Frugal AI Hub

Research Report · May 2026

LEADERSHIP IN THE AGE OF AI

Frugal AI: Practical Principles for Leaders

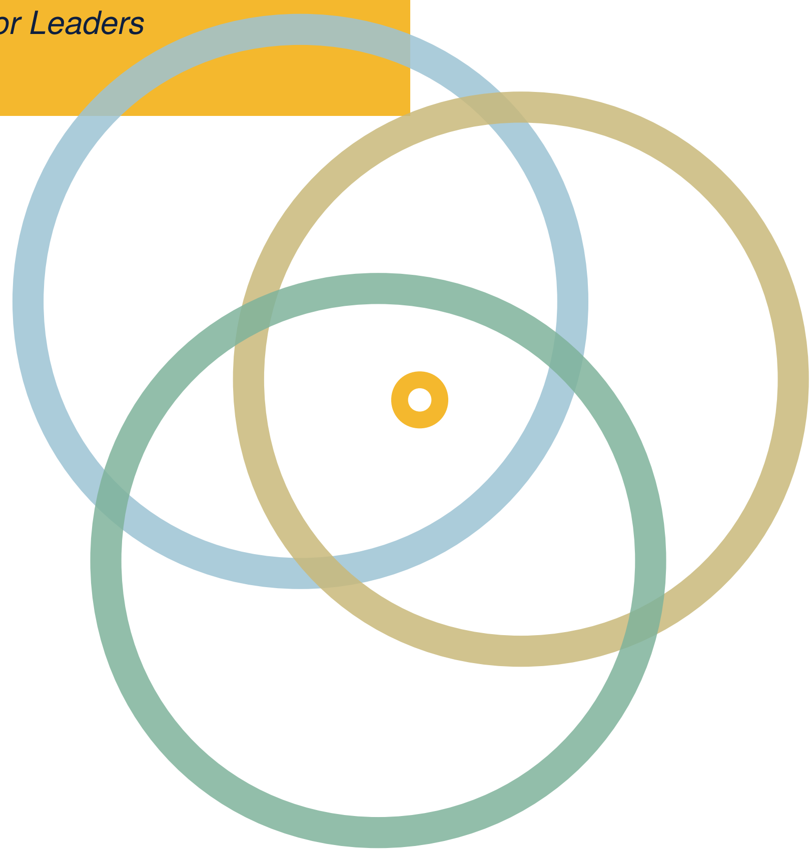

Three arcs of leadership: foundations, governance and innovation, strategy and boardroom, meeting at the point of judgement.

FRUGAL AI HUB at UNIVERSITY OF CAMBRIDGE Judge Business School

in collaboration with

global education lab

Frugal AI and the Executive Agenda

Authors

Serish Venkata Gandikota, Co-founder, Frugal AI Hub; Visiting Fellow, Cambridge Judge Business School

Elizabeth Osta, Co-founder, Frugal AI Hub; Visiting Fellow, Cambridge Judge Business School

Arjuna Sathiaseelan, Technology Lead, Frugal AI Hub; Visiting Fellow, Cambridge Judge Business School

Suyash Bhatt, Founder and CEO, Global Education Lab; Visiting Fellow, Cambridge Judge Business School

Jaideep Prabhu, Professor of Marketing and Nehru Professor of Indian Business, Cambridge Judge Business School

About the Frugal AI Hub

The Frugal AI Hub at Cambridge Judge Business School is a research and practice hub advancing the principles, methods, and adoption of Frugal AI: intelligent systems that are high performing yet resource efficient, and that can be understood and governed by the institutions that deploy them. The Hub works with industry, public sector, and civil society partners across fintech, healthcare, education, and sustainability.

About Global Education Lab

Global Education Lab is an innovation-driven knowledge transformation company shaping the next generation of global thinkers, leaders, and change-makers. GEL convenes executive learning experiences that combine academic depth with real-world application, working at the intersection of business education, entrepreneurial practice, and social impact.

Acknowledgements

This paper draws on the Global India Leadership Programme 2026, convened by Global Education Lab at Cambridge Judge Business School from 9 to 13 March 2026. The authors thank the 22 delegates who completed the pre-programme survey and contributed their reflections during the week. We acknowledge the Centre for India and Global Business (CIGB) at Cambridge Judge Business School for its support of research on Indian and global enterprise. We also thank Teck Ming Tan for editorial work on the companion Global Education Lab discussion paper, *Leadership in the Age of AI: Strategy, Innovation, and Decision-Making in a Transforming Economy*, which is referenced throughout.

Abstract

The Global India Leadership Programme 2026 brought together 22 senior leaders across nine industry sectors and four continents. The cohort skews toward founders, C-suite executives, and heads of function, and almost half lead organisations of fewer than 50 employees. They share a common situation: the question of whether to adopt AI has been settled, and the question of how to adopt it well has not.

A pre-programme survey of the cohort found that 21 of 22 leaders view AI as either highly or moderately transformative for their organisations. The same 22 leaders identify a tight cluster of organisational barriers as the binding constraint on adoption: a lack of internal AI skills (cited by 13 respondents), difficulty integrating AI into existing processes (10), data quality issues (7), and unclear return on investment (7). Eighteen of the 22 cite at least one of these four organisational barriers. Half the cohort, when asked how their organisation intends to adopt AI, explicitly select "frugal, modular AI solutions tailored to our context" over enterprise platforms, vendor-led approaches, or in-house development. None select large-scale enterprise AI platforms.

In short, many organisations do not yet have the institutional capacity required to absorb AI effectively, and the cohort already recognises frugality as the response. This paper applies the Frugal AI Hub framework to the cohort data and to the broader pattern visible across executive education.

First, for leaders it is important to set the level of aspiration and destination for AI use: from marginal task support through infrastructure-level design. This needs to address the level of investment required and the timing to reach the desired value.

Secondly, it is important to set the leadership intent on the use of AI. Frugal AI provides a framework based on:

(1) resource efficiency
(2) sustainability
(3) accessibility and inclusion
(4) impact and scalability

Finally, the paper touches on the need for leadership to govern AI — through measurement and constant optimisation of the AI portfolio.

1. Introduction

The Global India Leadership Programme 2026 took place at Cambridge Judge Business School in March 2026. Twenty-two senior leaders from nine sectors and four continents spent a week examining what it means to lead under conditions of rapid technological change, geopolitical fragmentation, and shifting stakeholder expectations. Cambridge faculty delivered the sessions. The Frugal AI Hub contributed a working session on how organisations can absorb AI under real-world constraints. Global Education Lab convened the programme and produced a companion discussion paper, *Leadership in the Age of AI*, which captures the broader leadership themes.

