacy and data ethics.

AI tools will increasingly streamline administrative workflows, from enrolment management to compliance reporting. Early adopters will achieve measurable efficiency and quality gains.

Quality assurance and integrity systems must evolve to manage AI-generated learning content, AI-assisted assessments, and sophisticated forms of academic integrity monitoring.

The Standards for RTOs 2025 will heighten expectations around demonstrating contemporary, technology-integrated training practices, including responsible AI usage.

Healthcare and Allied Health

The healthcare sector stands among AI's most promising domains, with the National AI Plan highlighting flagship initiatives that enhance diagnostics, service delivery, and clinical outcomes.

AI-powered diagnostics, such as Harrison.ai's $32 million investment in imaging analysis capable of detecting up to 124 conditions on chest X-rays, illustrate how AI augments rather than replaces human expertise.

Government-led AI trials have delivered measurable outcomes, validating AI's transformative potential in patient safety and efficiency.

For healthcare education providers, these advances necessitate the integration of AI literacy, ethics, and digital tool competency into health and clinical curricula. Future practitioners must be prepared to interpret, monitor, and collaborate with AI systems to deliver safe, effective, and ethical clinical care.

Agriculture and Resources

Australia's primary industries leverage AI to strengthen productivity, sustainability, and competitiveness, ranging from precision farming to biosecurity monitoring and mineral exploration.

AI applications in crop health analytics, livestock monitoring, and predictive maintenance enhance resource efficiency and reduce operational risk.

Regional disparities remain a barrier, with AI adoption at 40% in metropolitan areas versus 29% in regional enterprises, largely due to connectivity gaps and digital skill shortages.

The National AI Plan addresses this through commitments to expanded NBN infrastructure, satellite connectivity, and community-based AI upskilling programs.

Industry associations and training providers can reinforce these efforts through targeted outreach and regional capacity-building initiatives that drive equitable technology access.

Manufacturing and Construction

AI stands poised to revolutionise manufacturing and construction through improved safety systems, predictive maintenance, and process optimisation.

Infrastructure investments outlined in the plan accelerate AI readiness through smart manufacturing networks, digital twins, and sensor-enabled work environments.

AI-driven predictive quality control can significantly reduce material waste and improve compliance with safety and environmental standards.

Workforce transition must be proactive. The plan reinforces that AI reshapes job functions rather than eliminating them, yet reskilling is critical to ensure workforce inclusion. Transparent communication and joint consultation with employees and unions about AI adoption are essential to uphold both ethical and legislative frameworks while maintaining trust.

Financial and Professional Services

In financial, legal, and consulting domains, AI's growing application to data analysis, compliance monitoring, and client interaction drives transformative gains.

AI automates document analysis, risk detection, and regulatory reporting, freeing professionals to focus on higher-value advisory work.

The plan's emphasis on responsible AI governance supports firms in developing compliant, auditable AI frameworks aligned with privacy and ethics standards.

Copyright and intellectual property compliance emerges as a central focus. With the government ruling out text and data mining exceptions, firms must ensure that AI tools and datasets are licensed appropriately and that outputs respect creator rights.

Firms adopting AI must also embed ethical decision-making protocols that maintain transparency and client trust in algorithmic outputs.

Transport and Logistics

AI promises to reshape Australia's transport and logistics networks through smarter, safer, and more sustainable systems.

Immediate Opportunities:

The National AI Plan's focus on smart infrastructure creates funding and partnership opportunities for transport authorities and logistics companies.

AI-enabled route optimisation and predictive maintenance improve efficiency in freight and passenger transport, reducing fuel use and service disruptions.

Autonomous and semi-autonomous vehicle trials receive regulatory support under the plan, paving the way for scalable deployment in controlled environments.

Logistics firms can integrate AI demand forecasting to streamline warehousing, inventory management, and last-mile delivery systems.

Transformation Imperatives:

Workforce reskilling is essential for roles affected by automation, including drivers, controllers, and supply chain coordinators.

Data privacy and cybersecurity frameworks must protect sensitive transport and operational data as connectivity expands.

Collaboration between government, industry, and unions is necessary to ensure equitable transitions and uphold safety standards.

Defence and National Security

Australia's defence and security agencies are pivotal players in the responsible and strategic application of AI technologies.

Immediate Opportunities:

AI enhances threat detection, surveillance, and decision support capabilities through advanced analytics and real-time intelligence consolidation.

The plan supports collaborative R&D between Defence Science and Technology Group (DSTG), universities, and sovereign technology firms to ensure strategic autonomy in AI development.

Robotics and unmanned aerial systems benefit from AI integration for situational awareness, reducing risks to personnel.

Transformation Imperatives:

Ethical governance is paramount. Defence AI systems must align with Australia's obligations under international law and human rights conventions.

National security frameworks must manage dual-use risks and ensure transparency in AI deployment to avoid unintended escalations.

Continuous professional development in AI ethics, data management, and strategic foresight is essential for defence leadership and operations teams.

Energy and Utilities

In this sector, AI enables smarter energy management, renewable integration, and resource optimisation.

Immediate Opportunities:

AI-driven predictive maintenance can prevent equipment failures across power grids, pipelines, and renewable energy systems.

Real-time energy analytics help utilities balance supply and demand and increase renewable output efficiency.

The plan's investment in smart grids and digital infrastructure aligns with Australia's broader net-zero emissions strategies.

Transformation Imperatives:

AI integration requires robust cybersecurity to protect critical infrastructure from emerging digital threats.

Workforce capability building in data analytics and digital engineering will be vital to achieving operational resilience.

Transparent energy pricing models and consumer dashboards powered by AI can improve trust and empower customers in regulated markets.

Practical Recommendations for Industry Leaders

Translating policy into practice requires deliberate action. The following recommendations provide starting points for organisations across Australian industries.

1. Assess Your Current Position

Before embarking on AI initiatives, understand where you stand. What AI tools do your staff already use? What processes might benefit from AI enhancement? What risks does AI present to your operations? What opportunities might you be missing?

The National AI Centre provides free assessment resources that help organisations evaluate their AI readiness. Use these tools to establish a baseline before investing in specific technologies.

2. Invest in AI Literacy

AI literacy is no longer optional for leadership teams. Understanding AI's capabilities, limitations, and implications enables informed decision-making about adoption priorities, risk management, and workforce development.

TAFE AI microcredentials, university short courses, and industry training programs offer accessible pathways for building this understanding. Make AI literacy a professional development priority for your entire leadership team.

3. Start with High-Value, Low-Risk Applications

AI implementation need not begin with transformational projects. Administrative automation, customer service enhancement, and data analysis often provide significant returns with manageable risks.

Build organisational confidence through successful smaller initiatives before tackling complex applications. Document lessons learned to inform subsequent projects.

4. Establish Governance Frameworks

Responsible AI use requires explicit policies addressing transparency, accountability, fairness, and safety. The Digital Transformation Agency's Policy for Responsible Use of AI in Government provides a model that private sector organisations can adapt.

Consider establishing an AI governance committee that includes diverse perspectives, including technology, operations, legal, human resources, and worker representatives. Regular review of AI applications ensures ongoing alignment with organisational values and regulatory requirements.

5. Engage Your Workforce

The plan emphasises meaningful worker consultation about AI deployment. This is not merely a compliance requirement; it is good practice that improves implementation outcomes.

Workers often possess practical insights about how AI tools might enhance their roles. Their concerns about privacy, surveillance, and job security deserve respectful engagement. Building trust through transparent communication enables smoother transitions.

6. Access Available Support

The government has committed substantial resources to supporting AI adoption. The AI Adopt Program, NAIC guidance, R&D Tax Incentive, and National Reconstruction Fund all offer pathways for funded support.

Industry associations, business advisors, and training providers can help navigate these opportunities. Do not attempt to reinvent approaches that others have already refined.

Addressing Common Concerns

AI adoption generates legitimate concerns that deserve thoughtful responses rather than dismissal.

