Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

A Dialogue Across Three Worlds Artificial Intelligence, Sustainability, and Jainism

Prof. Saumya Sugata Dutta

VP

Executive Summary

Artificial intelligence is one of the fastest growing sources of demand for electricity on the planet right now. Global data-centre electricity consumption was estimated by the International Energy Agency to reach around 415 terawatt-hours (TWh) in 2024, rising to some 485 TWh in 2025 — a one-year increase of 17%, with AI-centred facilities alone jumping by around 50%. In the IEA's base case, total data-centre demand is projected to reach about 945 TWh by 2030 and 1,200 TWh by 2035, roughly equivalent to Japan's electricity consumption and then some. This growth is colliding with grid capacity constraints, water stress around cooling infrastructure and mounting public concern about the environmental footprint of a technology that is simultaneously reshaping economies, labour markets and daily life.

An ancient tradition almost at the same historical moment provides a curiously pertinent vocabulary for thinking about restraint, harm, and multiplicity of perspective. Jainism, one of the world's oldest continuously practised religions, structures its entire ethical system on three interlocking commitments: ahimsa (non-violence towards all forms of life), aparigraha (non-possessiveness and voluntary limitation of consumption), and anekantavada (the doctrine that truth is many-sided, and that no single viewpoint captures it completely). These are not abstract metaphysical curiosities. They are disciplines lived, honed for some 2,500 years, on how a conscious agent should act in a world of finite resources and interdependent life.

This white paper argues that these three Jain principles offer a coherent, non-arbitrary framework to assess and design sustainable AI systems – not as decoration or marketing language, but as genuine design constraints with technical analogues already emerging in the peer-reviewed literature on "Green AI" and "frugal machine learning." Ahimsa maps onto harm-reduction practices across the AI lifecycle, including energy sourcing, labour conditions in data annotation and content moderation, and the downstream social harms of biased or manipulative systems. Aparigraha is the new discipline of computational frugality; right-sizing models, minimising unnecessary training runs, resisting the industry's default assumption that bigger is better. Anekantavada aligns with epistemic humility in AI outputs, and pluralistic, multi-stakeholder governance structures that avoid interpretive authority being concentrated in a few firms or model architectures. Another thread woven through the paper as a concrete case study is the historic and ongoing dominance of Jain trading communities – above all, the Palanpuri Jains – in the world diamond and jewellery trade, an industry that is now being reconfigured by AI-driven grading, blockchain traceability, and lab-grown-diamond production, and one whose tension between aparigraha and commercial success is put into surprisingly sharp relief.

The paper is organised in seven sections. Part I quantifies the resource footprint of AI using the latest publicly available data. Part II provides an introduction to Jain metaphysics and ethics for a general audience, with some references to primary and secondary scholarly sources. Part III broadens

the three-principle mapping in technical detail, together with a proposed life-cycle framework; Part IV draws on real-world initiatives – from Jain temple solar programmes to frugal-AI research groups – as illustrative (not exhaustive) case material. Part V addresses a unique and unexplored intersection: the historic role of Jain communities in the global gems and jewellery trade, and how AI-driven grading, traceability, and lab-grown diamond technology now fit into the same ahimsa/aparigraha framework developed in Part III. Part VI is self-consciously self-critical: it foregrounds the tensions, category errors, and risks of appropriation involved in invoking a religious tradition in a secular technology-policy context—including the tension, explicitly articulated in Part V, between aparigraha and a community's historic success in a luxury trade. Part VII concludes with recommendations for technologists, policy makers, and faith communities.

This is a synthesis document, not original empirical research and does not speak on behalf of any Jain religious authority. Its aim is more modest and we hope more useful:

to show that an old ethical grammar and a new technical problem ask closely related questions, and that taking the older tradition seriously—on its own terms, not simply as a source of convenient slogans—can sharpen how we think about building AI systems that are less wasteful, less harmful, and less certain of their own infallibility.

1. The Resource Footprint of Artificial Intelligence

1.1 Why energy, water, and materials matter now

Ethics discussions around AI have tended to focus on issues such as bias, privacy, misinformation, and labour displacement. Environmental impact entered the mainstream discussion later on, mostly because the numbers became too big to ignore. But the sheer scale of capital being committed to AI infrastructure is in itself a signal: the five largest hyperscale cloud providers committed roughly $660-690 billion in AI-related capital expenditure in 2026 alone, a concentration of private infrastructure investment without precedent outside wartime economies.

This investment’s physical footprint is primarily measured in three dimensions: electricity, water, and embodied material (the silicon, rare-earth elements, and rare metals that make up chips and servers). Of the three, electricity has the best public data and is the focus of this section, although water and materials are briefly addressed at the end.

1.2 Electricity demand: current trajectory

In its 2025 special report Energy and AI and its April 2026 follow-up Key Questions on Energy and AI, the IEA estimates that the electricity demand of global data-centres was around 415 TWh in 2024, or about 1.5% of global electricity use. That number had already been growing at around 12% a year since 2017—more than four times the rate of growth in overall global electricity demand. In 2025, growth picked up again, with data-centre electricity demand increasing by 17% year-on-year to around 485 TWh, and electricity consumption in AI-specific facilities rising by around 50%.

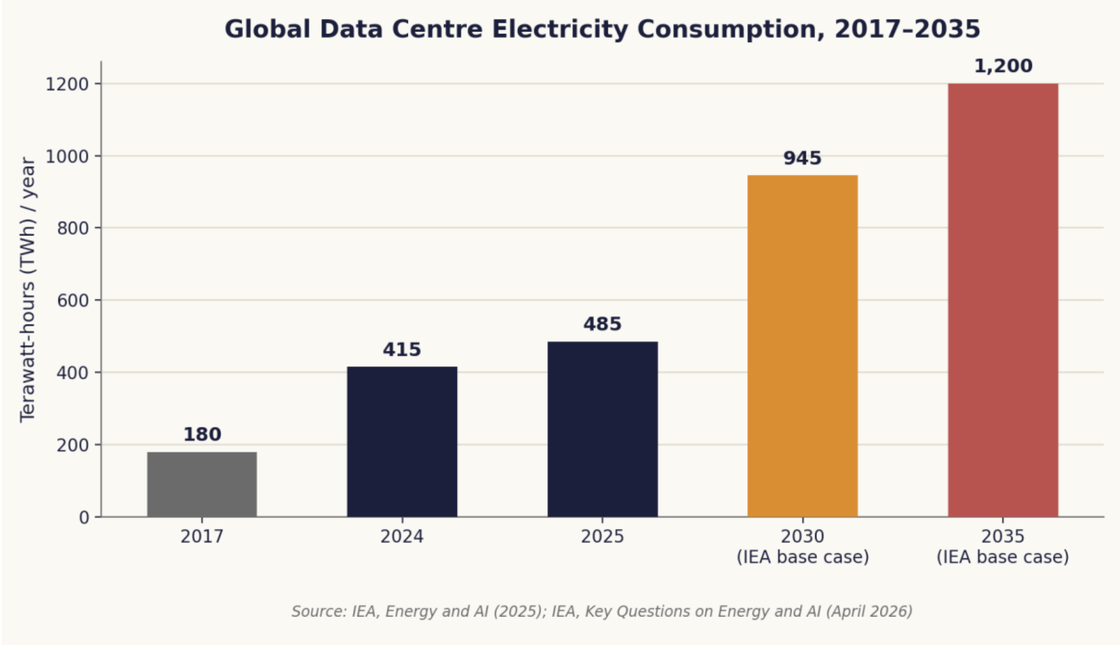

Figure 1. Global data-centre electricity consumption, historical and projected. Source: IEA, Energy and AI (2025); IEA, Key Questions on Energy and AI (April 2026).