This paper sits alongside that discussion paper and addresses a narrower question. The Global India Leadership Programme cohort, like most executive cohorts the Hub now encounters, has moved past the question of whether AI matters. The harder question is how to lead the AI transformation.

The shift in the AI question

For most of the past five years, the dominant question for executives has been one of capability: which AI systems are available, how powerful they are, and how quickly the frontier is moving. That question is increasingly settled at the level executives care about. Foundation models are available, performant, and commoditising. Open weights, hosted APIs, and specialised vendors mean that access to capability is no longer the constraint for any organisation that is willing to pay for it.

The first step for a leader is to set the aspiration for AI: what is the job that AI should do for my organisation?

The four levels of AI 'destinations'

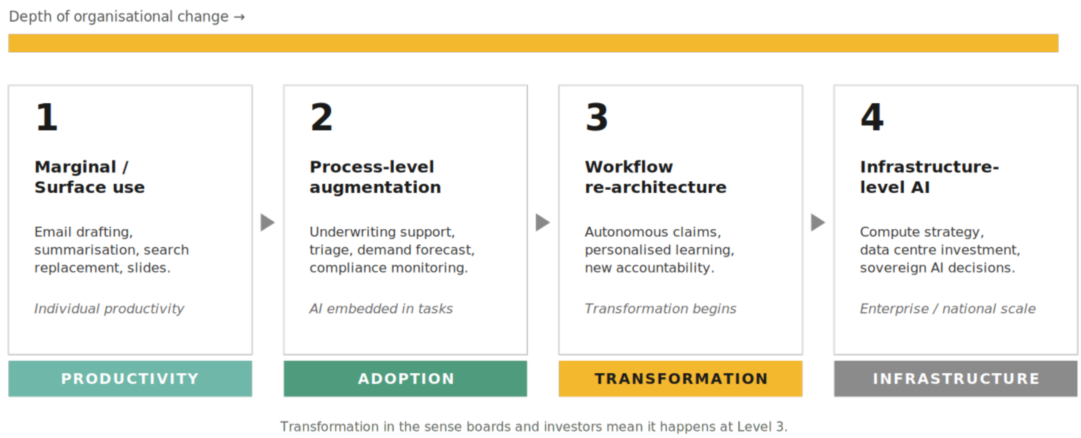

Figure 1. The four levels of AI use, from marginal individual productivity to infrastructure-level architectural decisions.

The four pillars describe what AI should achieve and the depth of organisational change.

In the figure above, Level 1, marginal or surface use, covers individual productivity applications: email drafting, summarisation, search replacement, presentation creation. Adoption is fast. Impact per user is real but limited. Even in surface use, real usefulness should be checked against productivity gains and employee satisfaction in using the tools.

Level 2, process-level augmentation, embeds AI into specific tasks within existing workflows: underwriting support, claims triage, demand forecasting, compliance monitoring. Adoption is moderate. Governance considerations become material. The workflow itself is unchanged.

Level 3, workflow re-architecture, redesigns the end-to-end process around AI capabilities: autonomous claims handling, personalised learning, role redefinition, new accountability structures. This is where measurable productivity transformation begins, and it requires organisational change at the operating-model level, with new governance structures and human-in-the-loop capabilities.

Level 4, infrastructure-level AI, involves compute strategy, clear architectures, and often new business models. It may be a complete redesign of how the company operates, becoming AI native. This is a fundamental shift, with consequences also on human capital and skillsets required.

Once the destination is set, implementation starts on how to get there. As businesses embark on this transformation, this paper advocates considering a Frugal AI lens, as it allows easier scaling if the design of AI has already taken into consideration the constraints in the environment.

2. The Frugal AI Lens

The Frugal AI Hub was established at Cambridge Judge Business School to advance the principles, methods, and adoption of intelligent systems that are high performing yet resource efficient, and that can be understood and governed by the institutions that deploy them. The Hub has produced four prior white papers establishing the framework's definition, technical and socio-economic foundations, and measurement methodology. Frugal AI is an intentional design and deployment approach for building AI systems that are high performing yet lightweight, cost-effective and energy-efficient, and accessible and scalable under resource constraints.

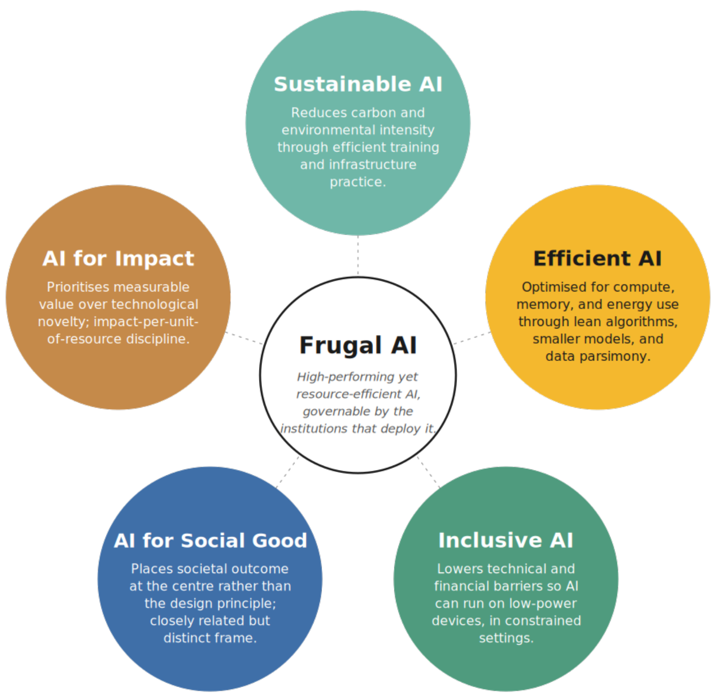

Figure 2. The Frugal AI framework. Frugal AI sits at the centre, surrounded by five adjacent frames: sustainable, efficient, inclusive, impact-focused AI, and AI for Social Good as a closely related but distinct lens.