Workforce Displacement

The plan explicitly states that AI reshapes tasks rather than eliminating jobs. Historical evidence supports this perspective. Technology transitions have consistently created more employment than they displaced, though the transition period requires active management.

Proactive investment in workforce development, meaningful consultation, and transition support enable organisations to navigate this change while retaining valuable human capability.

Privacy and Surveillance

AI-enabled monitoring raises legitimate concerns about worker privacy. The plan recognises these issues and emphasises transparency about AI tool deployment in workplaces.

Privacy Act modernisation will strengthen protections for personal information. Organisations should anticipate these requirements and establish practices that respect privacy while enabling beneficial AI use.

Bias and Discrimination

AI systems can perpetuate or amplify biases present in their training data. The plan addresses this through the AI Safety Institute's testing capabilities and responsible AI guidance.

Organisations deploying AI for decisions affecting individuals, including hiring, lending, and service provision, must evaluate systems for discriminatory outcomes and implement appropriate safeguards.

First Nations Considerations

The plan acknowledges specific harms that AI has caused to First Nations communities, including perpetuation of harmful stereotypes and misuse of cultural intellectual property.

Genuine engagement with First Nations communities and alignment with Closing the Gap reforms must inform AI initiatives that affect Indigenous Australians. The Kakadu AI project demonstrates respectful approaches that maintain Traditional Owner governance over cultural knowledge.

Looking Forward: The Plan as Beginning

The Ministers' foreword explicitly frames the National AI Plan as a beginning rather than a conclusion. As technology evolves and confidence grows, refinements and extensions will follow.

This adaptive approach reflects wisdom about the pace of AI development. Rigid frameworks would quickly become obsolete; responsive governance enables continuous improvement.

Industry stakeholders should engage actively with government consultations as the plan evolves. Our collective experience implementing AI in practice generates insights that can inform policy refinement.

Embracing the Opportunity

Australia's National AI Plan 2025 provides a thoughtful framework for navigating one of the most significant technological transitions in human history. Its balanced attention to opportunity, benefit distribution, and safety creates conditions for sustainable success.

For Australian industries, the plan represents both invitation and challenge. The invitation is to participate actively in shaping how AI transforms our economy and society. The challenge is to build the capabilities, governance structures, and partnerships that enable responsible participation.

As someone who has dedicated my career to building Australia's education and training sector, I see extraordinary potential in this moment. AI will not replace the human connections at the heart of quality education. But it will augment our capabilities, extend our reach, and enable outcomes previously beyond imagination.

The organisations that thrive will be those that embrace AI strategically, not chasing every new tool, but thoughtfully integrating technologies that genuinely advance their missions. They will invest in their people, establishing AI literacy as a foundational competency. They will govern AI use responsibly, building trust with stakeholders through transparency and accountability.

Most importantly, they will approach AI with the optimism it deserves. Words have power. The language we use about AI, whether fearful or hopeful, shapes the outcomes we create. The National AI Plan chooses hope, articulating a vision where AI serves Australian prosperity while reflecting Australian values.

Let us embrace that vision together.

About the Author

Sukh Sandhu is the Executive Director of CAQA (Compliance and Quality Assurance) and Editor-in-Chief of the VET Sector Magazine, Australia's premier publication for vocational education and training professionals. With over 30 years of extensive experience in the VET and Higher Education sectors, Sukh has successfully led international colleges as CEO and previously contributed to ASQA as a Business Systems Project Official.

Currently pursuing a doctorate, his research focuses on compliance and auditing challenges faced by Australian RTOs. Recognised for his influence and expertise, Sukh was one of the most visited profiles in 2011 and 2012 and has received over 250 prestigious global honours and awards for his outstanding contributions to education and industry.

> ***Disclaimer:****This article represents the author's independent analysis and does not constitute legal advice. Organisations should seek appropriate professional guidance regarding AI governance requirements applicable to their specific circumstances.*

Resources

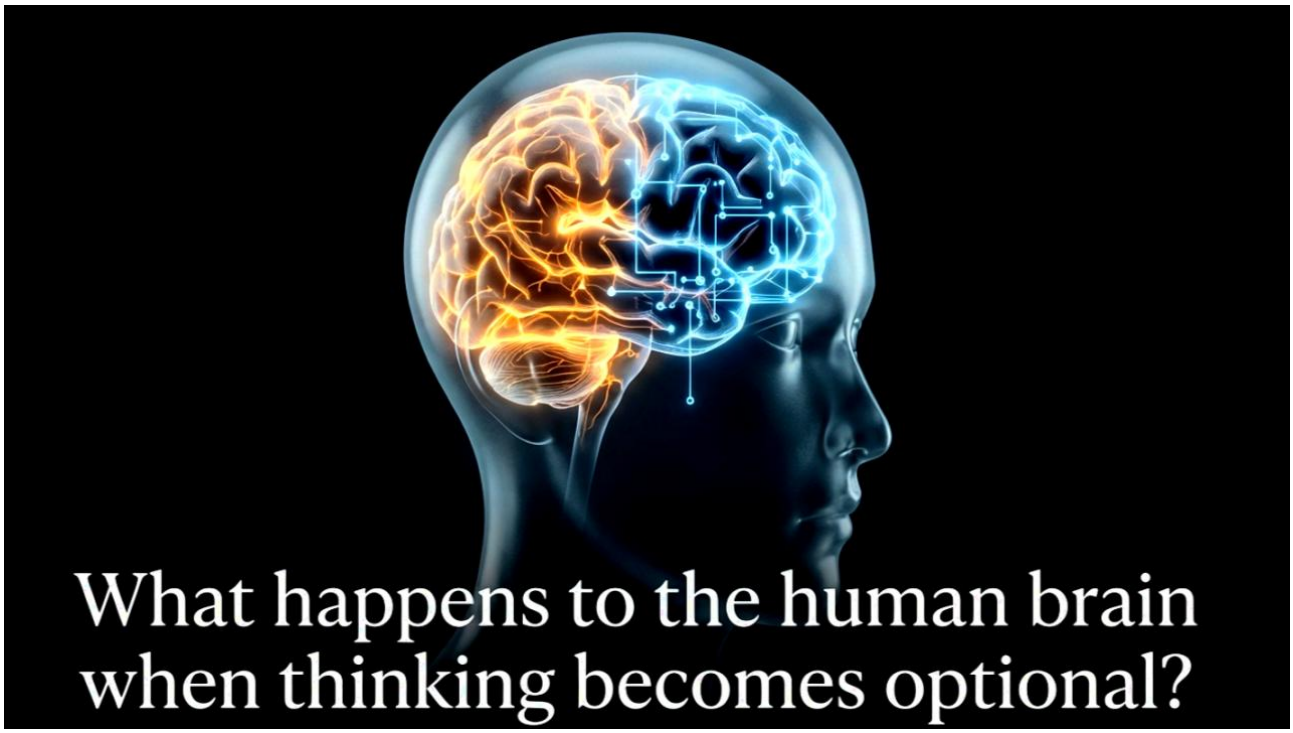
1. National AI Plan 2025: industry.gov.au/ai
2. National AI Centre: naic.gov.au
3. AI Adopt Program: business.gov.au
4. Skills and Training Information: skills.gov.au
5. CAQA Resources: caqa.com.au