Under the IEA's base-case scenario, the world's electricity consumption for data centres is projected to almost double to some 945 TWh by 2030, a little more than Japan's entire current electricity consumption, and to rise further to about 1,200 TWh by 2035. The United States accounted for the largest share of consumption in 2024 (45%), followed by China (25%) and Europe (15%). In the United States in particular, the U.S. Department of Energy's Lawrence Berkeley National Laboratory estimated data-center electricity use to increase from 176 TWh in 2023 to somewhere between 325 TWh and 580 TWh in 2028, which would be data centers making up 6.7% to 12.0% of total U.S. electricity consumption.

It is worth being precise about what is driving this growth, because the popular narrative sometimes confuses all digital infrastructure with AI specifically. The IEA's modelling suggests that "accelerated servers" – the GPU- and TPU-dense hardware mainly used for AI training and inference – will account for nearly half of the net increase in data-centre electricity use to 2030, with their electricity use expected to grow around 30% a year, versus roughly 9% annual growth for conventional (non-AI) servers. In other words, AI is not the sole reason for the growth of data centres, but it is by far the fastest-growing reason.

1.3 The energy cost of a single query

One comparison often made that does have real value is the electricity cost of a single AI interaction versus a traditional web search. Early, often-cited estimates suggested that a single advanced generative-AI query in 2024 consumed around 2.9 watt-hours (Wh) of electricity, compared with about 0.3 Wh for a standard search engine query – a difference of almost tenfold. More recent measurement work cited by the IEA in 2026 points to efficiency gains in model architecture and serving infrastructure bringing the median energy cost for typical short text queries down to somewhere in the region of 0.24–0.3 Wh, roughly converging with a conventional search.

Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

But that convergence is not the end of the story. The same IEA analysis highlights that newer, more energy-intensive application types – video generation, extended reasoning, and multi-step “agentic” tasks that call on tools and chain together many model calls – can burn hundreds or even thousands of times more energy per query than a simple text response. In other words, the average cost of a text query has been dropping, while the ceiling for the most demanding tasks has been rising sharply. This bifurcation is at the heart of any meaningful sustainability strategy for AI: efficiency gains at the low end can be overwhelmed by growth in usage of the most energy-intensive capabilities at the high end, a phenomenon sometimes known as a rebound or Jevons-paradox effect

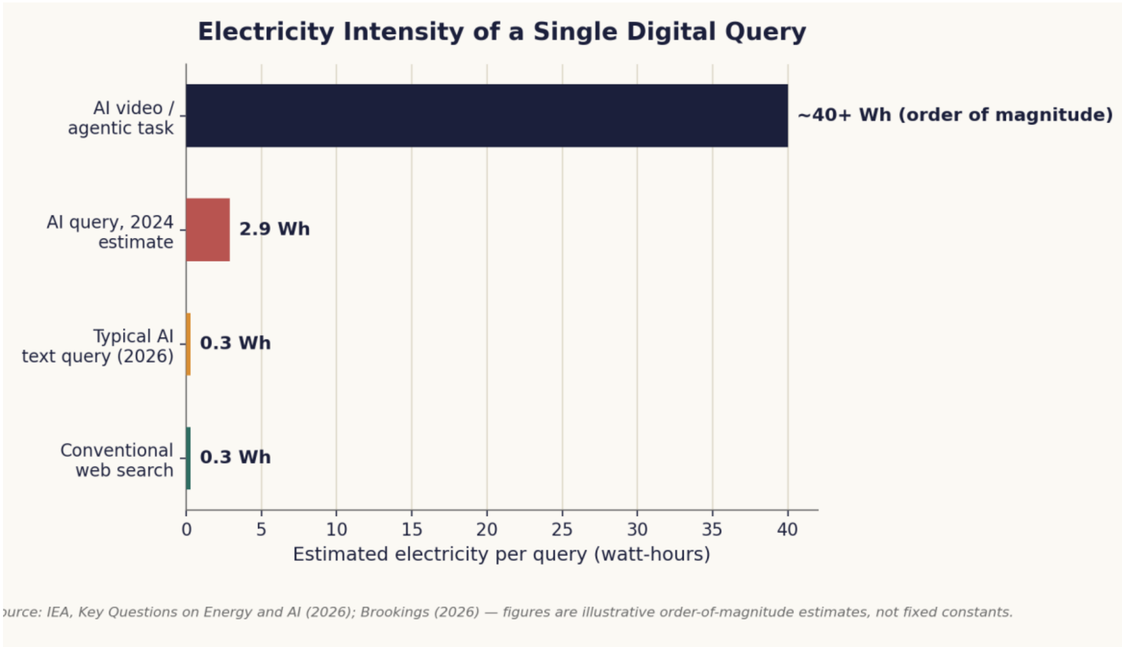

Figure 2. Illustrative order-of-magnitude electricity intensity across query types. Source: IEA, Key Questions on Energy and AI (2026); Brookings Institution (2026).

1.4 Water and embodied material

Cooling is still one of the largest non-compute electricity loads in a data centre based on IEA breakdowns of data-centre energy use, accounting for anywhere from around 7% of total consumption in efficient hyperscale facilities to over 30% in less efficient enterprise data centres. In the past, many facilities used evaporative cooling, which uses fresh water directly. Industry analysis shows that liquid, closed-loop cooling systems—which are increasingly required for the power density of modern AI accelerators—can reduce direct on-site water consumption by 70–90% compared with evaporative systems. However, liquid cooling does not eliminate the indirect water footprint associated with electricity generation itself, and it increases upfront capital costs.

The embodied material impacts – the mining, refining and manufacturing footprint of semiconductor supply chains – are relatively under-measured in public reporting compared with operational electricity, and are not modelled in detail in the IEA’s public analyses. This paper points to this as an area where transparency lags far behind the electricity conversation, and returns to it in Part VI.

1.5 Grid strain and the limits of efficiency alone

The IEA's 2026 report zeroes in on the idea that the AI energy problem is no longer just a matter of how much electricity a model consumes but increasingly about how fast new electricity supply, transmission infrastructure and grid connections can be built. An estimated 20% of planned data-centre projects are at risk of delay due to grid bottlenecks In advanced economies, transmission lines can take four to eight years to build, and lead times for critical components such as transformers and gas turbines have increased dramatically. In certain regional markets, such as Ireland, Northern Virginia, areas of Texas and Phoenix, data centres already represent a double-digit percentage of local electricity demand and the Irish data centre sector alone is said to represent more than 20% of national electricity consumption.

This is important for this paper's argument because it shifts sustainability from a purely technical optimisation problem ("make each model more efficient") into also a question of restraint and prioritisation ("should this capability be built, and at what scale, given what it displaces or strains?"). It is in this second, harder question that an ethical tradition based on restraint rather than a purely technical one based on optimisation has something distinctive to offer. This is the subject of Parts II and III.