The constraints Frugal AI solves: economic, environmental, adoption

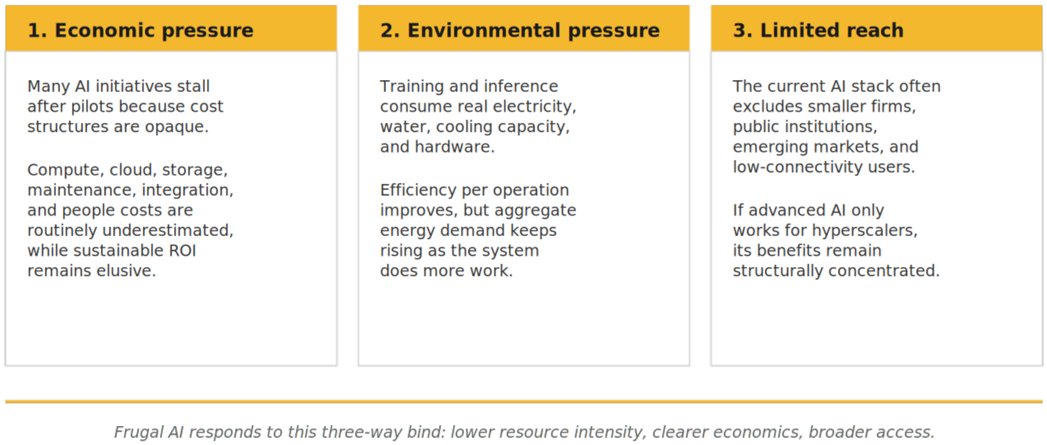

Figure 3. The Triple Challenge: economic, environmental, and reach pressures that Frugal AI is designed to address.

Many AI initiatives never get beyond pilots because the cost structure is opaque. Compute, cloud, storage, maintenance, integration, governance and people costs are routinely underestimated at the project approval stage, with the result that sustainable return on investment remains elusive. The cost denominator that organisations use to justify AI is usually narrower than the cost denominator they end up paying.

The second pressure is environmental. Training and inference consume real electricity, water, cooling capacity, and hardware. Per-operation efficiency continues to improve, but aggregate energy demand rises because the system is being asked to do more work all the time. This is the Jevons paradox playing out in AI infrastructure, and it places a ceiling on what efficiency improvements alone can achieve.

The third pressure is reach. The current AI stack often excludes smaller firms, public institutions, emerging markets, and low-connectivity users. If advanced AI continues to work only for hyperscalers and their immediate customer base, its benefits remain concentrated rather than diffused.

Frugal AI responds to all three at once: lower resource intensity, clearer economics, broader access.

3. Methodology

This paper draws on two sources. The first is a pre-programme survey of Global India Leadership Programme 2026 delegates. The second is the Frugal AI Hub framework as it has been developed through three prior white papers and the working session delivered to the cohort.

The cohort

The Global India Leadership Programme 2026 was convened by Global Education Lab and held at Cambridge Judge Business School from 9 to 13 March 2026. Twenty-two delegates participated. The composition skewed toward senior leadership: 9 founders, 6 C-suite executives, 5 directors and heads of function, and 2 others. By organisation size, 10 of the 22 lead organisations with fewer than 50 employees, with the rest distributed across mid-sized firms (6) and large enterprises with more than 1,000 employees (6).

The geographic base is predominantly India (14 of 22), with additional representation from the United Kingdom, Europe, Asia outside South Asia, and one respondent from North America. The sector spread covers manufacturing and industrial (4), services (4), education (3), agriculture and food systems (2), finance and FinTech (2), healthcare and life sciences (2), technology and digital (2), energy and climate (1), and two organisations whose business does not fit a single category. This spread matters for the analysis in section 4, where sector patterns prove informative.

The survey instrument

The pre-programme survey was administered through Qualtrics and covered six domains: digital transformation and technology, strategic growth and internationalisation, sustainability and ethics, leadership and culture and human capital, corporate governance and institutional frameworks, and financial stability and economic policy. Each domain combined Likert-scale self-assessment of leadership competence with open-ended questions on AI barriers, opportunities, and priorities for the next five years.

Domain-level competence scores are reported in the companion Global Education Lab discussion paper. This paper focuses on the AI-specific items: views on AI's transformative potential at the general, industry, and organisational level; current stage of organisational AI adoption; primary business objectives driving AI interest; barriers to faster adoption; opportunities seen; preferred adoption approach; and free-text reflections on context.

Analytical approach

The survey data is small (N=22) and the analysis is correspondingly modest in its statistical claims. We treat the survey as a structured baseline of how a particular cohort of senior leaders thought about AI at a particular moment, not as a generalisable population estimate. Where the paper uses the data, it does so to illustrate patterns visible in the responses rather than to test hypotheses.

The analytical approach has three steps. We first describe the cohort's views on AI's transformative potential and current adoption stage. We then map the cited barriers and opportunities onto the four Frugal AI pillars and the four levels of AI use, treating the mapping itself as a finding: the categories the framework anticipates are the categories the cohort actually populates. We finally read the free-text responses for sector-specific patterns, which we present as three recurring archetypes of Frugal AI practice rather than as exhaustive sector profiles.

The companion Global Education Lab discussion paper presents the same survey from a leadership-development perspective. Where this paper differs is in framing the responses as evidence about the adoption problem itself, rather than as input to the design of executive education.

4. What the cohort told us

The cohort's view of AI is consistent across three layers of generality but with a discernible step-down at each layer. Asked about AI in general, 14 of 22 respondents call it highly transformative and a further 6 call it moderately transformative; only 1 reports no meaningful impact. Asked about AI in their industry, the modal response shifts to "moderately transformative" (11 of 22), with 7 calling it highly transformative and 3 reporting limited impact. Asked about AI in their own organisation, the picture again moderates: 10 moderately transformative, 7 highly transformative, 2 limited, 1 no meaningful impact, with 2 declining the prescribed categories.

The pattern is informative. AI capability at the general level is rarely doubted; capability at the level of one's own institution is. The cohort is, however, actively engaging. Twelve of 22 describe their organisation as having started on the AI journey, 8 as considering starting, 1 as well underway, and 1 as not yet having considered AI. The cohort has moved past the question of whether to engage and is negotiating the harder questions of how.

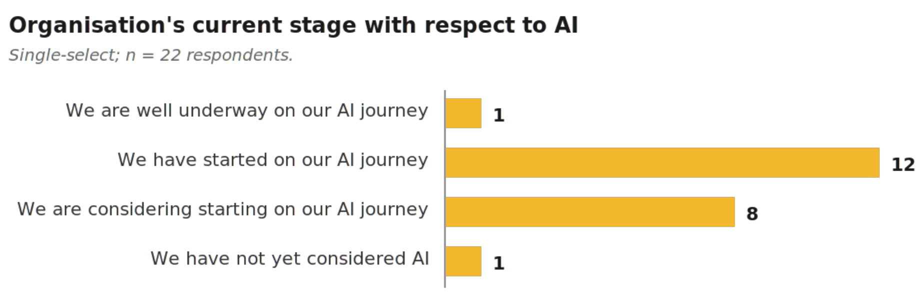

Chart 1. Current organisational stage with respect to AI (Q10). Single-select; n = 22.