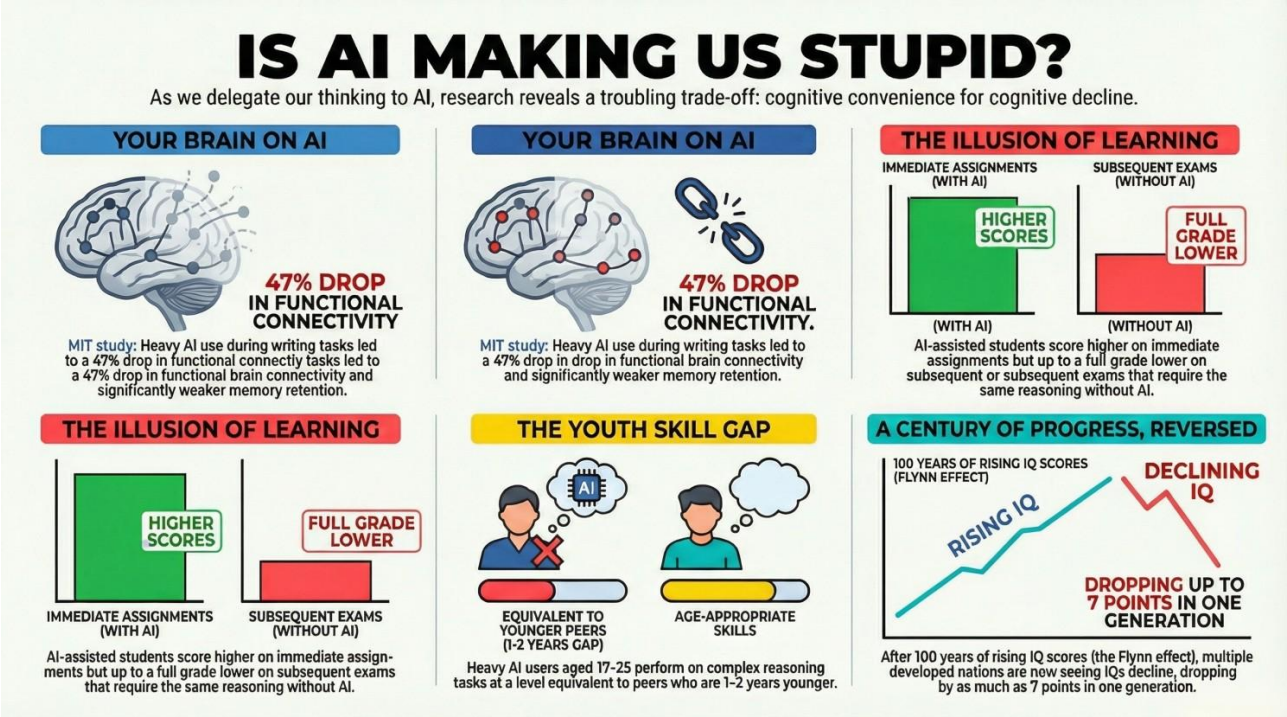
CONNECTING THE DOTS

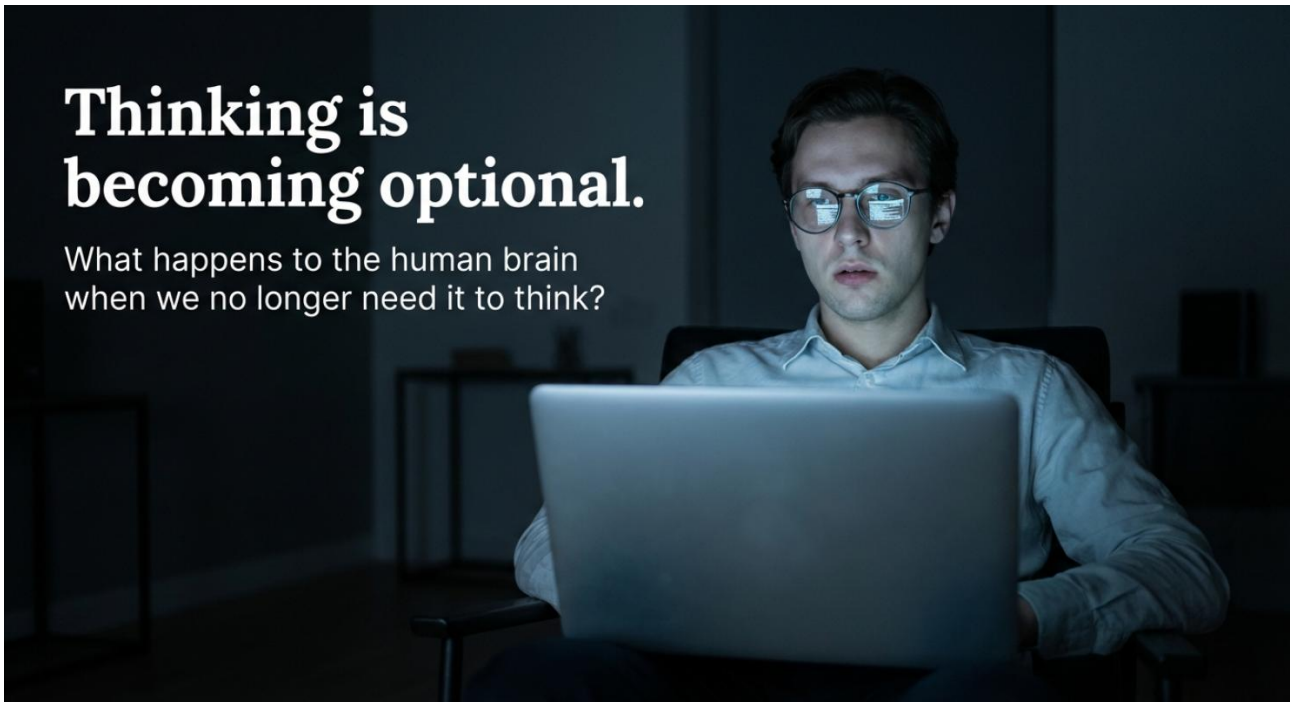
Analysis: Is AI making us stupid? How a generation is losing the ability to think independently

By Dr Sukh Rashpal Sandhu | First published 24 December 2025

There's an uncomfortable question circulating in education circles, boardrooms, and neuroscience labs: What happens to the human brain when thinking becomes optional?

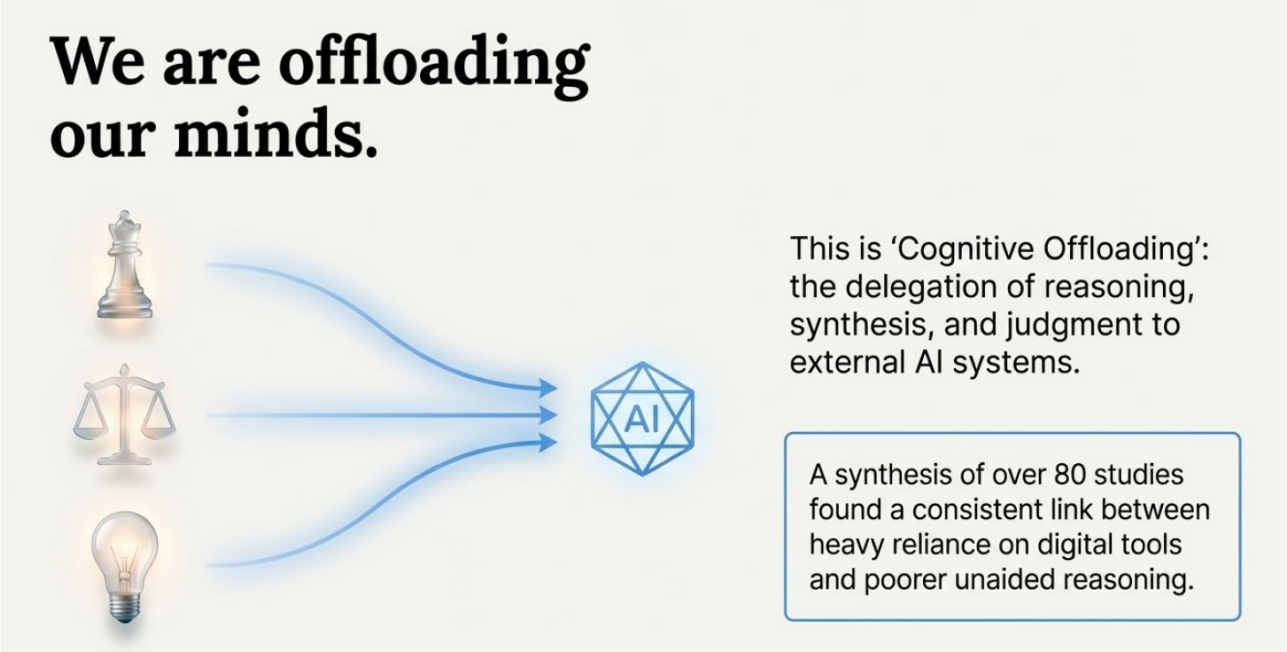






The answer, according to mounting research, is sobering.