2. Jainism: An Ecological Philosophy Older Than the Problem It Now Speaks To

2.1 A brief orientation

Jainism is one of the world's oldest living religious and philosophical traditions, and traces its lineage through a succession of twenty-four Tirthankaras ("ford-makers"). The last of these, Mahavira, lived in the sixth century BCE, and is often considered a contemporary systematiser of much older teachings associated with his predecessor, Rishabha (Adinatha). Today there are some four to six million Jains worldwide, the vast majority in India but with a well-organised diaspora in North America, the United Kingdom and East Africa. Despite this relatively small global population, however, Jain ethical vocabulary, not least of all ahimsa, has had outsized historical influence, most visibly through Mahatma Gandhi whose commitment to non-violent civil resistance was shaped in significant part through his friendship and correspondence with the Jain scholar Shrimad Rajchandra.

According to Jain metaphysics, the universe contains an infinite number of souls (jiva) in different states of embodiment, from animals down to the simplest organisms that Jain cosmology says inhabit earth, water, fire, air and plant life — what the tradition calls ekendriya, or one-sensed beings. This is not a metaphor for ecological interdependence, but a literal statement about the moral status of the physical world. The practical result is a religion whose central ethical discipline is not simple personal virtue, but a structured relation to the whole of the non-human world.

2.2 Ahimsa: non-violence as a total ethic

> There is nothing so small and subtle as the atom nor any element so vast as space. Similarly, there is no quality of soul more subtle than non-violence and no virtue of spirit greater than reverence for life. — Attributed to Mahavira, in the Jain Declaration on Nature (1990)

Ahimsa is the first, and in the traditional formulation, the pivotal vow among the five great vows (mahavratas) taken by Jain monastics and the more moderate vows (anuvratas) taken by lay householders. The Jain teaching importantly extends ahimsa to intention and not just to action: indirect harm, or the will to harm, carries moral weight even if no direct injurious act has occurred. Mendicant practice makes this concrete in ways that can seem extreme to an outside observer – some Jain monks and nuns sweep the ground before them to avoid stepping on insects, or wear cloth over the mouth to avoid inhaling small organisms – but the underlying logic is simply ahimsa taken to its most rigorous, uncompromising conclusion.

This has a direct ecological content. In the Jain Declaration on Nature, written by L. M. Singhvi (then Indian High Commissioner to the United Kingdom) and handed over to the World Wide Fund for Nature in 1990, ahimsa is explicitly spelt out as an ecological principle: elements of nature are described as living and to be protected from waste, overuse, abuse and pollution. A much-quoted dictum of Mahavira effectively states that one is what one would harm, thus destroying the difference between self-interest and the interest of the natural world upon which self depends.

The Federation of Jain Associations in North America (JAINA), through its Ahimsak Eco-Vegan Committee, published a supplemental Jain Declaration on the Climate Crisis in October 2019, updating the 1990 declaration with current climate science. This statement extends the logic of ahimsa to fossil-fuel consumption and industrial animal agriculture, and explicitly states that the karmic responsibility for climate harm extends to indirect contributors — including, notably, those who drive fossil-fuel vehicles or fly frequently, on the logic that indirect causation of harm (by means of greenhouse-gas emissions linked to wildfires and extreme weather) still falls within the scope of ahimsa. This reasoning has led some Jain temples in North America to install solar panels as a direct expression of this reasoning.

2.3 Aparigraha: non-possessiveness and voluntary limitation

Aparigraha (non-possessiveness) is the vow of non-acquisitiveness and non-accumulation of excess. In much Jain environmental writing it is paired with ahimsa, for the causal chain is direct: to have more than one needs is usually to take more than is necessary from the natural world, which in turn harms more of the living beings that are understood to inhabit that world. Jain commentators often cite Gandhi's own formulation – that the earth offers enough to meet each person's need, but not each person's greed – as a plain-language restatement of aparigraha. But the aphorism's attribution to Gandhi himself is disputed among scholars and it is best treated as broadly Gandhian in spirit, rather than as a verified verbatim quotation.

Aparigraha is not just a prohibition of luxury, but an active discipline of reduction. Lay Jains have traditionally put it into practice through voluntary limits on their possessions, conscious simplicity in their diet and dress, and through set periods of fasting. In its full expression it is practiced by Jain ascetics who usually have nothing but a robe, an alms bowl, and instruments for the observance of ahimsa in the course of daily movement. The 2019 JAINA climate declaration explicitly broadens aparigraha to include consumption choices with climate implications: energy choices, transport modes, and the material footprint of daily life.

2.4 Anekantavada: the many-sidedness of truth

Anekantavada literally 'non-one-sidedness' is the epistemological contribution of Jainism and arguably its most philosophically distinctive contribution. It is built on the idea that reality is complex and multi-faceted and that any single view or proposition can only capture a slice of the truth. The classic example is the parable of the blind men and the elephant, where each touches a different part of the animal and each comes up with a different, partial, and locally reasonable description -- trunk like a snake, leg like a pillar, ear like a fan. Nobody is wrong in what they see, everybody is incomplete.

Anekantavada is philosophically formalised by two associated doctrines: syadvada, the doctrine of conditional predication (which states that any statement about reality is true only from a particular standpoint and should be qualified accordingly) and nayavada, the doctrine of partial standpoints. Jain scholars like L. M. Singhvi, in the Jain Declaration on Nature, explicitly connect anekantavada to the idea of universal interdependence: because no single perspective reveals the whole truth, ethical and ecological reasoning must be open to correction from other perspectives, other life forms, and other ways of knowing. It is worth noting, in the interest of intellectual honesty, that the relation between Jain ecological teaching and contemporary environmentalism is itself a subject of scholarly debate rather than settled consensus. Non-Jain scholars of Jainism cited in Peter Flügel's 2005 review Jainism and Ecology such as John Cort and Paul Dundas have questioned the degree to which the tradition's historically monastic ethical system maps directly onto modern environmental politics, cautioning against reading contemporary ecological categories back into a tradition whose primary concern was individual liberation (moksha) rather than planetary stewardship per se. This paper takes such academic caution seriously and deals with it head-on in Part V.

3. Three Principles, Three Design Practices: A Proposed Mapping

Part I determined the size of AI's physical footprint; Part II presented Jain ethics on its own terms. This section connects the two, proposing, as an original synthesis for this white paper, rather than a claim about historical Jain doctrine, a structured mapping between the three principles and specific, already-emerging practices in sustainable AI research and governance.