Two further patterns frame what follows. When asked about the primary business objective driving AI interest, the cohort prioritises cost optimisation and operational efficiency (cited by 14 of 22), better decision-making and forecasting (13), revenue growth and customer acquisition (11), and innovation of products, services, or business models (10). Risk management and compliance rank lower (4). The cohort is approaching AI as a productivity instrument first and a transformation instrument second.

When asked which approach best reflects how the organisation intends to adopt AI, the dominant single choice is "frugal, modular AI solutions tailored to our context" (11 of 22). A further 8 have not yet defined an approach. Two have chosen outsourced AI through vendors and partners, one has chosen in-house development, and none have chosen large-scale enterprise AI platforms. The cohort is not, by and large, attempting to compete with hyperscalers on architecture. Half the cohort already names the approach this paper sets out to describe.

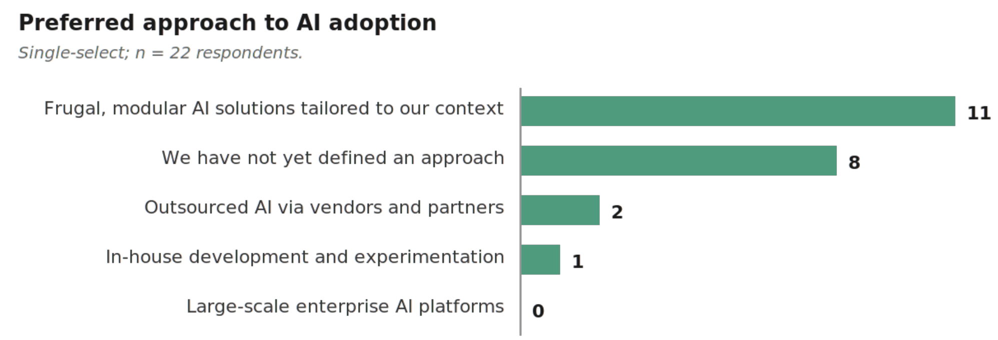

Chart 2. Preferred approach to AI adoption (Q16). Single-select; n = 22.

The barriers leaders face

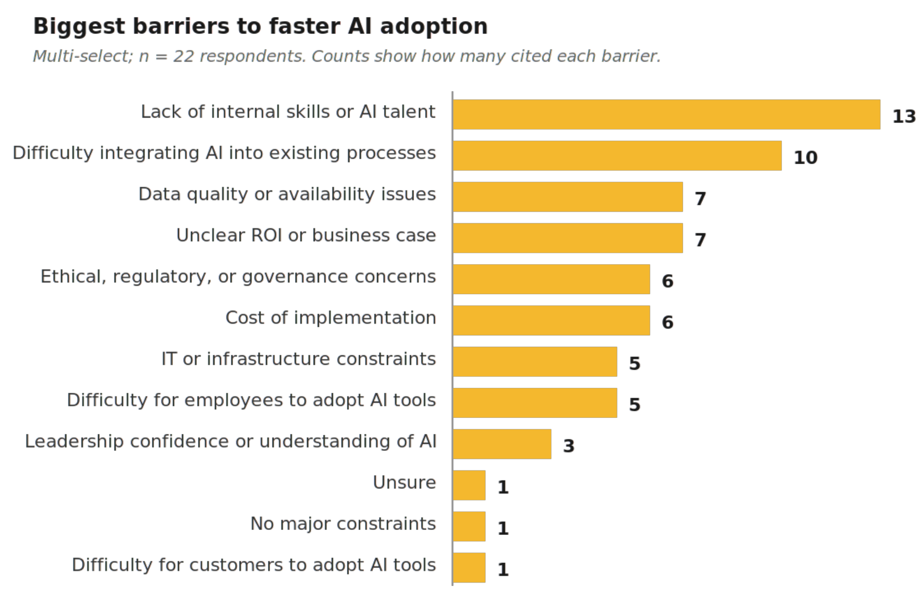

Chart 3. Biggest barriers to faster AI adoption (Q12). Multi-select; n = 22.

The most cited barriers to faster AI adoption, ranked by the number of respondents naming each (multi-select, n = 22), are:

- Lack of internal skills or AI talent (13)
- Difficulty integrating AI into existing processes (10)
- Data quality or availability issues (7)
- Unclear ROI or business case (7)
- Ethical, regulatory, or governance concerns (6)
- Cost of implementation (6)
- IT or infrastructure constraints (5)
- Difficulty for employees to adopt AI tools (5)
- Leadership confidence or understanding of AI (3)

The remaining options (difficulty for customers to adopt AI, "no major constraints," and "unsure") were each cited by a single respondent.

Four barriers (skills, integration, data, and unclear ROI) form a recognisable group describing the organisational gap between AI capability and AI deployment.

Eighteen of the 22 respondents, or 82 per cent, cite at least one of these four. That figure is the most direct quantitative finding the survey produces. It says, with little ambiguity, that the gap between capability and absorption is where executive attention is currently most needed.

A senior executive in manufacturing wrote that "use case identification as initial challenge, but barriers of implementation even harder to overcome, i.e. know-how, approval to use, IT set-up and security." A healthcare leader noted that "integrating AI into existing systems and workflows can be complex, and employees may need support and reassurance to confidently adopt new tools." An education leader pointed to "the shortage of skilled faculty and the high cost of GPU-heavy infrastructure" combined with "rigid regulatory frameworks" as the binding constraint. A financial services leader wrote of "resistance to change, no human capital ready to transform, infrastructure, prioritisation of other functions in AI adoption." The categories the survey instrument provided and the categories the respondents reach for in free text describe the same phenomenon.

The opportunities leaders pursue

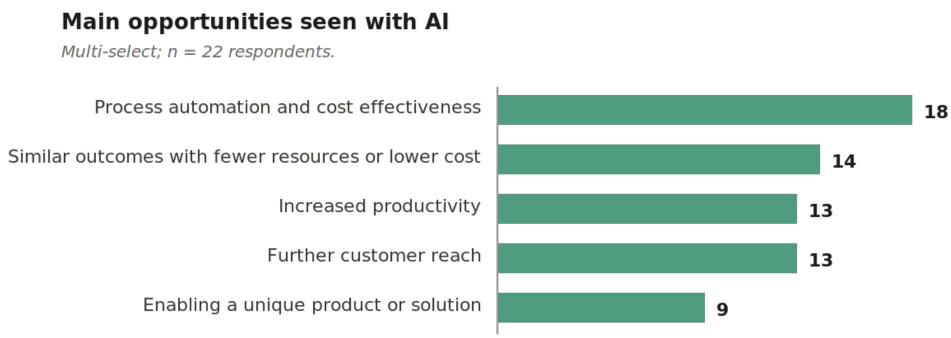

Chart 4. Main opportunities seen with AI (Q14). Multi-select; n = 22.

