

Frugal AI: Driving Sustainable, Inclusive, and Scalable Innovation

White Paper | April 2025





CONTENTS

Foreword	5
About The Authors	6
Executive Summary	8
Introduction: The Case for Frugal AI	9
Defining Frugal AI: Core Pillars & Principles	11
The Frugal AI Toolkit: Techniques & Methodologies	14
Measuring Frugality: Frameworks & Standards	16
Frugal AI in Action: Case Studies & Applications	17
The Cambridge Judge Frugal AI Hub Ecosystem	19
Addressing Misconceptions: Frugal AI Myths vs. Reality	20
Frugal AI vs. Traditional AI: A Comparative View	21
Conclusion: Towards a Sustainable & Inclusive AI Future	22
References	23

True innovation isn't measured just by compute power but by human impact. The most powerful AI in the future will be the one that's efficient enough to scale, inclusive enough to serve, and ethical enough to endure. Frugal AI redefines progress by asking how we can design intelligence that is open, accessible, and truly in service of all.

FOREWORD



The world today stands at the crossroads of unprecedented challenges and opportunities. Climate change, inequality, and resource scarcity are pressing us to rethink how we innovate and solve problems. In this context, the principles of frugal innovation—doing more with less for more people—are more relevant than ever. For over a decade, my colleagues and I at Cambridge have explored how individuals, organisations, and governments can leverage ingenuity to address complex issues sustainably and inclusively.

Frugal innovation is not just about cost-cutting; it is about reimagining processes, products, and systems to create value under constraints. It thrives on simplicity, adaptability, and a deep understanding of local needs. By fostering collaboration across sectors and embracing resourcefulness, frugal innovation has empowered communities in both developed and emerging economies to tackle their most pressing challenges.

Artificial Intelligence (AI), a transformative force of our time, holds immense potential to amplify the impact of frugal innovation. However, as we integrate AI into our lives and systems, we must ensure that it aligns with the principles of frugality: accessibility, affordability, and sustainability. This is where Frugal AI emerges as a groundbreaking concept—AI designed not just for efficiency but for equity.

The Frugal AI approach calls for designing intelligent systems that are lightweight, energy-efficient, and tailored to meet the needs of underserved populations. It emphasizes democratising access to AI technologies by reducing dependency on expensive infrastructure and fostering open innovation ecosystems. By doing so, Frugal AI can bridge digital divides, empower local innovators, and drive inclusive growth.

This whitepaper is a timely contribution to the ongoing discourse on responsible AI development. It underscores the importance of aligning technological advancements with societal goals and environmental sustainability. As you explore its insights and recommendations, I encourage you to reflect on how we can collectively harness the power of Frugal AI to create a fairer and more sustainable future.

Together, let's embrace this opportunity to innovate responsibly—doing more with less for more people—and ensure that technology serves humanity in its truest sense.

Jaideep Prabhu

**Professor of Marketing & Nehru Professor of Indian Business
Cambridge Judge Business School**

ABOUT THE AUTHORS



Venkata "Serish" Gandikota

Co-lead, Frugal AI Hub | Visiting Fellow, Cambridge Judge Business School

Serish is an innovation strategist, impact & climate fund adviser, entrepreneur, and researcher focused on frugal innovation, sustainability, impact investing, and inclusive technology. As Co-lead of the Frugal AI Hub at Cambridge Judge Business School, he works at the intersection of AI, systems thinking, and social impact. He has over a decade of experience designing and delivering frugal innovation programs across Europe, Asia, and Africa, and has been involved in Horizon 2020 and EU-backed initiatives on citizen science, sustainability, and impact investing. Venkata is also the co-founder of InnoFrugal, a platform advancing affordable and resource-efficient innovation globally.



Elizabeth Osta

Co-lead, Frugal AI Hub | Visiting Fellow, Cambridge Judge Business School

Elizabeth is a digital and data strategist and advisor working with boards and CXOs on innovation, data and AI strategy and digital transformation. As Co-lead of the Frugal AI Hub and founder of Digital Forward, she focuses on advancing inclusive and sustainable AI practices across sectors. Previously, Elizabeth served as Chief Data Officer at HEINEKEN, leading data transformation across 80 countries. She has held leadership roles at Barclays, dunhumby, Gillette, McKinsey, and Social Finance, where she explored the use of data to support vulnerable populations. Elizabeth also serves as a non-executive advisor to several fintech startups and is a Visiting Fellow at Cambridge Judge Business School.