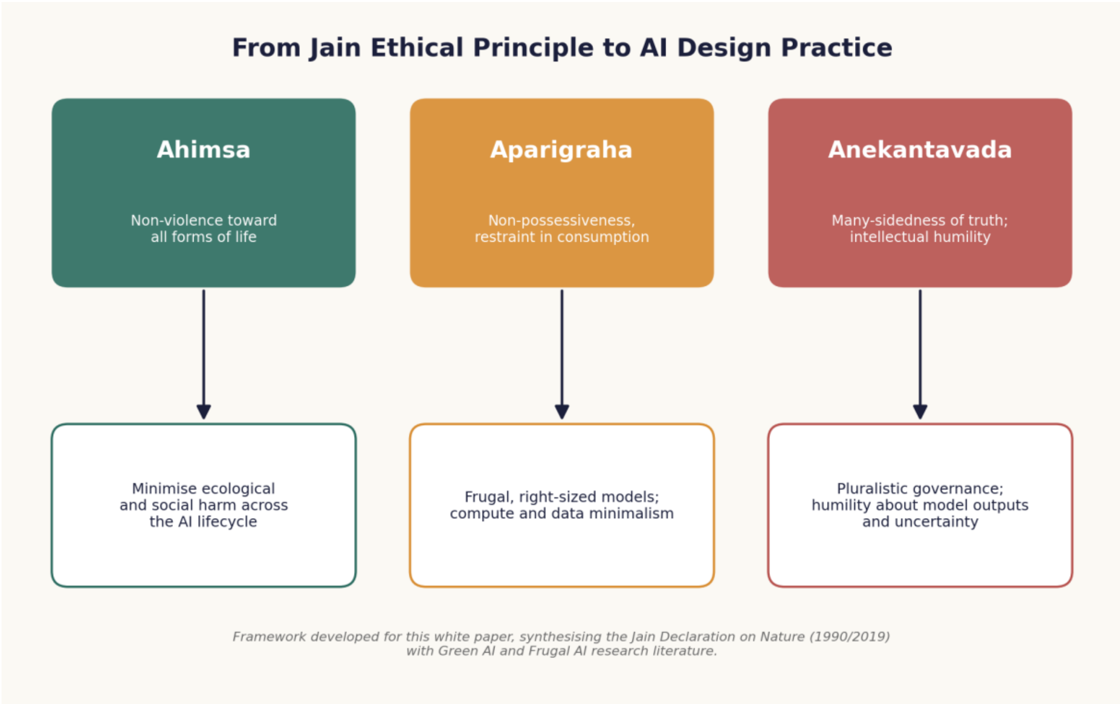

Figure 3. Proposed mapping from Jain ethical principle to AI design practice.

3.1 Ahimsa → harm-reduction across the AI lifecycle

Interpreted as a design constraint rather than a religious injunction, ahimsa poses a simple question, repeatedly, at each stage in a system's life: what harm does this action cause, directly or indirectly, and can that harm be reduced or avoided? Applied to AI, this reasoning traverses at least four different types of harm, often discussed separately, but sharing a common ethical structure.

● Environmental damage: the footprint of electricity, water and materials in Part I. The JAINA climate declaration's reasoning about indirect contributors to emissions having some moral responsibility maps directly onto modern Scope 2 and Scope 3 emissions accounting for AI infrastructure — the emissions a company causes indirectly through its energy purchases and supply chain, not just what it emits directly.

● Labour harm: the human labour of data annotation and content moderation, which underpins most modern AI systems and is often outsourced. Numerous investigative reports and academic studies have linked this work to low pay and psychological strain in parts of the global data-labelling supply chain ..

● Epistemic and social harm: downstream effects of biased, manipulative or deceptive AI outputs on individuals and communities. These are directly addressed in UNESCO's 2021 Recommendation on the Ethics of Artificial Intelligence under its "Proportionality and Do No Harm" and "Fairness and Non-Discrimination" principles.

● Non-human harm: a category almost missing from mainstream AI ethics frameworks but central to Jain reasoning – the effects of AI-driven land use, resource extraction and industrial agriculture optimisation on non-human animal welfare and ecosystems.

A version of the first category is operationalised in the research literature on 'Green AI', a term formalised by Schwartz et al. and since substantially developed.

In green AI research, a distinction is made between what is sometimes called "green-by-AI" (using AI to improve environmental outcomes elsewhere, e.g., smart-grid balancing and renewable-generation forecasting) and "green-in-AI" (making the AI systems themselves less resource-intensive to build and run). A 2024 review in the journal Neurocomputing frames green-in-AI around energy-efficient algorithm design, hardware choices and measurement tooling that allows researchers to see and reduce the energy cost of their own work — secular technical language echoing the Jain emphasis on awareness of the consequences of one's actions as a precondition for reducing harm.

3.2 Aparigraha → computational frugality

Aparigraha's ecological logic — that acquiring and consuming more than actual need causes avoidable harm — has a close technical cousin in what researchers increasingly refer to as "frugal machine learning" or "frugal AI." This growing body of work, evidenced by a number of papers and surveys in 2025–2026, explicitly reframes efficiency as more than a cost-saving measure, but as an ethical and accessibility commitment: models that require less data, less compute and less energy are more accessible to researchers and organisations without hyperscale infrastructure, echoing aparigraha's traditional emphasis on restraint as something that benefits the wider community, not just the individual practising it.

Frugal AI researchers often group techniques into three categories that align with the aparigraha framing of restraint used at different stages of a practice:

Frugality category	Representative techniques	Aparigraha analogue
Input frugality	Data selection, active learning, synthetic-data substitution	Consuming only what is genuinely needed
Process frugality	Early stopping, efficient hyperparameter search, transfer learning	Restraint in the means used to reach a goal
Model frugality	Pruning, quantisation, knowledge distillation, compact architectures	Owning only what serves the purpose, nothing in excess

In the words of Emma Strubell et al's 2021 paper that brought the carbon cost of training large NLP models to the general public, and Roy Schwartz et al's Green AI paper, the case for a research culture that does not take "bigger is better" as a default, but as a claim to be justified. In Jain vocabulary, this is simply aparigraha applied to a research community's incentive structure. A 2026 arxiv paper on Green AI architectures for circular economies reported that a multi-layered resource-conscious framework achieved a 25% reduction in energy consumption during real-world workflows and an 18% improvement in resource-recovery efficiency when applied to lithium-ion battery recycling and urban waste-management datasets – a concrete illustration that restraint-oriented design need not come at the cost of usefulness.

3.3 Anekantavada → epistemic humility and plural governance

Of the three principles, anekantavada is the least obviously 'environmental' and the most immediately relevant to AI governance and epistemics. There are two stand out application.

Output-level humility

Large language models produce text that sounds fluent and confident even when they are not certain about what they are saying. This property is well known, and leads to overreliance on and propagation of "hallucinated" claims. Anekantavada provides a philosophically coherent rationale — independent of purely engineering-driven calibration work — for viewing any given model output as one partial viewpoint among many possible ones, never as a complete or final answer to a complex question. This is not a metaphorical borrowing: syadvada, the related Jain doctrine of conditional predication, is a true centuries-old formal system for qualifying statements according to the standpoint from which they are made, and a small but growing body of interdisciplinary scholarship has begun to explore its relevance to the representation of uncertainty in computational systems.

Governance-level pluralism

At the institutional rather than output level, anekantavada bolsters an argument against the concentration of interpretive and design authority over AI systems in a few firms, countries, or model architectures. The UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), the first global standard-setting instrument on AI ethics, adopted by acclamation by 193 member states, includes sustainability as one of its ten core principles, defines it explicitly with reference to the UN Sustainable Development Goals and couples it with principles of human oversight, transparency and multi-stakeholder participation. The structural logic of that pairing is the logic of anekantavada: no single actor's standpoint of what AI should do is sufficient by itself, so governance must be constructed to incorporate multiple, sometimes conflicting, perspectives rather than optimise for one.

3.4 A proposed lifecycle framework

The paper proposes a four-stage lifecycle framework that combines the three mappings to apply a Jain-informed sustainability practice to AI systems. It is offered as a heuristic for practitioners and policymakers, not as a certification standard or a claim of religious authority.