The most cited opportunities, ranked by number of respondents (multi-select, n = 22):

- Process automation and cost effectiveness (18)
- Delivering similar outcomes with fewer resources or lower cost (14)
- Increased productivity (13)
- Further customer reach (13)
- Enabling a unique product or solution (9)

Process automation and cost effectiveness is cited by 82 per cent of the cohort, the highest mark any single response option achieves in the survey. The rhetoric of frontier AI (transformative business models, displacement of incumbents, technological singularity) is not where most of the cohort's stated opportunity sits. The cohort is asking for efficiency and reach.

The perception-capability gap

The most important finding in the survey is the gap between two adjacent self-assessments. The cohort considers AI to be highly transformative for their industry and organisation. The cohort also considers unclear ROI, integration difficulty, governance concerns, and skills gaps to be the primary barriers to faster adoption.

These two findings are usually presented as a contradiction, but they are the same finding seen from two angles. AI is transformative, and most of the cohort's organisations are not yet equipped to absorb the transformation. Capability exists but the organisational capacity to deploy it productively does not.

How should organisations approach AI absorption when capability is abundant and capacity is the constraint?

Three sectors' patterns

The free-text responses and the in-programme discussions surface three recurring patterns.

Governance. In banking and healthcare, leaders consistently frame AI adoption as a governance challenge before a technological one. The decisive question is whether the deployed system can be audited, explained, and supervised by the institution that bears the consequences. Frugal AI in this archetype manifests as deliberate choice of smaller, more interpretable models; investment in monitoring and validation infrastructure ahead of model deployment; and a willingness to forgo capability gains that cannot be governed. The risk these sectors guard against is regulatory action and reputational damage. The trade-off is slower adoption than vendors and consultants typically recommend.

Affordable reach. In FMCG, education, and parts of consumer-facing fintech, the dominant question is whether AI can extend the organisation's effective reach without proportional cost growth. The archetype shows up as voice-first interfaces for low-literacy users, multilingual support, low-bandwidth deployment, and partnerships with local distribution networks. Frugal AI in this mode is about lowering the per-user cost of high-quality service so that the addressable market expands. The risk these sectors guard against is irrelevance to the next billion users. The trade-off is that the technical work is harder than the marketing-led version of AI adoption suggests.

Operations-grade reliability. In manufacturing, logistics, and B2B services, the dominant question is whether AI can reduce the cost denominator of existing operations without disrupting reliability. The needs for these sectors are route optimisation, demand forecasting, predictive maintenance, and quality control augmentation. The risk these sectors guard against is operational disruption. The trade-off is that the productivity gains are bounded by the workflow design that contains them.

5. Measuring what matters

Setting targets, measuring outcomes at portfolio level, and optimising is a common path followed by organisations that would benefit greatly from AI implementations.

Why current AI cost frameworks fall short

AI has different cost structures and behaviours than traditional software, with potentially different implications for capex and opex. Prices are also highly variable as models change and their costs do too. Managing AI costs requires more visibility and telemetry than the traditional systems existing in companies. Moreover, good AI relies on good data and stronger governance capabilities.

Most organisations evaluate AI initiatives using one of two narrow cost lenses. The first is model licensing cost: the per-token, per-call, or per-seat price charged by the vendor. The second is project cost: the budgeted spend on a particular initiative over a defined period. Neither captures the full economic and environmental footprint of an AI system in production.

The omitted categories include compute and infrastructure cost beyond the licensing fee (autoscaling spend, idle capacity, storage, networking); data lifecycle cost (ingestion, preparation, governance, storage, maintenance); personnel cost across the lifecycle (development, support, training, change management); integration and orchestration cost (system integration, stakeholder coordination, eventual decommissioning); maintenance and monitoring cost (continuous improvement, model upkeep, alerting); governance, risk, and compliance cost (security, privacy, audit, sovereignty); and environmental cost (energy use and carbon emissions attributable to the workload).

Total Cost of Ownership (TCO) seven pillars

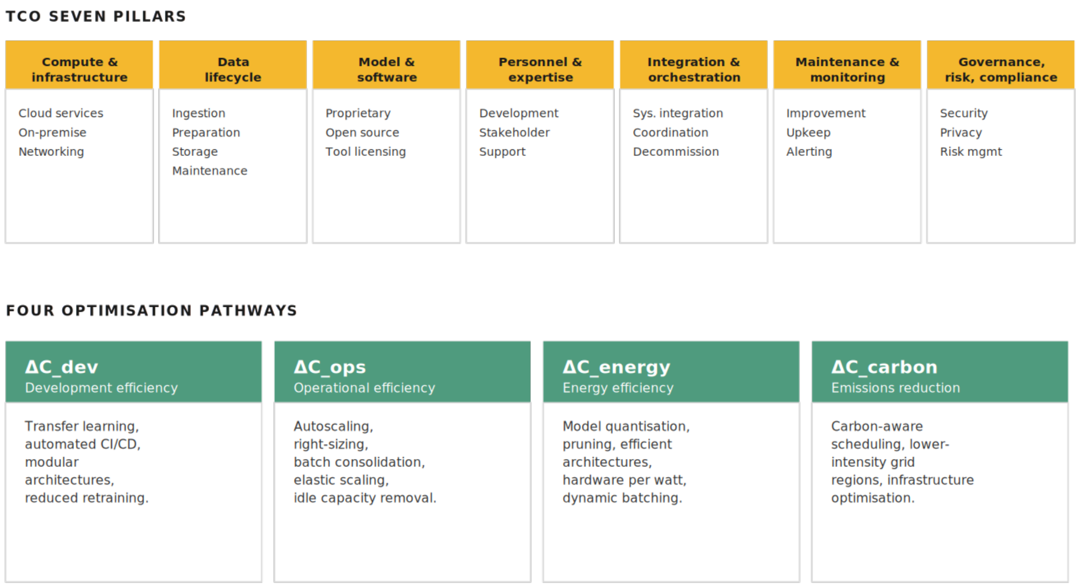

Figure 4. The TCO seven pillars and four optimisation pathways. The seven pillars structure where AI cost lives. The four ΔC pathways structure where it can be reduced.