EXECUTIVE SUMMARY

Artificial Intelligence (AI) is rapidly transforming industries and society, yet its expansion faces critical hurdles. Conventional AI development often demands massive computational power, vast datasets, and significant energy consumption, leading to escalating costs, environmental concerns, and limited accessibility. Frugal AI emerges as a necessary paradigm shift, defined by the principle of "doing more with less" – designing, developing, and deploying AI systems that are resource-efficient, affordable, accessible, and sustainable.

This whitepaper addresses the "Triple AI Challenge": Economic Pressure, where fewer than 20% of companies achieve sustainable ROI from AI, hampered by high costs and immature TCO methodologies; Environmental Impact, with large models consuming energy equivalent to hundreds of homes and contributing significantly to carbon emissions; and Limited Reach, where resource constraints prevent many potential beneficiaries, especially SMEs and those in emerging markets, from leveraging AI.

Frugal AI offers a strategic solution, moving beyond mere cost-cutting to enable responsible, ethical, and scalable innovation. Rooted in the philosophy of frugal innovation and Jugaad, it emphasizes lightweight models, data parsimony, energy efficiency, and a focus on measurable positive impact. The Frugal AI Hub at Cambridge Judge Business School serves as a nexus for research, collaboration, and adoption in this vital field. This paper outlines Frugal AI's core principles, explores key techniques and frameworks, showcases real-world applications, and invites stakeholders to engage with the Hub in building a more sustainable and inclusive AI future.

INTRODUCTION: THE CASE FOR FRUGAL AI

The transformative potential of Artificial Intelligence is undeniable, yet its current trajectory raises urgent questions. While AI drives innovation across sectors, its development is increasingly characterized by resource intensity. The training of large language models can consume energy equivalent to the annual consumption of over 100 homes, and the global electricity consumption of data centers rivals that of entire countries. This immense appetite for energy and computation translates into a significant environmental footprint and prohibitive costs, creating barriers to entry and hindering widespread adoption.

The True Impact of AI

AI is reshaping our world, but it comes at a cost—especially in terms of energy and carbon emissions. Here's what the numbers reveal:

Every AI Query Has a Carbon Cost

A single GPT-3.5 or GPT-4 query is estimated to produce 2–4.5g of CO₂, depending on hardware efficiency, energy source, and inference method.

That's significantly higher than a traditional Google search (~0.2g CO₂/query).

15–20 Prompts = One Kettle Boil

It takes roughly 70g of CO₂ to boil a kettle of water.

Just 15–20 ChatGPT queries can match that carbon cost.

130–150 Queries = One Laundry Load

One load of laundry emits ~600g of CO₂. Just 130–150 AI prompts can match that—small actions add up.

100,000 Prompts = SF to London & Back

Flying roundtrip from San Francisco to London emits ~1 ton of CO₂. You'd only need around 100,000 queries to match that.

30,000+ GPUs Running Daily

Models like GPT-4 require tens of thousands of GPUs, many running continuously. This contributes to rising global energy demand in data centers.

Summary



AI is powerful—but not weightless.

Its carbon footprint depends on model size, energy source, and usage volume.

Frugal AI offers a path to smarter, leaner, and greener intelligence.

Note: These estimates are based on public research and modeling by institutions like Hugging Face, MIT Technology Review, and Salesforce's AI Energy Score framework. Carbon cost can vary depending on model size, prompt length, deployment architecture (e.g., local vs. cloud), and whether green energy is used.

Compounding the environmental concerns are significant economic pressures. Many AI initiatives struggle to move beyond the pilot stage, with fewer than one in five companies successfully scaling AI with sustainable Return on Investment (ROI). Methodologies for calculating ROI and Total Cost of Ownership (TCO) remain underdeveloped, obscuring the true cost of AI deployment. Furthermore, the benefits of AI often fail to reach those who could gain the most, including small and medium-sized enterprises (SMEs), organizations in emerging economies, and applications requiring deployment on low-power devices or in areas with limited connectivity. This constitutes the "Triple AI Challenge" – a confluence of economic, environmental, and accessibility hurdles.