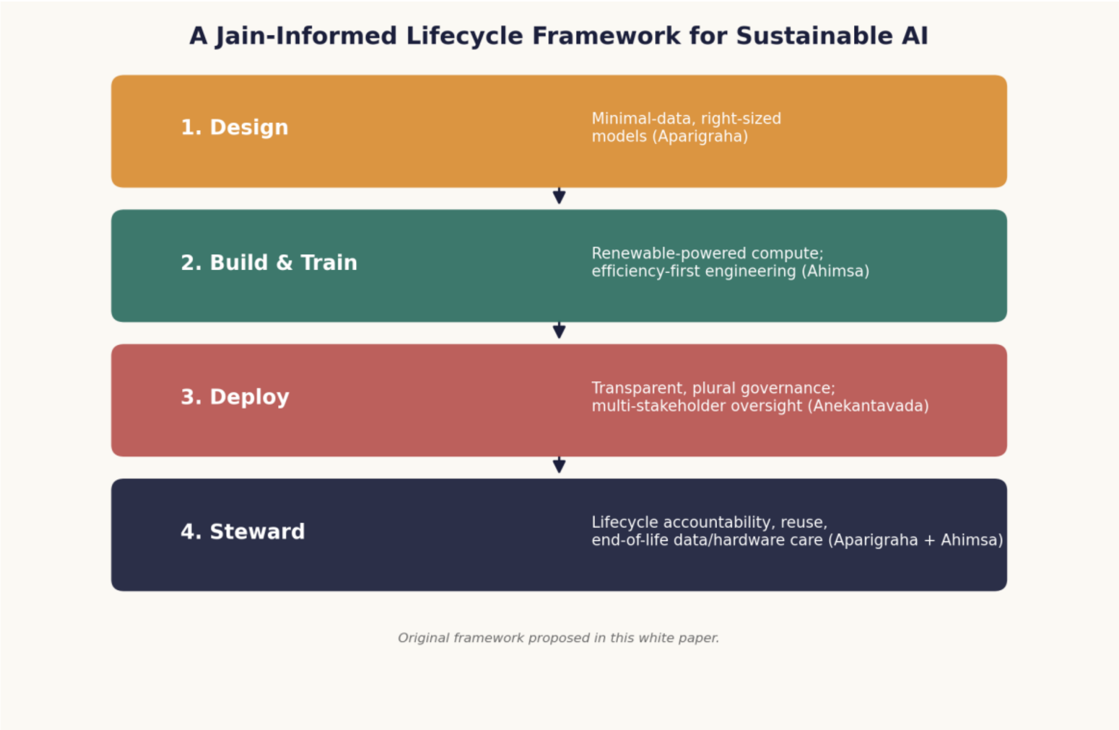

Figure 4. A Jain-informed lifecycle framework for sustainable AI, proposed in this white paper.

When designing, aparigraha recommends starting with the smallest, most data-efficient model that can perform the task, rather than automatically choosing the largest available foundation model. In the build-and-train phase, ahimsa argues for the procurement of renewable-powered compute and for the measurement and disclosure of the energy cost of training runs before they are undertaken and not just post hoc. At the point of deployment, anekantavada requires governance structures — advisory boards, red-teaming, external audit — that deliberately draw in perspectives from outside the developing organisation. At the stewardship phase, ahimsa and aparigraha together mandate lifecycle accountability: planning for model retirement, hardware reuse or recycling, and data deletion, rather than viewing deployment as the end of the ethical story.

4. Illustrative Case Material

The following examples are illustrative, rather than exhaustive, and are included to show that the mapping proposed in Part III is not purely theoretical – related practices already exist independently in both the Jain community and the AI research community, even where the two have not been in direct dialogue.

4.1 Jain community climate action

The clearest institutional articulation of Jain environmental reasoning on a contemporary issue is the 2019 Declaration on the Climate Crisis of the Federation of Jain Associations in North America. In addition to its doctrinal justification, the declaration calls for concrete action: advocacy for a plant-based diet on ahimsa grounds (given the reported emissions and animal welfare consequences of industrial dairy and meat production), consideration of choices regarding energy and transportation, and reduced use of chemical fertilisers and pesticides. Some Jain temples in North America have installed solar panel arrays, and temple leadership has explicitly framed this as a practical application of ahimsa to energy sourcing, rather than simply an economic decision..

4.2 The frugal-AI and Green AI research communities

A body of research has converged on the position that computational restraint is an environmental and an accessibility imperative, independent of any religious framing. This includes the 2025 Frugal Machine Learning survey by Violos et al., a 2025 SSRN paper on frugal approaches to scalable and green AI, and a 2025 SSRN paper by Lala Bouali, Ben Ahmed, Bradai, and Mazouzi that proposes a refined definition of frugal AI. They argue that the resource intensity of AI development favours well-funded organisations, excluding researchers and by extension whole regions and institutions, lacking access to hyperscale compute. This is a democratising argument that closely mirrors aparigraha's traditional claim that restraint by the resource-rich benefits the broader community, not just the individual practising it.

4.3 UNESCO's sustainability principle in practice

The UNESCO Recommendation on the Ethics of Artificial Intelligence (2021) operationalises its sustainability principle by requiring that AI technologies be evaluated by their impact on an evolving set of goals including the UN Sustainable Development Goals rather than a fixed, narrow environmental metric. This design choice is the embodiment of the central claim of anekantavada that no one, fixed perspective (say, carbon emissions) can provide the complete ethical story and that frameworks of assessment should accommodate the incorporation of additional, changing perspectives on what sustainability requires.

5. Gems, Jewellery, and the Jain Trade Tradition: AI, Ahimsa, and the Diamond Economy

No discussion of Jainism and material life would be complete without a discussion of an apparent paradox. For more than a century a tradition with a central vow of non-possession has produced some of the world's most successful traders in diamonds, gemstones and fine jewellery. This piece takes that paradox seriously, traces its roots, and considers how AI is now reshaping the very industry Jain communities helped build—via automated grading, blockchain-based provenance, and lab-grown diamond production that raises its own ahimsa and aparigraha questions..

5.1 Why a non-possessive tradition became a trading tradition

This link is not coincidence nor contradiction but an application of strict ahimsa to the choice of occupation. Traditional Jain teachings prohibit occupations that cause large-scale, direct harm to living beings. In the past this has meant that agriculture is out, because the classical Jain analysis is that tilling the soil inevitably kills the innumerable one-sensed organisms (ekendriya) thought to live in earth and water. Jain communities, historically excluded from the principal occupations of pre-modern South Asian economies, have focused instead on trade, finance and commerce, occupations which, on this reasoning, involve comparatively less direct violence to living beings. This category includes naturally gem cutting, diamond polishing and jewellery trading: all of these rely on inert, already extracted materials rather than the ongoing, repeated harm of farming or animal husbandry.

Figure 5. The occupational logic connecting ahimsa to Jain trading communities' historic concentration in gems and commerce.

5.2 The Palanpuri Jains and the making of the modern diamond trade

The best example of this pattern today is to be found among the Palanpuri Jains, named after their town of origin, Palanpur, in the Banaskantha district of Gujarat. For generations, Palanpuri Jains were merchants, bankers and trusted administrators to the Nawabs of Palanpur. They began to journey to the colonial cities of Bombay, Madras and Calcutta and establish trading businesses. By the 1920s there was a documented trade in rough and polished diamonds between Bombay and Antwerp, still and then the world's largest diamond trading centre.

The community was predominantly built in the 1960s. When the Indian diamond business sank after the India-China war of 1962, Palanpuri Jain traders looked to Antwerp, where they found an almost ignored niche: small, low-value rough diamonds and diamond ‘dust’ that the then Jewish-dominated Antwerp trade was not interested in cutting. The smaller stones were shipped by Jain traders to skilled, low-cost cutters and polishers in Surat and Navsari, Gujarat, and then re-exported the finished stones, polishing, in the words of one trader interviewed by the Economic Times, “in rupees and selling in dollars.” The model reached huge sizes. By the early 2020s, of every 100 diamonds traded through Antwerp, some 93 were cut and polished in India, a reversal of a trade relationship that was entirely the other way around a century earlier, according to industry sources quoted by the Gateway House research institute in its history of the Bombay-Antwerp trade.