The Frugal AI Hub measurement framework organises total cost of ownership into seven pillars:

- Compute and infrastructure (cloud services, on-premise, networking)
- Data lifecycle (ingestion and indexing, preparation and governance, storage, maintenance)
- Model and software (proprietary model usage, open-source alternatives, platform and tool licensing)
- Personnel and expertise (development and design, stakeholder management, support and training)
- Integration, orchestration, and decommission (system integration, stakeholder coordination, decommissioning)
- Maintenance and monitoring (continuous improvement, model and data upkeep, monitoring and alerting)
- Governance, risk, and compliance (security and compliance, data privacy and sovereignty, risk management)

The pillars are designed to be tractable. Each maps to existing cost categories that finance functions already track or can begin to track using standard cloud and operations tooling. The discipline the framework imposes is that all seven should be tracked at the level of the individual use case, as well as at portfolio level.

A worked example from fintech

Telemetry leads to visibility, measuring leads to optimisation.

The Hub's May 2026 white paper on Frugal AI in fintech applies the framework to a real-time credit card fraud detection use case at a large United States issuer. The numbers are illustrative rather than proprietary, but they are calibrated against realistic institutional parameters.

The Frugal AI optimisation applied across development, operations, energy, and emissions leads to reduced costs: roughly a 15% reduction in labour, 25% in operational compute, and 44% in both energy and carbon.

The baseline ROI stands at 5.03×; with Frugal AI optimisation the ROI rises to 6.27×, a 24.6% improvement, driven entirely by reduction in the cost denominator rather than by any change in the value produced.

The cost reductions are the cumulative effect of disciplined choices that produce a materially better ROI profile in aggregate. Frugal AI is precision applied consistently, rather than a single architectural breakthrough.

From financial ROI to societal impact

Measuring should not only be applied to ROI but also to societal impact at portfolio level.

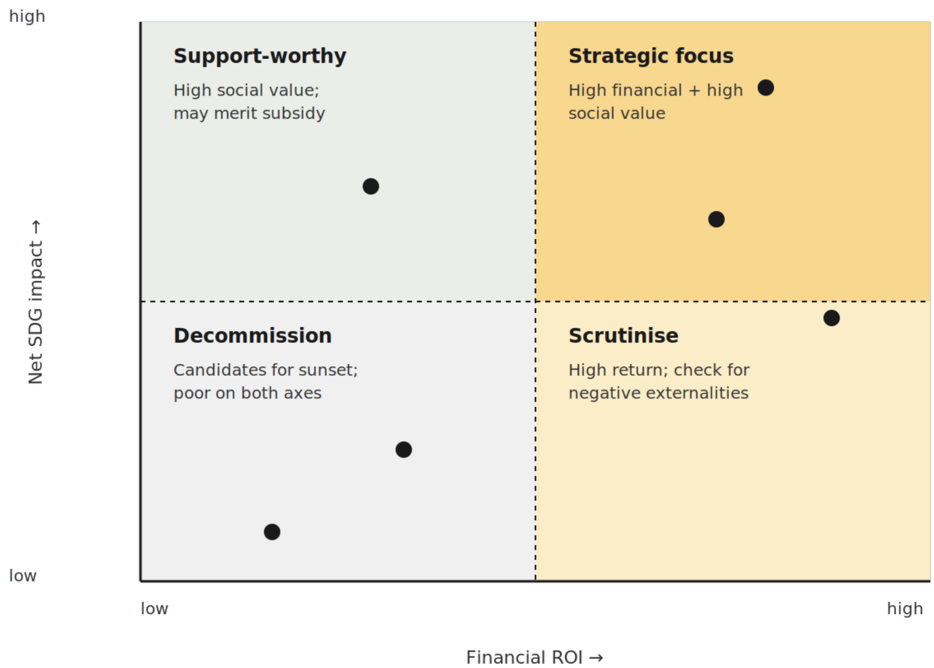

Figure 5. ROI × SDG portfolio quadrant. Projects are placed on a two-by-two of financial return against net social impact.

At the portfolio level, the relationship between financial return and social impact can be made visible by plotting initiatives on a two-by-two of ROI against net SDG (Sustainable Development Goal) score. High-ROI, high-SDG initiatives are the natural strategic focus. High-ROI, low-SDG initiatives warrant scrutiny for negative externalities the financial measure does not capture. Low-ROI, high-SDG initiatives may merit support for reasons beyond direct return. Low-ROI, low-SDG initiatives are candidates for decommissioning.

The portfolio view changes the boardroom conversation. AI initiatives are no longer evaluated only against revenue or cost targets. They are evaluated against the combined return that organisations are increasingly asked to deliver: economic, environmental, and social.

6. The India lens

The case for Frugal AI is sometimes read as an argument about emerging markets, about resource-constrained settings, or about the Global South. It is partly that. It is also more general. The dominant question of where AI is being built, and for whom, increasingly shapes what AI is being built.

The three paths

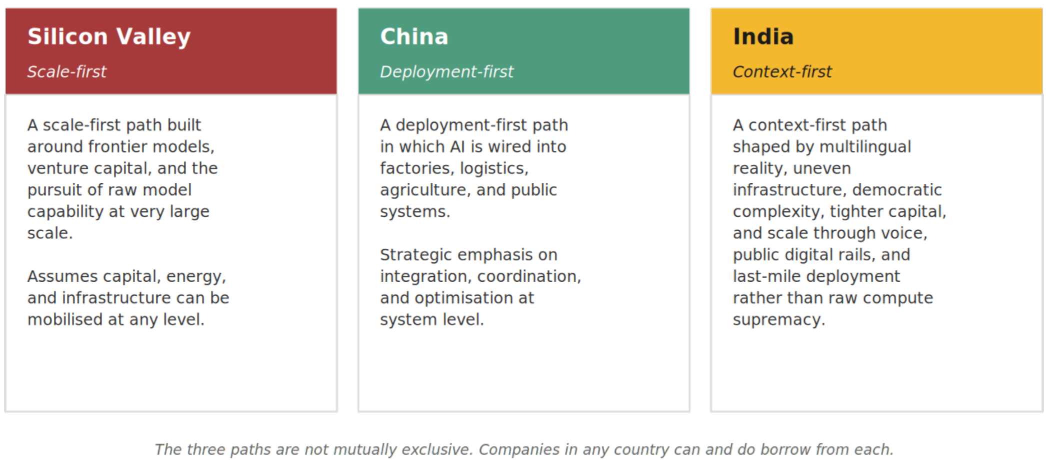

Figure 6. Three contemporary paths to AI: Silicon Valley scale-first, China deployment-first, India context-first.

A useful framing distinguishes three contemporary approaches to AI development at the national and ecosystem level.