Multiple reviews in cognitive science have found consistent links between heavy reliance on external digital tools and poorer performance on later unaided reasoning tasks, particularly in younger adults and frequent tech users.



[![Frontiers | The cognitive paradox of AI in education: between enhancement and erosion
Frontiers | The cognitive paradox of AI in education: between enhancement and erosion
problems in student retention and critical thinking in 206 vocational education students in the Akwa Ibom State, Nigeria. The results showed that AI posed si...
Frontiers](https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2025.1550621/full)

Recent reviews in cognitive science describe AI as a "cognitive offloading accelerator," warning that habitual delegation of analysis and judgment can lead to measurable declines in independent problem-solving over time. Importantly, research also shows that structured AI use, where learners critique, verify, or reflect on AI outputs, can enhance critical thinking; the risk arises primarily from passive, uncritical reliance.

[![AI tools may weaken critical thinking skills by encouraging cognitive offloading, study suggests
AI tools may weaken critical thinking skills by encouraging cognitive offloading, study suggests
Heavy AI tool usage is associated with weaker critical thinking, a study finds, as users increasingly offload cognitive tasks instead of engaging in independent analysis.
PsyPost, Psychology News](https://www.psypost.org/ai-tools-may-weaken-critical-thinking-skills-by-encouraging-cognitive-offloading-study-suggests/)

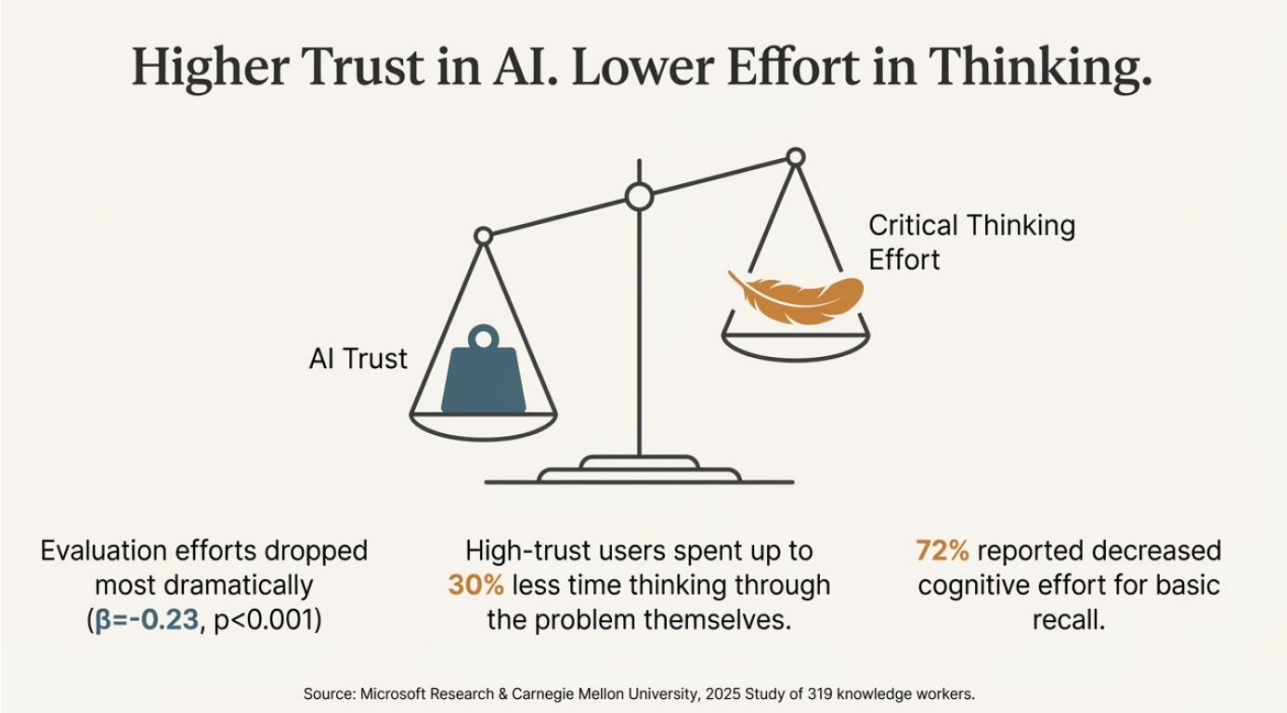
The Data We Can't Ignore

> *A 2025 study affiliated with Microsoft Research examined knowledge workers who use AI tools regularly. Analysing real-world examples of AI-assisted work, they found something troubling: higher confidence in AI capabilities was associated with reduced critical thinking effort.*

In parallel, a 2025 study published in the journal Societies examined 666 participants across age groups and found a significant negative correlation between frequent AI tool usage and critical thinking scores, with cognitive offloading mediating this relationship.

AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking
The proliferation of artificial intelligence (AI) tools has transformed numerous aspects of daily life, yet its impact on critical thinking remains underexplore
papers.ssrn.com

> *Both workers and students who relied heavily on AI engaged in less evaluation, analysis, and synthesis. The effects were observable across multiple cognitive activities measured using Bloom's taxonomy, with cognitive offloading behaviours mediating the relationship between AI use and lower critical thinking scores..*



Specifically:

- Evaluation efforts dropped most dramatically
- Analysis declined significantly
- Comprehension, synthesis, and basic knowledge activities all showed measurable reductions

In practical terms: most participants reported decreased cognitive effort for basic recall and comprehension tasks when using AI. Time-on-task logs showed that high-trust users spent substantially less time thinking through problems themselves and significantly more time crafting prompts, skimming AI outputs, and lightly editing them into final products. Even for higher-order evaluation tasks, where you'd expect humans to remain firmly in control, over half still reported reduced effort.

Neuroscience reviews now describe this pattern as a “cognitive engagement collapse,” where efficiency tools reduce the frequency and depth of the very mental operations that build durable expertise.

The cognitive paradox of AI in education: between enhancement and erosion
Artificial intelligence (AI) is rapidly transforming learning through unparalleled levels of personalization, efficiency, and scalability (Govea et al. 2023 ; Mahmoud and Sørensen, 2024). AI-based technologies such as adaptive learning systems and...
pmc.ncbi.nlm.nih.gov

The Young Are Most Vulnerable

The Societies study findings were stark: examined 666 participants across diverse age groups and educational backgrounds. The findings were stark:

Heavy AI users in the 17 to 25 bracket showed significantly lower critical thinking scores compared to light users in the same age group, with the gap particularly pronounced for complex reasoning tasks.

Younger participants (ages 17-25) showed significantly higher dependence on AI tools and scored lower on critical thinking assessments compared to older age groups.

- Across this and similar studies, younger participants (17 to 25) not only relied more heavily on AI, they also reported using it for everyday memory, decision-making, and even moral judgments, far more than older adults. Diary studies show that many young adults now consult AI for decisions like which courses to take, how to respond to social messages, and even whether to end relationships, essentially outsourcing value-laden thinking.

This wasn't about intelligence; it was about practice. Or rather, the lack of it.

The phenomenon researchers are documenting is called cognitive offloading: the delegation of mental tasks to external systems. While some offloading is natural (we've always used tools), AI represents a qualitative leap. Unlike a calculator that performs specific computations, generative AI can now handle reasoning, synthesis, and judgment, the very activities that build cognitive capacity.

Large majorities of students in recent surveys expressed concern that excessive AI reliance could weaken their ability to think independently and make decisions. They know. They just can't stop.