Today, Palanpuri Jains, and the related Katiawadi Patel community, are said to dominate Antwerp’s diamond exchanges, five of the six elected representatives to the 2012 elections to the Antwerp World

Diamond Centre board were Gujarati. India's Surat-based diamond cutting and polishing industry is estimated to employ around 500,000 people in Gujarat alone. According to scholarship in outlets such as The Conversation and the Indian Diaspora research network, it is primarily trust-based, reputation-driven business networks that have made the community successful. Diamonds are portable, easy to conceal and difficult to resolve in dispute through state courts, so the industry has traditionally been governed by close-knit community and family networks enforcing informal contracts through reputational consequence rather than through litigation. Some scholars attribute this governance structure to the ethical seriousness traditionally cultivated within Jain merchant communities .

5.3 AI enters the gem trade: grading, authentication, and provenance

Historically, the diamond and jewellery industry has relied on human expert judgement to grade and authenticate diamonds and jewellery, but there has been a significant shift towards AI-assisted methods in the past few years. In July 2020, the Gemological Institute of America (GIA), the industry's most widely-recognised grading authority, announced a partnership with IBM to use AI for diamond clarity grading, with the explicit aim of allowing faster, more consistent grading and ultimately on-site automated grading outside of GIA's own laboratories. Recent reporting from the lab-grown diamond sector reveals how AI systems are now monitoring and controlling temperature, pressure and chemical composition in real time in both High Pressure High Temperature (HPHT) and Chemical Vapour Deposition (CVD) lab-grown diamond production, with improvements in consistency and, producers claim, reduced energy waste in production.

Another major trend has been the adoption of blockchain-based traceability, often combined with AI-powered anomaly detection, said to be a way to give the consumer mine-to-market visibility of a stone's provenance, certification history and chain of custody. According to industry market research summarised in 2026 diamond-market reporting, the combination of AI grading, digital certification and blockchain traceability is now central to how both mined and lab-grown diamonds are competing for ever more sustainability-conscious buyers, especially among younger consumers.

5.4 The lab-grown diamond question: an ahimsa dilemma with no clean answer

Lab grown diamonds are chemically, physically and optically identical to mined stones and graded by the same institutions (GIA, IGI, GCAL) using the same cut, colour, clarity and carat criteria. They have grown from a marginal curiosity to, according to some estimates of the 2026 market, more than half of engagement-ring diamond sales in some markets. Their appeal is largely based on an implicit ahimsa argument. According to one frequently cited industry estimate, approximately 250 tonnes of earth must be moved to produce each carat of mined diamonds, and about 160 kilograms of greenhouse gases are emitted for every polished carat, which has an impact on ecosystems, watersheds and, in some documented historical cases, entire communities living near mine sites. Lab-grown production does not involve the same land disruption and eliminates the lingering, albeit reduced, risk of "conflict diamonds" that the UN-backed Kimberley Process Certification Scheme has only partially addressed since 2003.

But the ahimsa accounting is in fact more complicated than the marketing indicates, and this paper takes that complication as an illustration of anekantavada rather than a problem to be glossed over. Both HPHT and CVD lab-growth methods are energy-intensive, and a 2024 industry analysis found that more than 60% of lab-grown diamonds are produced in regions, notably China and India, where the electricity grid is largely coal-powered, meaning a coal-grid lab-grown stone can have a larger carbon footprint per carat than a mined one, even as a solar-powered lab-grown stone from a producer like Diamond Foundry can approach carbon neutrality. Or, put another way, "lab-grown" isn't a single environmental category anymore than "mined" is: the actual damage depends on the specific energy source behind a specific rock, a fact that can get lost in industry marketing on both sides.

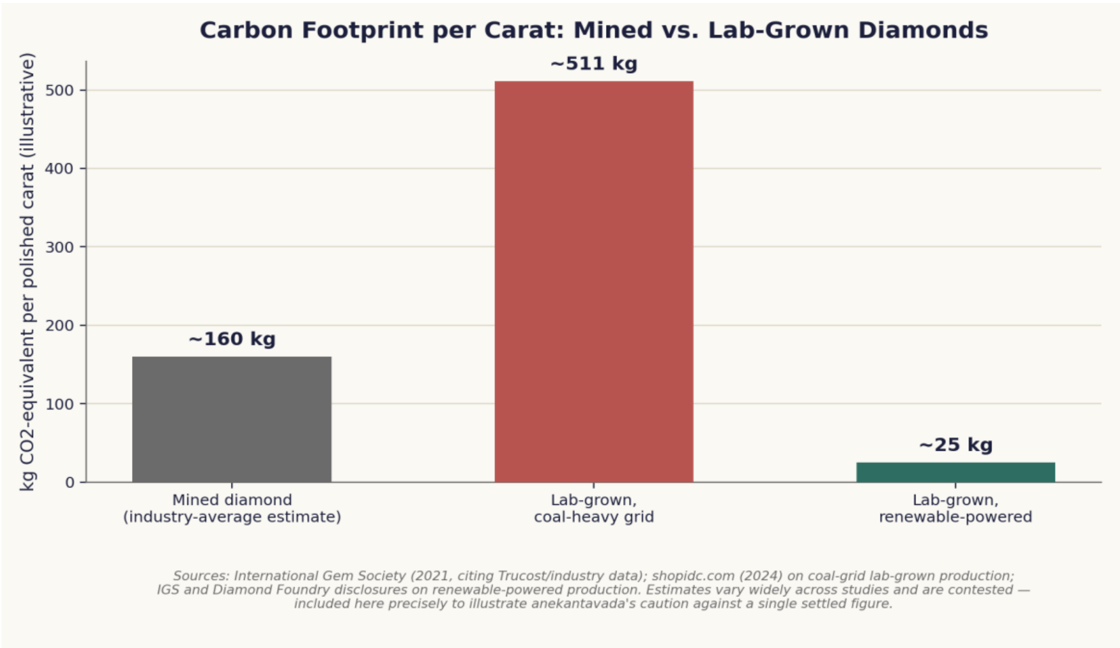

Figure 6. Contested per-carat carbon estimates across production methods, presented deliberately as a range rather than a single figure — consistent with anekantavada's caution against one-sided claims to settled truth.

5.5 Aparigraha and the trade: an honest tension, not a resolution

This paper cannot tell the story of Jain involvement in the diamond and jewellery trade as a simple case of aparigraha, intellectually honest, and if anything this is the most acute internal tension this paper faces. A commitment to non-possession and restrained consumption does not sit easily with a community’s long and well-documented record of success in a global luxury trade, largely based on accumulation, high margins and, at the consumer end, discretionary adornment – arguably among the least materially necessary categories of consumption there is.

The tension, though not entirely eliminated, is at least partially explained in Jain ethical writing as understood by the tradition itself. First, aparigraha is traditionally understood as an internal discipline of non-attachment, not as a flat prohibition on commerce or wealth generation as such—the vow constrains one’s relationship to possessions (freedom from grasping, willingness to give, absence of ostentation) more than it constrains participation in trade per se, which is why the tradition has historically paired successful life as a lay merchant with extensive religious giving, temple endowment, and community charity, a pattern well documented among Palanpuri Jain trading families. Second, the commodities being traded – cut and polished diamonds – are in the strict Jain analysis of harm a step removed from the traders themselves: the ecological and labour harms of mining are upstream, mostly beyond the direct control of Jain communities. This is a separate moral position from that of causing that harm directly, even if (as the

2019 JAINA climate declaration argues in the adjacent case of fossil fuels) indirect contribution to harm is not thereby exonerated from all responsibility.