Leading countries by number of data centers as of March 2025

Source: Statista

Frugal AI offers a compelling response. Drawing inspiration from the principles of Frugal Innovation and the concept of Jugaad (resourceful improvisation), Frugal AI advocates for designing AI systems that achieve high performance with minimal resources – compute, data, energy, and capital. It represents a strategic shift away from a "scale-first, resource-heavy" approach towards one prioritizing efficiency, sustainability, and democratization. This approach recognizes differing global priorities: while environmental responsibility and energy efficiency may be primary drivers in the Global North, Frugal AI serves as a critical enabler of socio-economic empowerment, digital inclusion, and leapfrogging innovation in the Global South. It is not merely about cheaper AI; it is about smarter, leaner, more responsible, and ultimately more impactful AI for everyone.

DEFINING FRUGAL AI: CORE PILLARS & PRINCIPLES

Frugal AI is an intentional design and deployment approach for building AI systems that are:

- High-performing but lightweight
- Cost-effective and energy-efficient
- Accessible and scalable under resource constraints

It stems from frugal innovation and jugaad—a Hindi term meaning resourceful improvisation. Frugal AI translates this spirit into AI development: solving complex problems using minimal resources.

Innovation Through Constraints

Making AI work with limited data, computing power, and resources.

Enables deployment in low-power environments, edge devices, and developing regions.

It prioritizes simplicity, interpretability, and reuse of existing models to reduce costs

Efficient AI

Computational efficiency through techniques like model pruning, quantization, and knowledge distillation.

Ensures AI solutions run faster and consume less energy without sacrificing performance

Environmental Sustainability

The "Fair, Few, Frugal" (3Fs) framework emphasizes minimizing resource use for a more sustainable AI ecosystem.

This aligns with initiatives like AI Energy Scoring, which evaluates models based on energy consumption.

Sustainable Data Centres focus on optimised power usage, innovative cooling technologies and in some cases, use of AI to drive environmental efficiency

Democratization and Accessibility

Cost-effective AI can democratize access by lowering technical and financial barriers.

Makes advanced capabilities available to a broader range of users and use cases, particularly benefiting small and medium enterprises and society in constrained environments

Fast Scaling

Frugal AI enables fast scaling by optimizing models for low-cost deployment, leveraging lightweight architectures, and maximizing resource efficiency.

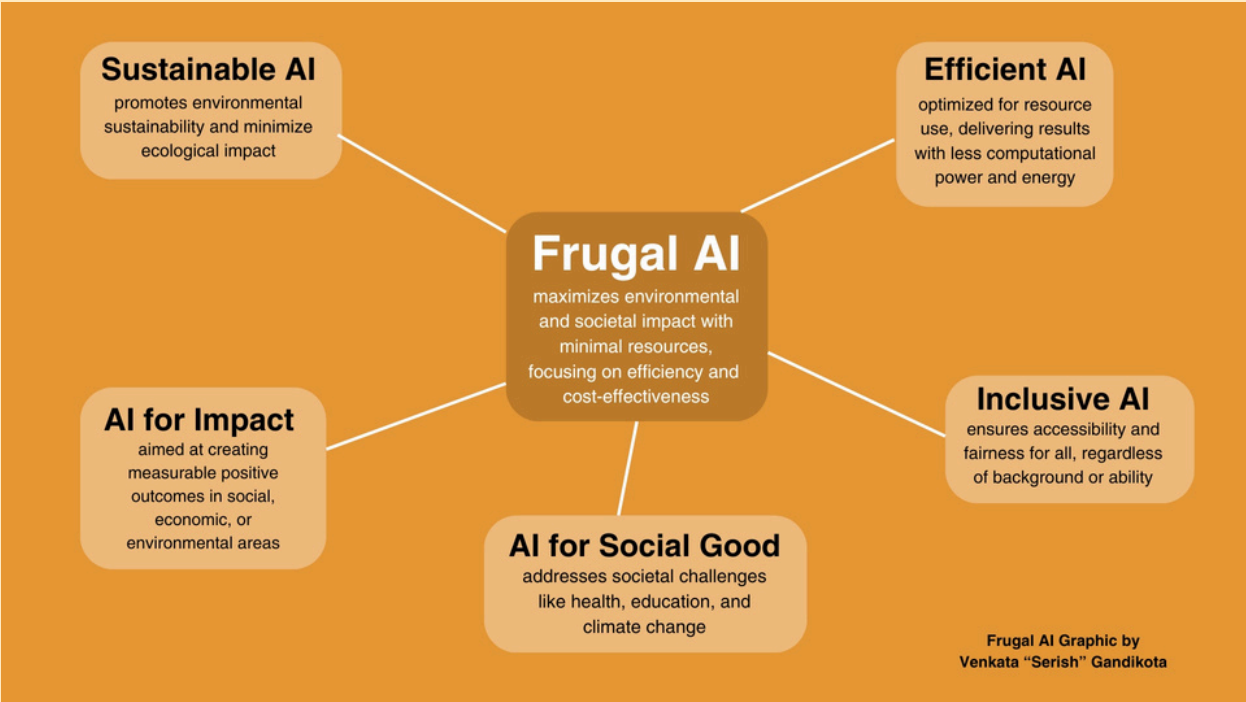
Allows organizations to expand AI solutions rapidly without significant infrastructure investment.