Strikingly, substantial proportions of students in surveys admitted they had already used AI to avoid doing reading or problem-solving they could have done themselves, even while believing this was bad for their learning.

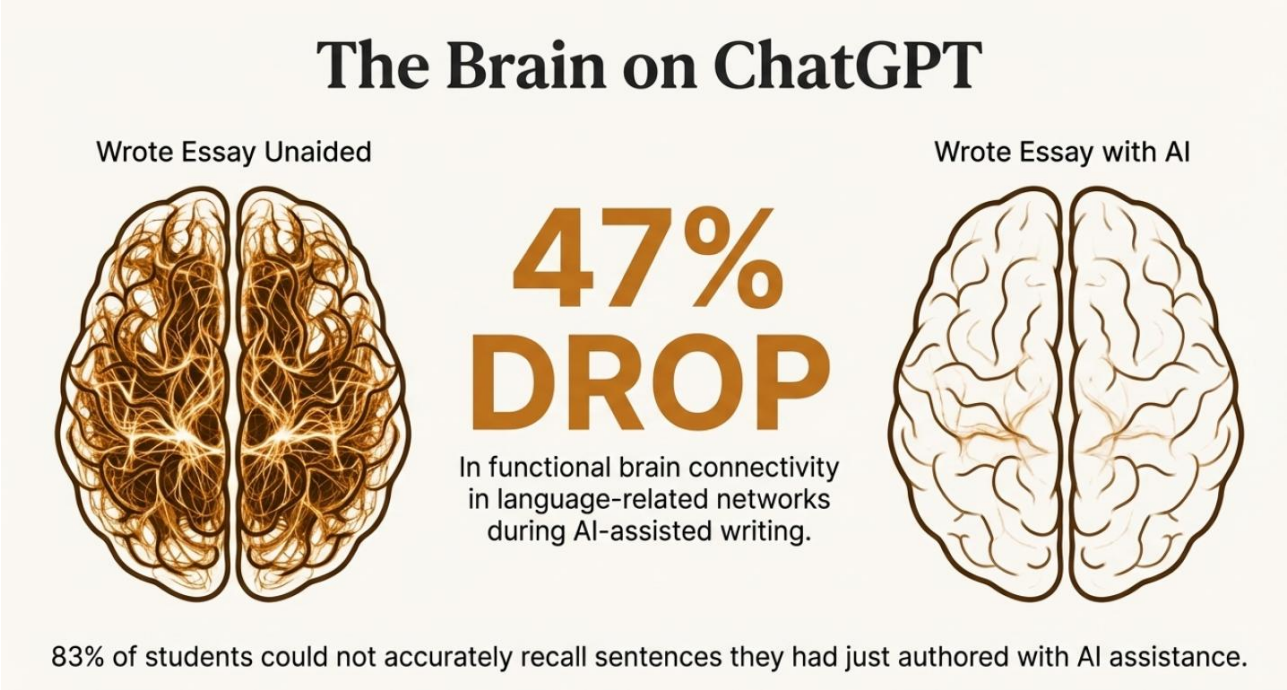
In UK surveys, majorities of pupils reported concerns that AI was harming their skill development and encouraged copying rather than original thinking.

The Brain on ChatGPT

MIT Media Lab's 2025 study wired 54 students with EEG headsets to measure brain activity during essay writing.

The tasks were not trivial: students had to produce 800 to 1,000-word essays under time pressure on unfamiliar topics, mimicking real high-stakes coursework conditions.

Participants were divided into three groups: those using ChatGPT, those using search engines, and those using no tools at all.



Key findings included significantly reduced functional connectivity in language-related networks and weaker engagement in attention regions. Notably, the majority of ChatGPT users failed to correctly recall sentences they had just "authored" with AI assistance, compared with far better recall in the unaided group.

![Using ChatGPT to write essays may harm critical thinking, MIT study finds
Using ChatGPT to write essays may harm critical thinking, MIT study finds The 'Your Brain on ChatGPT' study will make you consider the consequences. Mashable](https://mashable.com/article/chatgpt-writing-study-mit-cognitive-cost)

The results were remarkable:

Students who wrote without AI showed stable or increased brain connectivity. They could recall their main arguments and key sentences a few minutes later.

In contrast, students who only used AI had weaker brain connectivity and lower memory retention. They also felt less ownership of their work. Their performance on new writing tasks without AI was significantly worse than the non-AI group. When these students later completed tasks without AI, the patterns of reduced engagement persisted. The researchers observed weaker cognitive engagement and poorer transfer to new tasks, suggesting that habitual AI use may affect subsequent independent performance.

This supports broader evidence about heavy, uncritical AI use in learning. It is associated with lower motivation, less effort, and weaker long-term understanding.

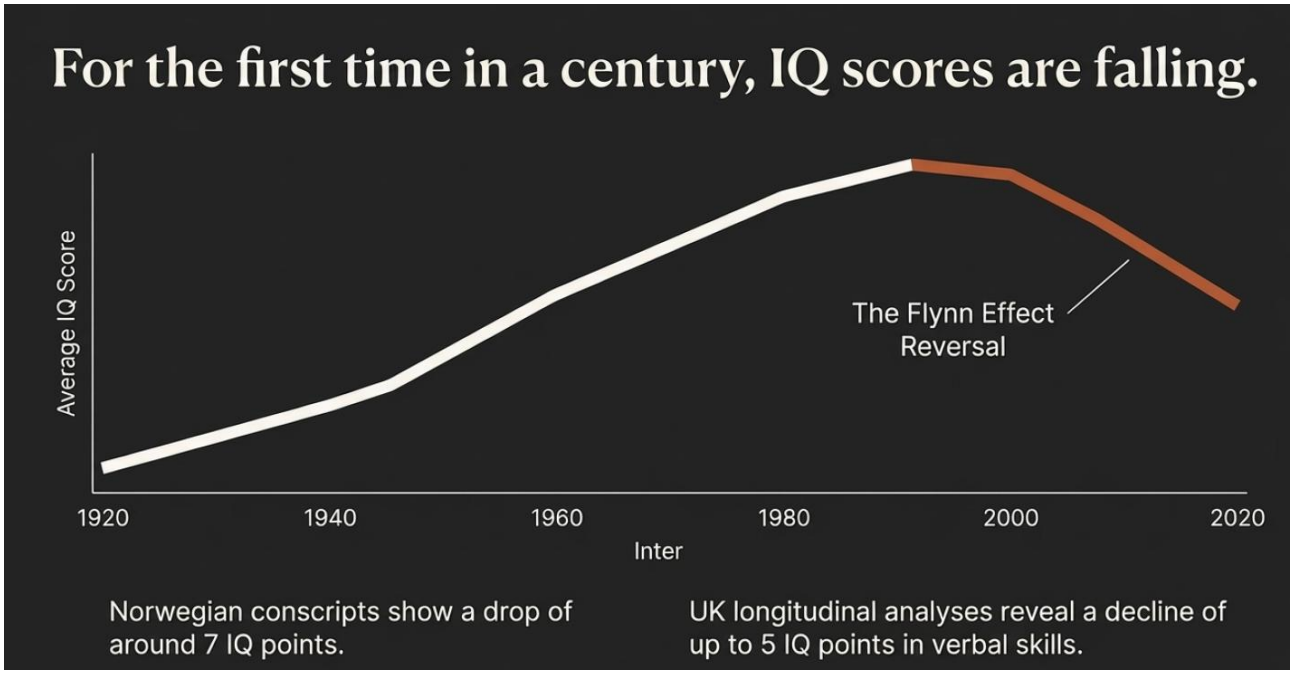
How AI quietly undermines the joy and effort of learning: a call for rebalancing education in the digital age The integration of artificial intelligence (AI) into the learning process has undoubtedly reshaped education, providing learners with access to resources, tailored experiences, and efficient tools to tackle problems. However, an important question... pmc.ncbi.nlm.nih.gov

Randomised controlled trials confirm this pattern. Students perform better when using generative AI tools, but perform worse when those tools are taken away. In one reported experiment, AI-assisted students scored higher on immediate assignments but notably lower on subsequent closed-book tests that required similar reasoning without AI. They've bypassed the cognitive processes essential for building mental capacity. The tool solved the problem; the brain didn't learn to.