The first is the path most visible in Silicon Valley: scale-first, built around frontier models, venture capital, and the pursuit of raw model capability at a very large scale. The operating assumption is that capital, energy, and infrastructure can be mobilised at whatever level the frontier requires. The competitive logic is that whoever builds the most capable model wins the largest share of the value. This requires large capex investment, often fronted by VCs and enabled by large-scale capital recycling across the AI ecosystem.

The second is the path most visible in China: deployment-first, with AI wired into factories, logistics, agriculture, and public systems. The strategic emphasis is on integration, coordination, and optimisation at system level rather than on raw model supremacy. The competitive logic is that the country that deploys AI most widely captures the productivity gains, regardless of whose model sits underneath. This requires significant patient government investment.

The third is the path most visible in India: context-first, shaped by multilingual reality, uneven infrastructure, democratic complexity, tighter capital, and the opportunity to scale through voice, public digital infrastructure, and last-mile deployment rather than through raw compute supremacy. The competitive logic is that distinctive constraints (language diversity, low connectivity, cost sensitivity, public-sector adoption at scale) produce distinctive architectures that travel well to other contexts with similar constraints.

The three paths are not mutually exclusive at the firm level. Companies in any country can and do borrow from each.

India as a context-first capital

India's conditions (a population approaching 1.5 billion, more than twenty official languages, public digital infrastructure that already reaches the majority of citizens, a cost-sensitive consumer base, and an administrative state that increasingly delivers services digitally) make it a setting where deployment discipline matters greatly. Building systems that work across large populations, complex institutions, and the messy last mile is a different competitive game from building the most parameter-heavy model or production-integrated systems.

The discipline this imposes is useful for executives in any geography. The use cases that survive the conditions of Indian deployment (low bandwidth, multilingual, cost-sensitive, regulator-aware, institutionally embedded) tend to be the use cases that also survive at scale in other markets.

Voice-first systems and linguistic reach

A specific implication of the context-first path is the strategic role of voice as a primary interface. In contexts where the keyboard is a barrier (whether through literacy, language, or device constraints), voice-first interfaces lower the threshold to AI access and expand the addressable user base substantially.

The strategic point goes beyond accessibility. Designing for voice in non-Western languages requires investment in linguistic representation, in community-verified data, and in culturally grounded deployment. The resulting systems perform better not only for the originally targeted users but, in many cases, for the broader global user base whose interaction with AI is also moving toward speech.

Frugal AI as a corridor, not a category

The broader implication is that Frugal AI is best understood as a corridor rather than a regional category. The corridor connects the Global South's deployment-heavy constraints to the Global North's growing concerns about cost, energy, governance, and reach. Many of the design moves that respond to Indian constraints (smaller models, edge deployment, multilingual capability, voice interfaces, public-digital-infrastructure integration) are increasingly relevant to firms in any geography facing pressure on the same three dimensions: cost, energy, and reach.

The cohort discussions made this concrete. UK and German executives in the room found that the design patterns being discussed in the Indian deployment context applied directly to their own operations, because the discipline of designing under constraint produced systems that were also more economical, more sustainable, and more governable in their domestic contexts.

This is the strategic case for Frugal AI as a global proposition. The countries and firms that learn to deploy under constraint will be better positioned for an environment in which constraint becomes the default condition.

7. From insights to actions

The cohort survey asked what delegates hoped to achieve through the programme. The responses, ranked by number of respondents naming each (multi-select, n = 22), were: identifying realistic AI use cases for their context (18), a better understanding of AI opportunities and risks (17), learning from peers in similar contexts (15), understanding costs and trade-offs and ROI of AI (14), and developing next steps or a roadmap (13). The Frugal AI framework offers a direct response to each.

The starting point needs to be the level of aspiration set in Figure 1. This defines the longer-term view of what we need to achieve. Then, the next level is to start with use cases that are aligned to that vision and represent clear value for the organisation, albeit also considering the absorption constraints.

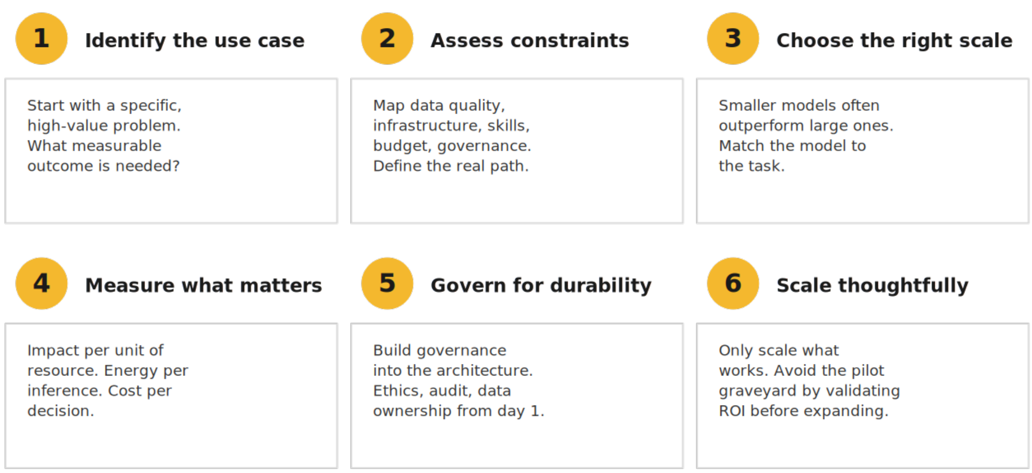

Figure 7. The six-step Frugal AI adoption framework.

The six steps are intended as a disciplined sequence for any AI initiative.

Step 1: Identify the use case. Start with a specific, high-value problem within the reach of the organisation, rather than AI for its own sake. What is the measurable outcome the organisation needs, and what is the cost of not achieving it?

Step 2: Assess constraints. What is the real state of data quality, infrastructure, skills, budget, and governance readiness? The frugal path is defined by actual constraints.

Step 3: Choose the right scale. Modular, smaller models often outperform large ones in specific domains. Match the model to the task. Resist the assumption that the largest available model is the appropriate one.

Step 4: Measure what matters. Impact per unit of resource. Energy per inference. Cost per decision. Accuracy is necessary but not sufficient.

Step 5: Govern for durability. Build AI governance into the architecture from day one. Ethics, data ownership, audit trails, decommissioning criteria. Retrofitting governance is more expensive than building it in.

Step 6: Scale thoughtfully. Only scale what works. Avoid the pilot graveyard by validating ROI before expanding, and by deciding in advance what would justify decommissioning.