Frugal AI also aims at sharing best practices for effective scaling, reducing the hurdle that many face in driving value from AI

Frugal AI is built upon a foundation of core principles that guide the development and deployment of resource-conscious intelligent systems.

- **Resource Efficiency:** This is the central tenet, focusing on minimizing the consumption of computational power, energy, data, and financial capital throughout the AI lifecycle. It involves employing lean algorithms, optimizing models for size and speed, and practicing data parsimony – using the minimal amount of high-quality, relevant data necessary to achieve the desired outcome, rather than assuming "more data is always better."
- **Sustainability:** Frugal AI directly addresses the environmental impact of technology. This pillar emphasizes reducing the carbon footprint of AI through energy-efficient model training and inference, promoting the use of sustainable data centers powered by renewable energy, and utilizing tools like AI Energy Scoring to benchmark and incentivize lower energy consumption. It aligns AI development with broader ecological goals and net-zero targets.
- **Accessibility & Inclusion:** A core goal of Frugal AI is to democratize access to AI capabilities. This involves designing lightweight models deployable on low-power edge devices (smartphones, IoT sensors) and in resource-constrained environments. It aims to lower technical and financial barriers, making AI tools available and beneficial to SMEs, non-profits, emerging markets, and diverse user groups, thereby bridging the digital divide and ensuring fairness.
- **Impact & Scalability:** Frugal AI prioritizes delivering tangible, measurable value – whether social, environmental, or economic. It necessitates robust frameworks for evaluating ROI and TCO to ensure financial viability. Furthermore, by optimizing models for low-cost deployment and leveraging lightweight architectures, Frugal AI enables rapid scaling of solutions without requiring massive infrastructure investments, allowing impactful innovations to reach a wider audience quickly.

These pillars and principles are interconnected, collectively defining an approach to AI that is not only technologically sound but also economically viable, environmentally responsible, and socially equitable.



THE FRUGAL AI TOOLKIT: TECHNIQUES & METHODOLOGIES

Model Pruning

Removes unnecessary weights & connections in neural networks.

Reduces model size and computation while maintaining accuracy.

Example: ResNet pruning reduces parameters by 50–90% with minimal accuracy loss.

Quantization

Converts high-precision (e.g., 32-bit) weights into lower precision (e.g., 8-bit).

Reduces model size, memory footprint & speeds up inference.

Example: INT8 quantization speeds up inference by 4x with minimal accuracy loss.

Knowledge Distillation

Trains a smaller “student” model to mimic a larger “teacher” model.

Achieves near-teacher accuracy with fewer parameters.

Example: BERT distillation reduces model size by 60% while retaining 97% of accuracy.

Sparse & Low-Rank Approximation

Converts dense weight matrices into sparse representations.

Reduces computational complexity while preserving key information.

Example: GPT-3 sparsity techniques cut energy use by 50% with minor accuracy drop.

Efficient Architectures & Hardware Optimization

Uses AI models designed for efficiency (e.g., MobileNet, EfficientNet).

Optimized execution using specialized AI chips (TPUs, FPGAs).

Example: EfficientNet achieves state-of-the-art performance with 5x fewer computations.

Data Parsimony & Selection

Focuses on smaller, high-impact datasets using active learning or smart curation.

Lowers labeling costs and computational demand.

Example: Active learning in medical imaging selects only 20% of data for labeling with near full-data performance.

Synthetic Data & Augmentation

Expands datasets using artificial examples or transformations.

Improves generalization without new data collection costs.

Example: Image augmentation boosts model robustness with minimal new data.

Transfer Learning

Fine-tunes pre-trained models for new tasks using smaller datasets.

Reduces training time, data needs, and compute resources.

Example: Fine-tuned GPT or ViT models outperform scratch-trained models with 10x fewer examples.