The Flynn Effect Reverses

For most of the 20th century, IQ scores rose steadily, approximately 3 points per decade. Psychologist James Flynn documented this phenomenon so thoroughly that it bears his name. Over a century, that 3-points-per-decade rise translated into roughly 30 IQ points, enough to shift the entire bell curve of human problem-solving upwards.

That trend has reversed.



Large-scale cohort studies in Norway, Denmark, Finland, the UK, France, and the United States now document stagnation or decline in IQ scores, particularly in verbal reasoning and complex problem-solving, among cohorts born after the mid-1970s. In Norway alone, analyses of over 730,000 conscripts show measurable IQ declines between the mid-1970s cohorts and those born in subsequent decades.

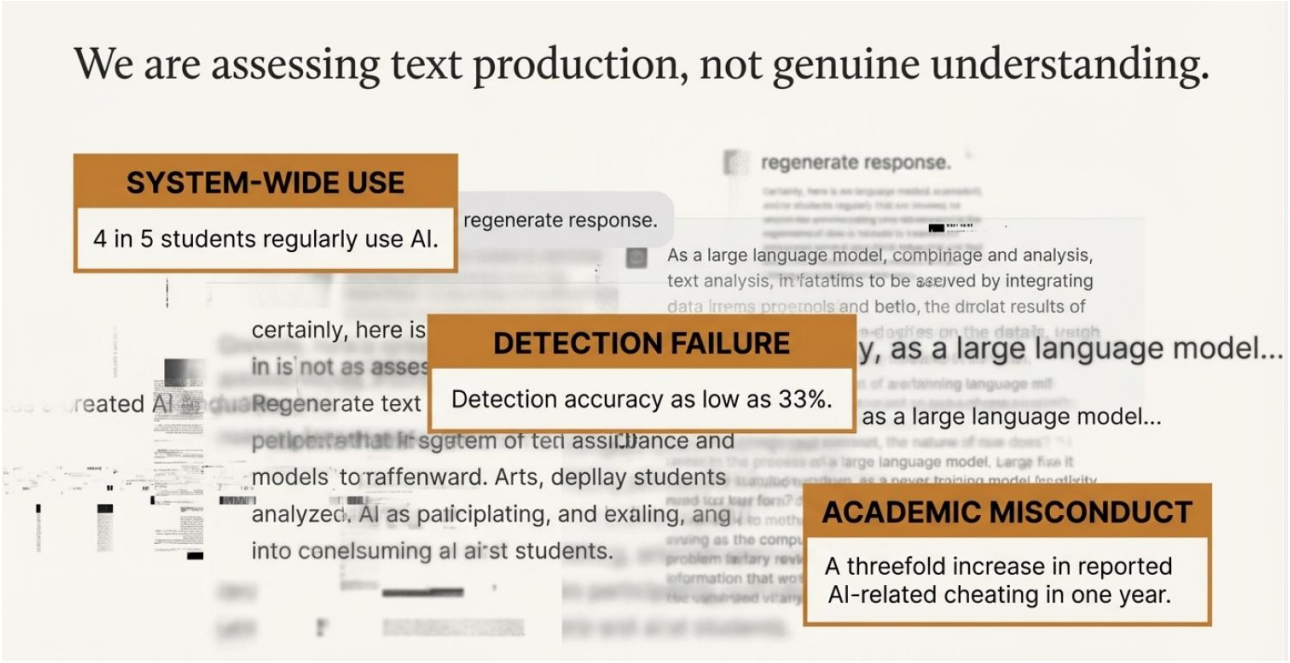
The cognitive paradox of AI in education: between enhancement and erosion Artificial intelligence (AI) is rapidly transforming learning through unparalleled levels of personalization, efficiency, and scalability (Govea et al. 2023 ; Mahmoud and Sørensen, 2024). AI-based technologies such as adaptive learning systems and... pmc.ncbi.nlm.nih.gov

Norwegian researchers Bratsberg and Rogeberg documented declining IQ scores among men born after 1975, and critically, the decline appeared *within families*, eliminating genetic explanations. Similar reversals have been observed in Finland, Denmark, the UK, France, and the United States.

Recent cohort research has found something revealing: while some specific cognitive scores continued to rise, there is evidence that the "positive manifold" of intelligence, the pattern where performance in one cognitive area predicts performance in others, may be weakening. Some researchers suggest people are becoming cognitively specialised rather than generally capable.

In the UK, longitudinal analyses suggest declines in verbal skills and structured reasoning among post-1970 cohorts. Studies of American adults have shown declines in vocabulary, reading comprehension, and mathematical reasoning in recent decades.

Researchers have proposed several contributing factors, including educational shifts away from memorisation and retrieval practice, and the rise of cognitive offloading to digital tools, though multiple factors are likely at play, and causation remains debated.



Some researchers warn that if these trends continue, future cohorts may display narrower islands of expertise but weaker general reasoning, a pattern some employers already report.

Some analyses suggest that reduced retrieval practice, combined with increased digital cognitive offloading, may contribute to observed declines in general cognitive performance, though establishing direct causation remains challenging.

How AI quietly undermines the joy and effort of learning: a call for rebalancing education in the digital age The integration of artificial intelligence (AI) into the learning process has undoubtedly reshaped education, providing learners with access to resources, tailored experiences, and efficient tools to tackle problems. However, an important question... pmc.ncbi.nlm.nih.gov

The Assessment Crisis

Here's where the classroom reality becomes stark:

Surveys suggest that over half of college students have used AI to complete assignments or exams. Global survey work suggests this may be conservative, with multiple reports

finding that large majorities of students regularly use AI in their studies for homework or revision. In some disciplines, such as business and computer science, reported usage rates are particularly high, making non-AI-assisted work increasingly the exception rather than the rule.

Large proportions of students now use AI tools for coursework, with smaller but significant numbers using them to answer assessments or generate substantial portions of submitted work

UK universities have reported substantial increases in AI-related academic misconduct, with some institutions seeing multi-fold increases year-on-year

Student discipline rates for AI-related misconduct have risen sharply in recent academic years

Yet detection remains inadequate. AI detection tools show variable accuracy, with substantial false negative rates. Students have become sophisticated at combining tools, editing outputs, and avoiding detection. In trials, paraphrasing AI-generated text through other tools has proven highly effective at evading detection, illustrating how easily determined students can defeat current systems.



A 2025 analysis of over 10,000 samples found that popular detectors produced substantial false positive rates for human-written work, with non-native English writers disproportionately misclassified, in some cases at rates of 30% or higher.

Even vendors that advertise "99% accuracy" acknowledge that these figures depend heavily on test conditions and tolerate non-trivial false positive and negative rates in practice.

As one researcher noted: "AI hasn't made more students cheat, it's changed how they do it." The baseline was already concerning. AI just gave misconduct a polished interface.

The Viva Voce Solution

For 2,000 years, oral examinations, *viva voce*, Latin for "living voice", were the standard for academic assessment. Medieval scholars defended their knowledge through public debate. The tradition persists in PhD examinations, Oxbridge tutorials, and European postgraduate programmes.

Now educators are rediscovering why.

You cannot outsource your reasoning during a live conversation. When a professor asks follow-up questions, probes your logic, and challenges your assumptions in real-time, no AI can intervene. The written submission becomes merely a foundation; the viva exposes actual comprehension.

Some institutions have piloted replacing written exams with structured oral assessments in select subjects. Early reports suggest not only fewer suspected misconduct cases but also richer insight into students' understanding, as live questioning reveals where comprehension breaks down, though peer-reviewed outcome data remains limited.