This paper does not resolve that internal debate. Rather it notes that the tension is the sort of multi-sided question that anekantavada actually tells us to leave open rather than to resolve too quickly — and that AI-powered traceability, if well-generated and maintained, offers a pragmatic way to bridge the gap between a trader's formal distance from upstream harm and real, provable knowledge of that harm, which is the precondition for aparigraha and ahimsa to be more than good intentions.

5.6 AI, blockchain, and anekantavada-consistent supply chains

This case study's most positive contribution to the framework of Part III is precisely to anekantavada. The Kimberley Process has been credited with reducing the trade in conflict diamonds since 2003, but is widely acknowledged – even by ethical-jewellery retailers themselves – to leave meaningful gaps, as its certification scope is narrower than the full range of ethical and environmental concerns buyers now raise. In effect, a single-source, single-authority scheme of certification is a one-sided view of a many-sided problem, the same structural weakness anekantavada finds in any claim to complete, final truth.

The multi-party blockchain traceability enabled by AI indicates a structurally different model: distributed verification among miners, cutters, certifiers and retailers, cross-checked by machine-learning anomaly detection, rather than any single actor's say-so. This is not a solved problem – industry reporting in 2026 still describes blockchain traceability in the diamond sector as an emerging trend rather than a mature standard, and ironically, lab-grown diamonds still lack the standardised, independently verified environmental certification that would allow a buyer to actually compare a coal-grid stone to a solar-grid one. But the direction of travel – plural, cross-verified provenance rather than a single certifying voice – is the clearest concrete example in this whole paper of anekantavada's abstract epistemic caution becoming an actual piece of supply-chain infrastructure.

6. Tensions, Limitations, and the Risk of Appropriation

Any paper proposing such a synthesis ought, however, to be intellectually honest enough to make explicit where the analogy overreaches, where it threatens to distort the source tradition, and where it is not to be over-claimed. It highlights a further tension raised in the discussion of the diamond trade in Part V.

6.1 Jainism's core concern is liberation, not sustainability

As discussed in Part II, scholars like John Cort and Paul Dundas have warned against projecting contemporary environmentalist categories back into Jain doctrine. The central soteriological goal of the tradition is the liberation of the individual soul (moksha) from the cycle of rebirth through the shedding of karma, primarily achieved through ascetic practice. On this more guarded reading, ecological benefit is best seen as a consequence of ahimsa rigorously practiced for its own reason—the moral status of other living beings—rather than as the primary purpose of the doctrine. Hence, the mapping in this paper should be interpreted as an application of Jain ethical reasoning to a contemporary problem that the tradition did not originally address in its own terms, rather than a claim that Jain scripture anticipated data-centre electricity policy.

6.2 Monastic ahimsa and aparigraha are not directly transferable to industrial technology

The most intense Jain practice of ahimsa and aparigraha is exhibited by a small monastic population living under vows of extreme material renunciation – a life structured around minimising all possible harm and possession, including in some traditions the renunciation of electricity and mechanised transport altogether. To apply this standard literally to the global AI industry, or to the lay technologists and policymakers that this paper addresses, would be neither realistic nor arguably what most lay Jain ethical teaching itself prescribes. The tradition has always drawn a distinction between mahavratas (the great vows of monastics) and anuvratas (the more moderate vows appropriate to householders). The mapping proposed in Part III is pitched deliberately at the householder, or anuvrata, level of ambition: meaningful restraint and harm reduction within an ongoing engagement with technology, not renunciation of it.

6.3 The risk of extractive or superficial appropriation

There is a pattern of cherry-picking isolated concepts (mindfulness from Buddhism is a frequently cited example) from Western and secular appropriations of Asian religious traditions, and discarding the doctrinal, communal, and soteriological context that lends those concepts their original meaning and authority. This paper is open to the same objection unless read for what it professes to be: a secular policy synthesis drawing structural inspiration from Jain ethics. This is not a theological document, not an endorsement by any Jain religious body and not a claim that Jain practitioners are obliged to accept this application of their tradition's reasoning.

6.4 AI as a form of accumulation and control — a harder problem

A more difficult tension, not fully explored in the earlier sections of this paper, is that large-scale AI systems are not only energy-hungry, they are also, in their commercial form, vehicles for data collection, market concentration, and, in some cases, surveillance and domination, outcomes that sit uneasily with aparigraha’s suspicion of accumulation for its own sake and with ahimsa’s concern for the dignity and autonomy of other beings. A white paper that uses Jain vocabulary to make AI development more energy-efficient, but does not examine the concentration of power that efficient, large-scale AI systems can enable, would be following the letter of the tradition’s restraint principle while missing a great deal of its spirit. A more complete Jain informed critique of AI would require extending aparigraha to data and market power as well as compute and electricity—a direction for future work rather than a problem solved by this paper.

6.5 Measurement uncertainty

Finally, a purely empirical caveat: as multiple sources cited in Part I note, public data on AI-specific energy consumption remains a patchwork of vendor disclosures, industry estimates and modelling assumptions rather than comprehensive, independently audited measurement. The figures in this paper, from the IEA, Lawrence Berkeley National Laboratory, and Brookings, among others, are the most credible public estimates available as of mid-2026, but should be taken as order-of-magnitude indicators subject to revision, not precise constants.

7. Recommendations

7.1 For AI developers and technologists

- Think of model scale as a choice to be justified, not the default. Scale up (aparigraha) by benchmarking the smallest model that satisfies a real requirement.
- Report energy & water use for each major training run and, where possible, per unit of inference, using a consistent, comparable methodology (ahimsa, transparency).
- Construct mechanisms that intentionally bring to the surface perspectives outside the immediate development team, prior to major releases (anekantavada) – advisory input, external red-teaming, dissenting-view logging.
- Use lifecycle thinking in hardware and data: plan for reuse and recycling of equipment and deletion or retirement of data and models no longer having a justified purpose (aparigraha, ahimsa).

7.2 For policymakers and standard-setters

- Standardised, independently verifiable disclosure of data-centre electricity and water consumption specific to AI workloads is needed to close the gap between aggregate industry figures and accountability for individual systems.
  Scale up (aparigraha) by benchmarking the smallest model that satisfies a real requirement.
- Report energy & water use for each major training run and, where possible, per unit of inference, using a consistent, comparable methodology (ahimsa, transparency).
- Construct mechanisms that intentionally bring to the surface perspectives outside the immediate development team, prior to major releases (anekantavada) – advisory input, external red-teaming, dissenting-view logging.
- Utilise multi-criteria sustainability assessment frameworks, as advocated by UNESCO's 2021 Recommendation, rather than single-metric (carbon-only) assessments, to avoid optimising away one harm while worsening another.
- As a matter of industrial and research-access policy, not only environmental policy, support research into frugal and Green AI, recognising the democratising argument raised by frugal-AI researchers in Part IV.

7.3 For faith communities and interfaith dialogue

- Jain groups that are already working on environmental issues, such as JAINA and its Ahimsak Eco-Vegan Committee, are well positioned to explicitly expand their existing climate declaration models to cover digital and computational consumption, which current declarations seem to address only indirectly.
- Scale up (aparigraha) by benchmarking the smallest model that satisfies a real requirement.
- Report energy & water use for each major training run and, where possible, per unit of inference, using a consistent, comparable methodology (ahimsa, transparency).