A 90-day Frugal AI roadmap

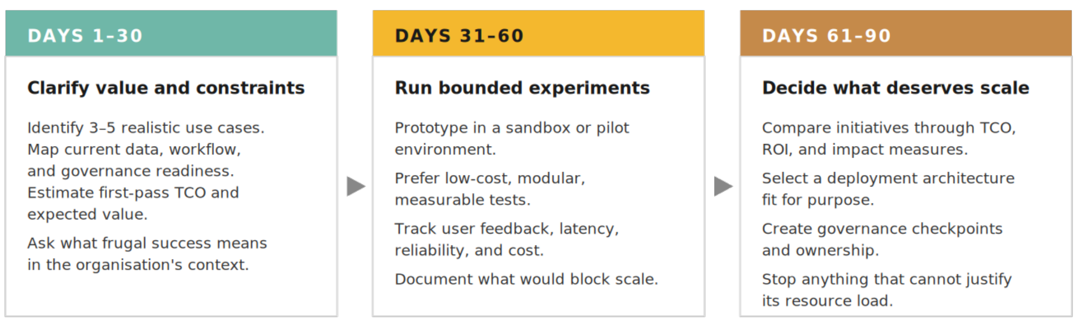

Figure 8. A 90-day Frugal AI roadmap for organisations beginning serious AI deployment.

The six steps map naturally onto a 90-day cadence for organisations beginning serious AI deployment.

Days 1–30: Clarify value and constraints. Identify three to five realistic use cases. Map current data, workflow, and governance readiness. Estimate first-pass TCO and expected value. Ask what frugal success would mean in the organisation's context.

Days 31–60: Run bounded experiments. Prototype in a sandbox or pilot environment. Prefer low-cost, modular, measurable tests. Track user feedback, latency, reliability, and cost. Document what would block scale.

Days 61–90: Decide what deserves scale. Compare initiatives through TCO, ROI, and impact measures. Select a deployment architecture fit for purpose. Create governance checkpoints and ownership. Stop anything that cannot justify its resource load.

Pathways for scaling

Capability to absorb AI will vary from organisation to organisation. First scaled use cases give incredible insights for playbooks for the next deployment. What worked according to plan? What did not? Measurements can offer a lot of insight on how to approach the next use-case scale-up.

The boardroom role

The role of the boardroom is critical in the AI journey. The boardroom should play both offence (creating the platform and culture to innovate and scale) and defence (ensuring the right governance is in place).

Developing AI capability requires patient investment. Early failures should not disappoint; measurements should not be read as a constraint but as a clear learning for constant optimisation.

Boardrooms should be concerned about ensuring the organisation has the right skillset and capacity to scale AI, that the vision for AI is aligned to the organisation's strategic objectives, and that the potentially disruptive, transformative impact on business models and other elements is clear and that scaling pathways and learnings are discussed.

8. Towards a longitudinal research programme

The Global India Leadership Programme 2026 cohort is a snapshot. The value of the snapshot increases substantially if it becomes the first observation in a panel.

The Frugal AI Hub and Global Education Lab propose, through this paper, a longitudinal research programme that uses successive cohorts of Global Education Lab executive programmes as a panel for tracking how organisations move from AI capability to AI operationalisation between 2026 and 2030.

The programme would have four elements.

A standardised pre-programme instrument. The 2026 survey provides a baseline. Future cohorts would receive a refined version of the same instrument, allowing comparison across cohorts and over time on AI views, adoption stage, barriers, opportunities, and preferred approach.

A consistent framework for analysis. The Frugal AI pillars and the four levels of AI use provide the analytical categories. As the panel grows, sector-level and geography-level patterns become more reliably visible, and the archetypes proposed in section 4 can be tested and refined.

Cohort-specific deep work. Each year's programme produces a thematic application: this year's paper applies the framework to leadership and adoption; the May 2026 fintech paper applies it to financial services; future papers could apply it to healthcare, public sector, climate technology, or education. The thematic work feeds back into refinement of the framework.

An annual publication. A short annual report would track movement in the cohort data, surface emerging patterns, and update the playbook in section 7. The annual cadence matters because the underlying question (how organisations absorb AI) changes faster than five-year research cycles can capture.

The research programme has three intended audiences. For executives, it provides a continuous source of benchmarking against peers. For policymakers, it provides evidence about where executive practice is and is not progressing on questions of AI governance, sustainability, and inclusion. For the academic community, it provides a structured longitudinal dataset on a phenomenon that is currently under-measured.

The proposal is open to additional institutional collaborators who share the analytical agenda. Expressions of interest can be directed to the Hub.

References

Frugal AI Hub and partner publications

Frugal AI Hub, 2025. *Accelerating a Frugal AI Ecosystem.* Cambridge Judge Business School.

Frugal AI Hub, 2025. *Frugal AI: Technical and Socio-Economic Foundations.* Cambridge Judge Business School.

Frugal AI Hub and United Nations International Computing Centre (UNICC), 2025. *From Total Cost of Ownership to Social Impact: A Frugal AI Framework to Measure AI Portfolios as Strategic Assets.* Cambridge Judge Business School and UNICC.

Frugal AI Hub, 2026. *Applications of Frugal AI in the Fintech Industry: Measuring Total Cost of Ownership and Return on Investment in Real-Time Fraud Detection.* Cambridge Judge Business School, May 2026.

Global Education Lab and Cambridge Judge Business School, 2026. *Leadership in the Age of AI: Strategy, Innovation, and Decision-Making in a Transforming Economy.* Companion discussion paper.

Further reading on Frugal AI and AI sustainability

Radjou, N. and Prabhu, J., 2015. *Frugal Innovation: How to do more with less.* The Economist Books.

Schwartz, R., Dodge, J., Smith, N.A. and Etzioni, O., 2020. Green AI. *Communications of the ACM*, 63(12), pp.54–63.

Strubell, E., Ganesh, A. and McCallum, A., 2019. Energy and Policy Considerations for Deep Learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp.3645–3650.

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M. and Dean, J., 2021. Carbon Emissions and Large Neural Network Training. arXiv:2104.10350.

Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D. and Pineau, J., 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *Journal of Machine Learning Research*, 21(248), pp.1–43.

Standards and reports

AFNOR, 2024. *AFNOR SPEC 2314: Frugal Artificial Intelligence.* Guidelines for the Design and Deployment of Frugal AI Systems.

ISO/IEC 42001:2023. *Information Technology: Artificial Intelligence Management System.*

World Economic Forum, 2026. *The Future of Jobs Report 2026.*

Frugal AI Hub

at Cambridge Judge Business School

Trumpington Street
Cambridge CB2 1AG
United Kingdom

https://frugalai.org