Pilot programmes suggest that even low-intensity sampled vivas, with spot-checking of a small percentage of submissions, may substantially deter contract cheating and undisclosed AI use. Early case reports indicate meaningful reductions in misconduct referrals, though robust peer-reviewed evaluations are still emerging.

Early reports indicate drops in suspected academic misconduct and generally positive student perceptions of fairness, despite initial anxiety about oral grading.

As researchers from UCL and Cambridge argue: "We must ask ourselves whether we are assessing text production or genuine understanding."

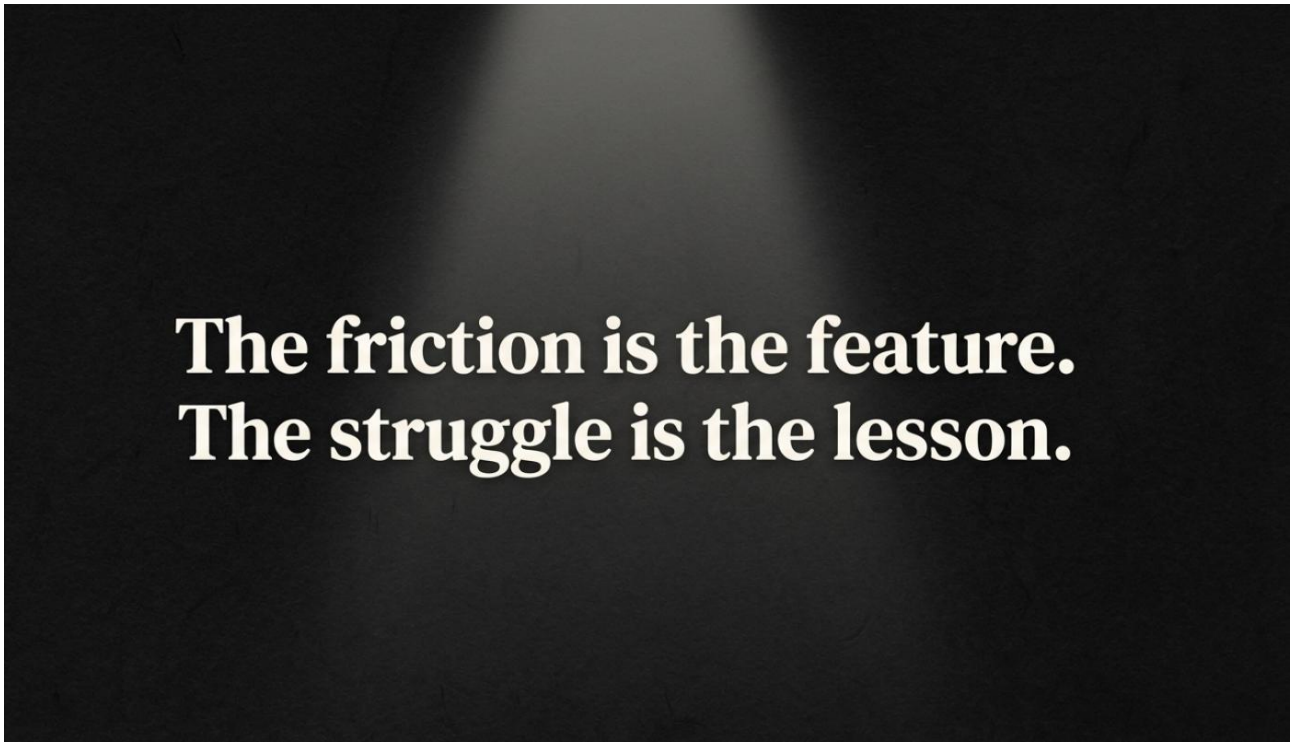
The Real Irony

The Microsoft study documented what they call "The Ironies of Generative AI", a callback to Bainbridge's classic observation about automation:

By mechanising routine tasks and leaving exception-handling to humans, automation deprives users of routine opportunities to practise judgment and strengthen cognitive faculties.

In other words, the very efficiency that makes AI attractive is what makes it cognitively corrosive. We improve through practice. Remove practice, and capacity atrophies.

There's a cruel symmetry here. AI promises to free us for "higher-order thinking." But higher-order thinking is built on the foundation of lower-order activities, retrieval, comprehension, application. Skip the foundation, and the higher floors have nothing to stand on.



What Must Change

The path forward isn't banning AI, that battle is already lost. It's redesigning how we learn, work, and assess understanding:

For educators:

Shift from output-based to reasoning-based assessment: Research on "explanation-first" assessment designs shows higher gains in transfer and critical thinking compared to traditional problem-set grading, especially when students must justify or critique AI-generated solutions.

Implement oral examination components, even if sampled

- Design assignments that require AI tools but evaluate the human reasoning applied to AI outputs
- Courses that embed structured reflection on 'how I used AI and why I trusted or rejected its outputs' report stronger gains in critical thinking than courses that either ban AI or allow it without reflection.
- Teach metacognition explicitly: how to verify, critique, and build upon AI-generated content: Metacognitive training, teaching students to plan, monitor, and evaluate their thinking, has one of the highest effect sizes in the learning sciences and is a key protective factor against uncritical AI reliance.

For organisations:

Recognise that AI-boosted productivity may come with hidden cognitive costs: Field experiments with knowledge workers show that AI assistance can improve short-term task speed while degrading independent performance on similar tasks days later, especially when workers skip manual practice.

In workplace experiments, teams that alternated between AI-assisted and ‘no-AI’ practice sessions retained more independent problem-solving ability than teams that used AI continuously on all tasks.

Create intentional "no-AI" spaces for skill development

Train workers to maintain critical thinking despite (not because of) AI assistance

For individuals:

Delay offloading: struggle with problems before seeking AI help

Practise retrieval: actively recall information rather than re-looking it up

Maintain cognitive friction: the discomfort of thinking hard is the feeling of learning

The Stakes

Barbara Oakley's 2025 research on the "Memory Paradox" synthesises decades of cognitive science into a simple insight: when we repeatedly retrieve information, our brains form durable memory structures and neural patterns essential for intuitive reasoning and error-checking.

When we default to "just ask AI," those processes never fully develop.

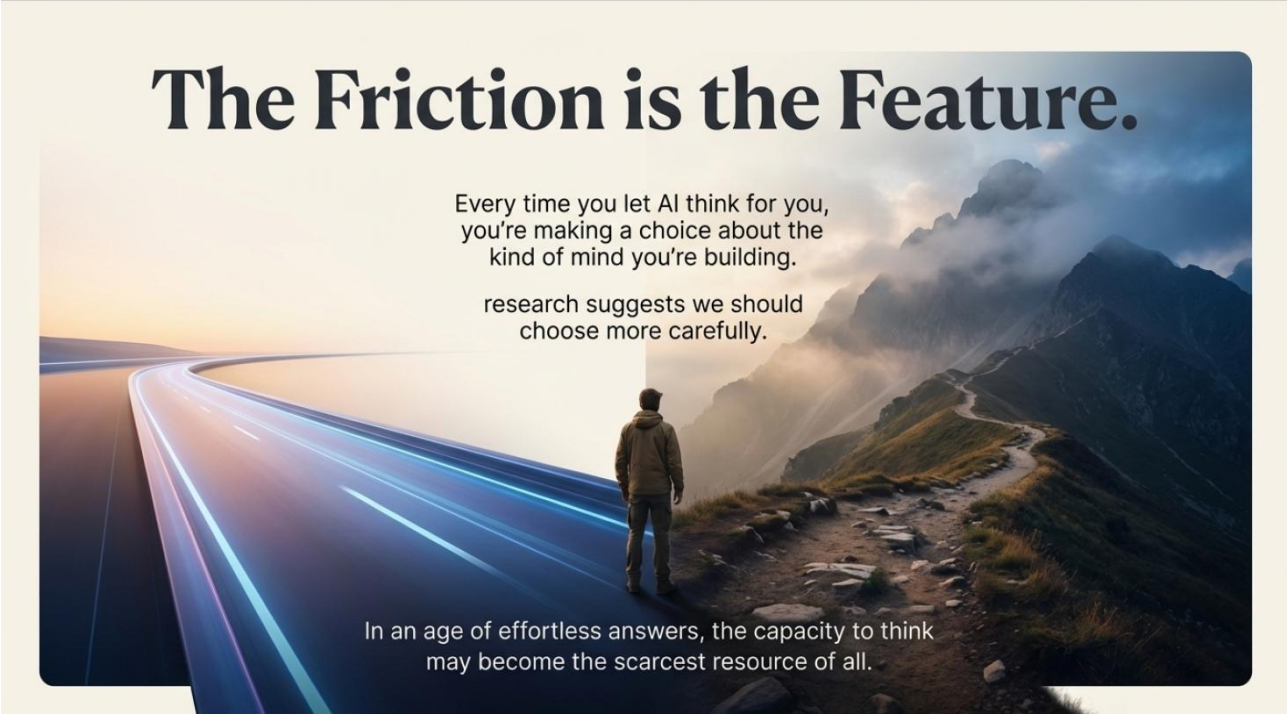
These advantages appear to persist over time: students trained with high retrieval and low offloading not only remember more facts, they tend to be better at spotting flawed arguments and misleading information in new contexts.