Edge AI & TinyML

Runs models directly on low-power hardware like sensors or microcontrollers.

Minimizes latency, energy use, and data transfer costs.

Example: Wake word detection on a \$2 microcontroller with <50KB RAM.

Low-Power Hardware

Uses AI chips optimized for energy efficiency (e.g., TPUs, neuromorphic processors).

Delivers faster, greener inference at the same or lower cost.

Example: Google's Edge TPU enables 4 TOPS with <2W power draw.

Neurosymbolic & Swarm Intelligence

Combines logic-based reasoning with learning or distributed agent coordination.

Enables lightweight, explainable, and distributed intelligence.

Example: Neurosymbolic AI in robotics reduces data dependency while improving interpretability.

Sustainable Data Centers

Deploys AI workloads in facilities powered by renewables & efficient cooling.

Supports green innovation goals and long-term TCO reduction

New advancements are also brought by using AI to design energy efficient data centres.

Applying these techniques, often in combination, allows developers to build powerful AI systems that adhere to the principles of frugality.



MEASURING FRUGALITY: FRAMEWORKS & STANDARDS

Applying these techniques, often in combination, allows developers to build powerful AI systems that adhere to the principles of frugality.

Frugal AI Self-Assessment Framework: This framework allows organizations to evaluate their maturity across key dimensions of Frugal AI, including sustainable data centers, AI efficiency techniques, energy scoring, handling constraints, ROI/TCO calculation, scalability, ecosystem engagement, inclusivity, and impact focus. It provides a structured way to identify strengths and areas for improvement.

Value Clarification Matrix: This tool helps prioritize AI projects by assessing them against multiple criteria: Business Value (revenue/cost impact), Environmental Impact (resource reduction), Social Impact (accessibility/inclusion), Implementation Risk, and Total Cost of Ownership (lifetime resource needs). It encourages a holistic view beyond pure technical performance.

Total Cost of Ownership (TCO) & ROI: Frugal AI necessitates moving beyond rudimentary ROI calculations. A comprehensive TCO analysis must include not only initial development costs but also ongoing infrastructure, compute, energy, data storage, and maintenance expenses over the AI system's lifecycle. Accurately measuring TCO and linking it to tangible ROI is vital for demonstrating the economic benefits of frugal approaches.

AI Energy Scoring: This involves measuring and benchmarking the energy consumed by AI models during training and inference relative to their performance (e.g., accuracy, latency). Tools like ML CO2 impact calculators and standardized scoring systems help compare models based on their energy efficiency, guiding developers towards more sustainable choices.

Standardization (e.g., AFNOR Spec 2314): Initiatives like the French AFNOR specification provide concrete benchmarks and guidelines for evaluating the environmental impact of AI services. Such standards encourage transparency and accountability by requiring justification for using AI over less intensive alternatives, adherence to best practices throughout the lifecycle, and consideration of planetary boundaries. Further development and international adoption of such standards are key.

FRUGAL AI IN ACTION: CASE STUDIES & APPLICATIONS

The principles of Frugal AI are not merely theoretical; they are being actively applied across diverse sectors, yielding tangible results:

Startup Innovations:

Sqwish:

What they do: Compresses AI prompts to reduce input size, making GenAI faster, cheaper, and more efficient.

Why it's Frugal AI:

- Reduces computational load and carbon emissions.
- Lowers the cost barrier for running LLMs — enabling access for smaller players.
- Promotes reuse and optimization over brute-force scale.

VoltUp :

What they do: VoltUp is a battery-swapping platform for electric two- and three-wheelers. Their AI and data analytics optimize battery station placement, performance monitoring, and user demand prediction.

Why it's Frugal AI:

- Applies AI to physical infrastructure — uses data and ML to optimize battery swaps
- without needing massive grids or centralized charging.
- Frugal by design — enables low-cost, high-speed EV adoption using modular systems.
- Accessibility & Scale — Serves urban India with minimal downtime and localized insights.

Orbital Materials:

What they do: Designs advanced materials to radically improve cooling and energy use in data centers.

Why it's Frugal AI:

- Targets infrastructure-level efficiency, reducing water and power needs.
- Enables greener, more cost-effective AI deployment at scale.
- Supports sustainability goals at the systems level, not just model level.