- Construct mechanisms that intentionally bring to the surface perspectives outside the immediate development team, prior to major releases (anekantavada) – advisory input, external red-teaming, dissenting-view logging.
- Interfaith environmental networks, such as those coordinated through the Yale Forum on Religion and Ecology and the Alliance of Religions and Conservation, offer an existing infrastructure for Jain perspectives on technology to be placed in dialogue with other traditions' environmental teachings, broadening the plural, anekantavada-consistent character of the conversation this paper has tried to model.

7.4 For further research

- Empirical work quantifying the actual adoption and impact of frugal-AI techniques at industrial scale, beyond the research-paper stage.
- Scale up (aparigraha) by benchmarking the smallest model that satisfies a real requirement.
- Report energy & water use for each major training run and, where possible, per unit of inference, using a consistent, comparable methodology (ahimsa, transparency).
- Construct mechanisms that intentionally bring to the surface perspectives outside the immediate development team, prior to major releases (anekantavada) – advisory input, external red-teaming, dissenting-view logging.
- A broader Jain-inspired critique of the data-accumulation and market-concentration dynamics of AI, beyond aparigraha to include compute and electricity as noted in Section 6.4, and to extend the gems-and-jewellery case study in Part V to market power in the gem trade specifically.
- Comparative work that places this Jain-informed framework in the context of similar frameworks from other traditions — including Buddhist, Indigenous, Islamic, and Christian environmental ethics — consistent with the pluralistic spirit of the anekantavada itself.

8. Conclusion

Artificial intelligence and Jainism have radically different histories, timescales and epistemological starting points – the former a series of statistical and engineering techniques not much more than a human generation old in its current form, the latter an ethical and metaphysical tradition honed over two and a half millennia. Scale up (aparigraha) by benchmarking the smallest model that satisfies a real requirement.

● Report energy & water use for each major training run and, where possible, per unit of inference, using a consistent, comparable methodology (ahimsa, transparency).

● Construct mechanisms that intentionally bring to the surface perspectives outside the immediate development team, prior to major releases (anekantavada) – advisory input, external red-teaming, dissenting-view logging.

Yet both, in their own idiom, are posing a version of the same question: what should a powerful, resource-consuming agent do in a finite, interdependent world, when its own certainty about the consequences of its actions is bound to be partial?

In this paper, I have argued that the three central principles of Jainism—ahimsa, aparigraha, and anekantavada—constitute a coherent, and on inspection, technically translatable vocabulary for that question, one that already resonates with emergent practice in Green AI, frugal machine learning, and pluralistic AI governance, even where those fields have developed independently of any religious framing. It has also tried to be honest about where that translation strains: about the difference between monastic renunciation and industrial policy, about the danger of taking a tradition's vocabulary and discarding its context, and about the more difficult, unresolved question of what aparigraha would require of an industry that is built largely on the accumulation of data and market power, not just energy.

This gives you no finished formula. Rather, it offers a discipline of the sort anekantavada itself proposes: a means of balancing the scale of AI's physical footprint, the specificity of Jain ethical reasoning, and the real difficulty of constructing sustainable technology, all at once, without reducing any of them to the others. That, perhaps, is the most faithful thing a technology white paper can do with a philosophy built around the many-sidedness of truth.

References

1. All sources below were accessed in June–July 2026 and reflect the most current publicly available data at time of writing.
2. Brookings Institution. (2026). Global energy demands within the AI regulatory landscape. https://www.brookings.edu/articles/global-energy-demands-within-the-ai-regulatory-landscape/
3. Federation of Jain Associations in North America (JAINA), Ahimsak Eco-Vegan Committee. (2019). Jain Declaration on the Climate Crisis. https://veganjains.com/jain-declaration-on-the-climate-crisis-oct-2019/
4. Gateway House. (2024). A Century of Bombay-Antwerp Diamond Trade. https://www.gatewayhouse.in/a-century-of-bombay-antwerp-diamond-trade/
5. Grown Diamond Corporation. How Is Technology Advancing the Lab-Grown Diamond Industry? (GIA–IBM AI grading collaboration, 2020). https://www.growndiamondcorp.com/blog/how-is-technology-advancing-the-lab-grown-diamond-industry/

6. International Gem Society. (2021). Is Growing Diamonds a Sustainable Alternative to Diamond Mining? https://www.gemsociety.org/article/sustainable-alternative-to-diamond-mining/
7. Lum, K. (2014). The Rise and Rise of Belgium's Indian Diamond Dynasties. The Conversation. https://theconversation.com/the-rise-and-rise-of-belgiums-indian-diamond-dynasties-32332
8. Shopidc.com. (2024). The Environmental Impact of Naturally Mined Diamonds vs. Lab-Grown Alternatives. https://shopidc.com/blogs/news/the-environmental-impact-of-naturally-mined-diamonds-vs-lab-grown-alternatives
9. Flügel, P. (2005). Jainism and Ecology. Reviewed/summarised via ResearchGate. https://www.researchgate.net/publication/269791797_Jainism_and_Ecology
10. Institute of Jainology / Interfaith Center for Sustainable Development. Jain Ecology. https://interfaithsustain.com/jain-ecology/
11. International Energy Agency. (2025). Energy and AI. IEA, Paris. https://www.iea.org/reports/energy-and-ai
12. International Energy Agency. (2026). Key Questions on Energy and AI. IEA, Paris. https://www.iea.org/reports/key-questions-on-energy-and-ai
13. International Energy Agency. (2026). Electricity 2026. IEA, Paris. https://www.iea.org/reports/electricity-2026
14. Lala Bouali, Y., Ben Ahmed, O., Bradai, A., & Mazouzi, S. (2025). Frugal AI: The New Era of Sustainable Deep Learning. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5334656
15. Ranpara, R. (2025). Energy-Efficient Green AI Architectures for Circular Economies Through Multi-Layered Sustainable Resource Optimization Framework. arXiv:2506.12262. https://arxiv.org/pdf/2506.12262
16. A review of green artificial intelligence: Towards a more sustainable future. (2024). ScienceDirect / Neurocomputing. https://www.sciencedirect.com/science/article/pii/S0925231224008671
17. Satyagraha Foundation. The Jain Declaration on Nature (commentary and full text). https://www.satyagrahafoundation.org/the-jain-declaration-on-nature/
18. Singhvi, L. M. (1990). The Jain Declaration on Nature. Presented to the World Wide Fund for Nature. Yale Forum on Religion and Ecology archive. https://fore.yale.edu/sites/default/files/Jain-Declaration-on-Nature.pdf
19. The Anuvrat Movement: A Case Study of Jain-inspired Ethical and Eco-conscious Living. (2019). Religions, 10(11), 636. MDPI. https://www.mdpi.com/2077-1444/10/11/636
20. The Jain Principle of Ahimsa (Nonviolence) and Ecology. Totetu.org academic archive. https://www.totetu.org/assets/media/paper/j023_166.pdf
21. UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. Adopted by 193 Member States, November 2021. https://www.unesco.org/en/artificial-intelligence/recommendation-ethics
22. Violos, J., et al. (2025). Frugal Machine Learning for Energy-efficient, and Resource-aware Artificial Intelligence. arXiv:2506.01869. https://arxiv.org/abs/2506.01869
23. Yale Forum on Religion and Ecology. Overview Essay: Jainism. https://fore.yale.edu/World-Religions/Jainism/Overview-Essay