Research suggests that learners who regularly engage in retrieval practice and effortful problem-solving develop stronger foundations for error detection and flexible reasoning than peers who rely heavily on external aids.

The question isn't whether AI is good or bad for productivity. It's whether we're willing to trade cognitive capacity for cognitive convenience, and whether we even notice the exchange as it happens.

Every time you let AI think for you, you're making a choice about the kind of mind you're building. The research suggests we should choose more carefully.

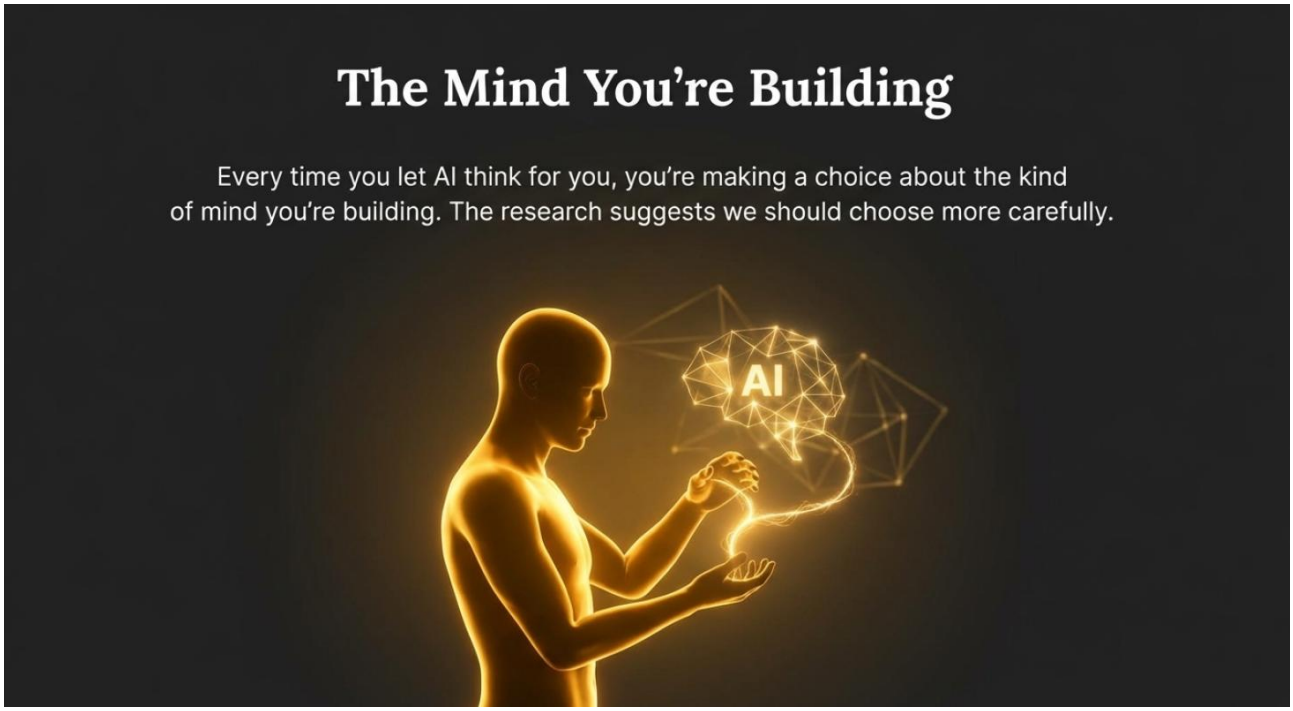
The friction is the feature. The struggle is the lesson. And in an age of effortless answers, the capacity to think may become the scarcest resource of all.



> ***What's your experience? Have you noticed changes in your own thinking patterns since adopting AI tools? I'd love to hear your perspective in the comments.***

Surveys suggest many people have noticed such changes, with substantial numbers of students and professionals reporting that AI makes it "too easy" to get answers and reduces their inclination to struggle with difficult tasks.

In recent surveys, substantial proportions of respondents said AI tools make them less patient with difficult material, with many admitting they are now more likely to give up quickly and consult AI when stuck.



[![Increased AI use linked to eroding critical thinking skills
Increased AI use linked to eroding critical thinking skills A study by Michael Gerlich at SBS Swiss Business School has found that increased reliance on artificial intelligence (AI) tools is linked to diminished critical thinking abilities. It points to cognitive offloading as a primary driver of the decline.
Phys.org](https://phys.org/news/2025-01-ai-linked-eroding-critical-skills.html)

Image Credits:

All images included in this article have been created using NotebookLM, Fotor and Gemini.

#ArtificialIntelligence #Education #CriticalThinking #FutureOfWork #CognitiveScience #VocationalEducation #HigherEducation #LearningAndDevelopment #Leadership

STAY CONNECTED

Follow us on LinkedIn

Enjoyed this edition? Each of these newsletters publishes new pieces on LinkedIn. Subscribe, like and share to read them first, and join the conversation in the comments.



Connecting the Dots You are reading this newsletter

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/connecting-the-dots-7397392757360648192/



Letters of the Heart

Real stories, timeless wisdom, and soul-deep inspiration.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/letters-of-the-heart-7329273590376210432/



The Hope Dispatch

Uplifting news, inspiring stories, and positive change from around the world.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-hope-dispatch-7338912777081638914/



The HealthTech Herald

The latest breakthroughs at the intersection of health and technology.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-healthtech-herald-7338913774701027328/



Truth, Translated

Plain language for a complicated world: education, geopolitics, political accountability and Big Tech.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/truth-translated-7450724196525035520/

Follow the author: Dr Sukh Rashpal Sandhu on LinkedIn

BACK PAGE

About Connecting the Dots

Connecting the Dots is a LinkedIn newsletter written by Dr Sukh Rashpal Sandhu. Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Read and subscribe

New editions continue to be published on LinkedIn. Search for Connecting the Dots on LinkedIn to subscribe.

Disclaimer

The views in this collection are the author's own. Pieces are reproduced as first published, and facts, figures and links are current at their original publication date. Nothing in this collection is legal, financial, medical, migration or regulatory advice.

Copyright

Copyright Dr Sukh Rashpal Sandhu 2026. All rights reserved. Images are reproduced as they appeared in the original posts.

Connecting the Dots | Dr Sukh Rashpal Sandhu | sukh@sukhsandhu.com

Masthead

Author: Dr Sukh Rashpal Sandhu

Publisher: Dr Sukh Rashpal Sandhu

Edition: 3, September 2026

Contact: sukh@sukhsandhu.com

Connecting the Dots is published on LinkedIn by Dr Sukh Rashpal Sandhu. This edition collects the pieces in the order they were first published, with their original images.