Frugal AI Application Domains:

- **Healthcare:** Developing lightweight diagnostic tools deployable on mobile devices for disease detection (e.g., image analysis) in remote or under-resourced areas; optimizing personalized treatment recommendations using limited patient data.
- **Education:** Creating cost-effective AI tutoring systems accessible on low-end devices; using efficient models for automated grading to reduce teacher workload without high computational demands.
- **Environment:** Employing AI for optimizing energy distribution in microgrids; analyzing environmental data (e.g., pollution, deforestation) from limited sensor networks (e.g., IA.rbre project using minimal data for urban cooling analysis; Predict AI'r using mobile data for air quality insights).
- **Industry:** Building robust machine learning models for predictive maintenance or process optimization on the shop floor, even with limited or variable data quality, running on edge devices.
- **Agriculture:** Enabling precision agriculture solutions (e.g., crop monitoring, disease detection) using AI on low-power devices or analyzing limited drone/satellite imagery in remote farming communities.
- **Customer Experience:** Implementing efficient chatbots and virtual assistants (e.g., using FrugalGPT approaches) to reduce operational expenditure while maintaining performance.

These examples demonstrate the versatility and real-world potential of Frugal AI to deliver value efficiently and sustainably across numerous fields.



THE CAMBRIDGE JUDGE FRUGAL AI HUB ECOSYSTEM

The Frugal AI Hub at Cambridge Judge Business School serves as a catalyst and connector for advancing Frugal AI principles and practices. It fosters a vibrant ecosystem bringing together diverse stakeholders committed to responsible and resource-efficient AI innovation.

- **Mission & Activities:** The Hub promotes knowledge sharing, research, and collaboration through various channels. This includes an active LinkedIn Knowledge Hub for professional connections and discussion, alongside regular academic events, workshops, and seminars hosted at Cambridge Judge Business School to explore cutting-edge research and applications.
- **Interdisciplinary Ecosystem:** A key strength of the Hub is its ability to convene a wide range of voices – including investors, startups, students, academics, large and small corporations, consultants, and technology providers. This interdisciplinary approach ensures that Frugal AI development is informed by diverse perspectives and grounded in real-world needs.
- **Frugal Adoption Labs:** This flagship initiative provides a structured environment for piloting and validating Frugal AI solutions in real-world settings. Labs focus on specific domains (e.g., healthcare, sustainability, education) and match AI developers with host organizations to test prototypes under genuine constraints, applying Frugal AI frameworks to benchmark efficiency, TCO, and impact, ultimately accelerating the path from concept to scalable deployment.
- **Measure What Matters:** Apply frameworks like the Value Clarification Matrix, TCO analysis, and AI Energy Scoring within your organization.
- **Join the Community:** Participate in the LinkedIn group, attend events, share insights, and find collaborators ((<https://www.linkedin.com/groups/13123893/>)).
- **Partner With Us:** Engage in joint research, training programs, policy dialogues, or support startup initiatives through the Hub.

ADDRESSING MISCONCEPTIONS: FRUGAL AI MYTHS VS. REALITY

As Frugal AI gains traction, several misconceptions have emerged. It's crucial to distinguish the reality from the myths:

Common Myth	Strategic Reality	Takeaways
Frugal AI = Cheap, Low Quality	Frugal AI prioritizes efficient, high-impact solutions, not cutting corners.	Reframe cost-saving as value-maximization through smarter, leaner AI investments.
Low-tech can't scale	Simple tools (WhatsApp, SMS, low-code AI) power platforms used by millions globally.	Embrace context-appropriate scaling – not every solution needs enterprise-grade infrastructure.
Only big tech can do AI	AI is increasingly democratized; startups and small teams build effective, lean models.	Tap into ecosystem partnerships and pilot frugal solutions for niche or underserved markets.
Frugal = Outdated Tech	Frugal AI leverages modern techniques (edge AI, model compression, efficient models).	Integrate Frugal AI metrics (TCO, energy score) into procurement and solution architecture.

Understanding these distinctions is key to appreciating Frugal AI not as a compromise, but as a sophisticated and strategic approach to innovation.

