VEDRAI OBSERVATORY

The AI that doesn’t waste energy, eecisions, and measurable value

AI energy consumption 2020–2025: growth, composition, and operational impact. Sustainability is not a declared value, it is a direct effect of governed AI.

by Vedrai Observatory
Vedrai S.p.A.

April 2026

From 2020 to 2025, the global energy consumption of artificial intelligence systems rose from roughly 65 TWh to more than 415 TWh. A sixfold increase in five years — one that transformed AI from a marginal presence in the world's energy ledger into critical infrastructure with an impact comparable to the annual energy needs of entire national grids.

This figure, taken on its own, says nothing about sustainability. It only speaks to scale. The real question is different: how much of this consumption produces better decisions? How much of it reduces real operational waste? How much translates into measurable value for those who commission it?

The answer depends entirely on the type of AI adopted and the governance framework used to deploy it.

THE IDEA IN BRIEF

THE PROBLEM
Global AI consumes six times more energy than it did in 2020. But not all of that energy produces measurable operational value. AI without governance consumes without reducing waste.

THE RESEARCH
Analysis of global AI energy consumption from 2020 to 2025, with decomposition by system type and a comparison between generic AI and decisional AI in operational contexts.

THE FINDING
AI applied to the governance of operational decisions reduces operating costs by 15–30%. Sustainability is not a declared value; it is a direct effect of adopting AI with a method.

THE ENERGY PARADOX

Between 2020 and 2025, the global AI ecosystem underwent a radical transformation. Large language models, multimodal architectures, and distributed inference systems pushed AI energy consumption to levels that no pre-2022 projection had fully anticipated.

The share attributable to generative AI alone (a marginal 2 TWh in 2020) reached 165 TWh in 2025, or 40% of the total. This expansion was fueled by the race to build foundational models, enterprise adoption of generative systems, and democratized API access.

A data center powering an LLM to answer general queries consumes the same energy as one governing the purchasing decisions of a mid-sized manufacturing company. But the second one reduces excess inventory by 22%. The first measures nothing.

At the same time, energy efficiency per query improved significantly: next-generation chips and inference framework optimizations cut per-task consumption by roughly twelve times compared to 2020. Yet total consumption grew anyway, driven by an explosive surge in request volume.

> Technical efficiency and environmental impact are not equivalent when volume grows faster than efficiency.

THE AI THAT CREATES MEASURABLE VALUE

There is a category of AI systems for which energy consumption is directly and documentably justified: AI applied to decision governance. Supply chain, procurement, production, demand planning. In these domains, every wrong decision generates physical waste: idle inventory, over-purchased raw materials, poorly calibrated production shifts, oversized stock.

When AI governs these decisions with method, the system's computational consumption translates directly into a reduction of operational waste. The relationship is causal and measurable, not declarative.

> Sustainability is not a value to insert in an ESG report. It is a direct effect of having adopted AI with a precise objective.

GENERIC AI VS. DECISIONAL AI: A COMPARISON

Dimension	Generic AI (Ungoverned)	Decisional AI
Objective	Respond to queries	Govern decisions
Measurable output	Text, images, code	Waste reduction, operational KPIs
Energy ROI	Not measurable	15–30% reduction in operating costs
Real ESG impact	Declarative	Direct, documentable effect
Decision governance	Absent	Structural, with method

THE DATA: AI CONSUMPTION 2020–2025

The historical series of AI energy consumption between 2020 and 2025 shows a compound annual growth rate above 45%. A pace without precedent in the evolution of any other technological infrastructure in recent decades.

GLOBAL AI CONSUMPTION 2020–2025 (TWh)

Year	Total (TWh)	YoY Change	of which Gen. AI	Query Efficiency (index)
2020	65	baseline	2 TWh	100
2021	95	+46.2%	6 TWh	88
2022	140	+47.4%	18 TWh	72
2023	200	+42.9%	60 TWh	54
2024	295	+47.5%	115 TWh	38
2025	415 (est.)	+40.7%	165 TWh	25

The per-query efficiency figure (rightmost column) is particularly significant: in 2025, an equivalent task consumes roughly one quarter of the energy it required in 2020. Yet total consumption grew sixfold. This is the empirical proof that technical efficiency alone is not enough. Governance of consumption is required to ensure that every watt produces value.

PRACTICAL IMPLICATIONS

Every investment in AI infrastructure carries a real energy cost. The question every business leader should ask (on Earth Day as in every board meeting) is not “are we using AI?” but “does the consumption our AI generates produce better decisions, or does it produce only output?”

> Not wasting the consumption AI generates is the real operational responsibility of business leaders in 2026.

COO — The energy cost of AI is a real operating cost

It is not enough to know that the company “uses AI.” Leaders must know whether that consumption translates into a measurable reduction of waste: inventory, procurement, production. If there is no operational KPI linked to the AI system in use, there is only an additional cost that justifies itself neither economically nor environmentally.

CEO — Decision governance is the true lever of sustainability

ESG declarations about AI must be backed by verifiable operational data. An AI that governs supply chain decisions with method reduces physical waste in a documentable way. This is a concrete competitive advantage, not a narrative for sustainability reports.

CIO — Not all AI infrastructure has the same energy ROI

The selection of AI systems must include a criterion of decisional efficiency: what is the ratio between computational consumption and operational value produced? Systems with structured decision governance maximize this ratio. Infrastructure that supports specific operational decisions produces a measurable energy ROI.

CDO — Data must feed decisions, not just models

The value of enterprise data is measured by the quality of the operational decisions it enables. A data architecture oriented toward decision intelligence transforms the energy consumption of models into measurable operational value, closing the loop between AI and genuine sustainability.

CONCLUSIONS

AI has become critical infrastructure, with an energy footprint that can no longer be ignored. But like any infrastructure, its value depends not on installed power but on how it is used and to what ends.

Ungoverned generative AI consumes energy, cognitive bandwidth, and IT budget without producing any output measurable in terms of real waste reduction. Decisional AI, applied with method to the governance of operations, turns every watt consumed into a better decision — and every better decision into less physical waste.

> Sustainability is not declared. It is governed. Now is the right moment to ask, for every enterprise AI system: how much value does every kilowatt-hour produce?

ABOUT THE RESEARCH

This analysis is produced by Vedrai Observatory, the research and thought leadership center of Vedrai S.p.A., dedicated to the intersection of artificial intelligence, operational decisions, and enterprise value. Data on AI energy consumption are derived from international research estimates (IEA, Goldman Sachs Global Investment Research, Lawrence Berkeley National Laboratory) covering the period 2020–2025. Values are estimates based on data center consumption models, computational load distribution, and global adoption trends. Per-query efficiency is calculated as an index relative to 2020, accounting for improvements in chips (from A100 to H100 and successors) and inference framework optimizations. Data on the operational impact of decisional AI are based on documented evidence from supply chain, procurement, and production planning contexts.