FRUGAL AI VS. TRADITIONAL AI: A COMPARATIVE VIEW

Frugal AI represents a distinct paradigm compared to traditional, often resource-intensive AI development. Key differences include:

Feature	Frugal AI	Traditional AI
Resource Use	Minimal compute, energy, data; lightweight models	High compute, energy, data requirements; large, complex models
Sustainability	Environmentally conscious by design; lower carbon footprint	Often carbon-intensive; environmental impact a growing concern
Cost	Lower operational expenses (OpEx), reduced TCO	High infrastructure and operational costs; potentially uncertain ROI
Accessibility	Democratized; deployable on edge/mobile; lower barriers to entry	Often limited to well-resourced firms; requires significant infrastructure
Performance	Optimized for efficiency; may involve trade-offs; can excel in specific tasks	Potential for state-of-the-art results, especially on complex, large-scale tasks
Data Requirements	Focus on quality/relevance; effective with smaller datasets	Typically relies on massive datasets for training
Speed to Deploy	Faster prototyping and deployment, especially with lean models	Can involve long cycles due to infrastructure setup and complex training
Design Approach	Constraint-driven, problem-first, value-focused	Often scale-first, data-heavy, performance-maximization focused
Primary Beneficiaries	SMEs, emerging markets, resource-constrained settings, specific use cases	Large enterprises, tech hubs, research-intensive applications

CONCLUSION: TOWARDS A SUSTAINABLE & INCLUSIVE AI FUTURE

Frugal AI is more than a set of techniques; it is a necessary evolution towards a more responsible, sustainable, and accessible future for artificial intelligence. In an era defined by both the immense potential of AI and the pressing realities of resource constraints and environmental limits, the principle of "doing more with less" is paramount. By prioritizing resource efficiency, sustainability, accessibility, and measurable impact, Frugal AI offers a pathway to unlock the benefits of AI for a broader range of organizations and individuals, without incurring unsustainable environmental or economic costs.

The advantages are compelling: significant reductions in operational expenses and carbon footprint, the democratization of advanced capabilities, faster scaling of impactful solutions, and the ability to deploy AI effectively in diverse, real-world environments. While challenges exist, ongoing research, the development of new tools and frameworks, and growing industry adoption demonstrate the increasing viability and importance of this approach.

The Frugal AI Hub at Cambridge Judge Business School is committed to championing this shift. We believe that by fostering collaboration between academia, industry, and policymakers, we can accelerate the development and deployment of Frugal AI solutions that address critical global challenges. We invite researchers, innovators, investors, businesses, and public sector leaders to join us in this mission. Let us work together to ensure that the future of intelligence is not only powerful but also prudent, equitable, and sustainable for all.



REFERENCES

AFNOR Spec 2314. (2023). Evaluation Environnementale des Modèles d'Intelligence Artificielle. Association Française de Normalisation.

Cambridge Centre for Smart Infrastructure and Construction (CSIC). (2023). Low-carbon digital infrastructure. University of Cambridge.

Clever Republic. (2024). Why Frugal AI needs AI Data Governance. Retrieved from <https://www.cleverrepublic.com/blog/why-frugal-ai-needs-ai-data-governance/>

Forbes. (2023, September 8). What Large Models Cost You – There Is No Free AI Lunch. Retrieved from <https://www.forbes.com/sites/craigsmith/2023/09/08/what-large-models-cost-you--there-is-no-free-ai-lunch>

LeCun, Y. (2022). A Path Towards Autonomous Machine Intelligence. Meta AI Research. Retrieved from <https://openreview.net/pdf?id=BZ5a1r-kVsf>

Open Data Institute. (2021). Data Stewardship and Responsible Innovation. Retrieved from <https://theodi.org>

Salesforce. (2023). Introducing the AI Energy Score. Retrieved from <https://www.salesforce.com/blog/ai-energy-score>

Scribd. (2024). Afnor-spec Norme Ia | PDF | Artificial Intelligence. Retrieved from <https://es.scribd.com/document/835753369/Afnor-spec-Norme-Ia>

Smile. (2025). Frugal AI: sustainable, responsible and high-performance artificial intelligence. Retrieved from <https://smile.eu/en/publications-and-events/frugal-ai-sustainable-responsible-and-high-performance-artificial>

Teneo.Ai. (2024). Embracing Frugal AI for a 98% save in AI OpEx in Customer Experience in 2024. Retrieved from <https://www.teneo.ai/blog/embracing-frugal-ai-for-a-98-save-in-ai-opex-in-customer-experience-in-2024>

TinyML Foundation. (2022). Why TinyML is Key to Sustainable AI. Retrieved from <https://www.tinymml.org>

UKRI. (2024). Responsible AI: Programmes and Impact Stories. Retrieved from <https://www.ukri.org/what-we-do>

UNESCO. (2022). Recommendation on the Ethics of Artificial Intelligence. Retrieved from <https://unesdoc.unesco.org>

