

Reassessing AI's Carbon Footprint: Toward a Sustainable Ethical Framework for Responsible Innovation.

Agordi Paul Oluwatosin*

ABSTRACT

With the way artificial intelligence (AI) is rapidly expanding today, when looked at from the angle of the structural engineering of such development, we come to realize that such innovation is contributing a large quota to giving environmental sustainability the stereotype of “a particular concern”. It should be noted that while AI systems offer significant societal and economic benefits, their underlying infrastructures, especially the engineering of data centers, which make AI work, contribute greatly to their growing carbon footprints. Using both qualitative and quantitative research methods, this paper critically examines the environmental implications of AI technologies, especially in the Global South. It argues that traditional ethical theories, such as utilitarianism and deontology, are inadequate for guiding the progress of AI innovations due to their rapidly evolving nature. It proposes a Sustainable AI Responsibility Framework (SARF) that justifies meeting the demands of AI innovations. It avers that the uniqueness of this framework is evident in its scope to incorporate and harness character formation, shared responsibility, regulatory clarity, and environmental justice. The paper concludes by recommending strategic policy structures for reducing AI's carbon footprint, which ensures that technological progress is not achieved at the detriment of environmental sustainability.

Keywords: Artificial Intelligence (AI), Carbon Footprint, Data Centers, Environmental Sustainability.

INTRODUCTION

It was Ramos Brandao who noted that AI has moved from a niche technology to a force reshaping how economies function, how people interact, and what ethical questions societies now have to answer.¹ This shift shows up concretely in automation, decision-making systems, and data analysis, areas where AI has genuinely changed how industries operate. This development does not come on a platter. As AI models grow more complex and “data-hungry”, the energy needed to train, run, and maintain them grows right along with it, and such growth leads to carbon emissions and resource strain. Haoyang Liu and Jiangtao Zhai make a similar point by contending that even AI computing powered by renewable energy still produces meaningful emissions, simply because of how intensive high-performance workloads are.²

It is premised on this growing concern that this paper argues that AI's carbon footprint has stopped being a purely technical problem and has become a global moral one. Ethical responsibility for AI development has to include environmental stewardship, especially in the Global South, where AI infrastructure is expanding fast even as many of these regions already

¹ Pedro Ramos Brandão, "The Impact of Artificial Intelligence on Modern Society," *AI* 6, no. 8 (2025): 190, <https://doi.org/10.3390/ai6080190>.

² Haoyang Liu and Jiangtao Zhai, "Carbon Emission Modeling for High-Performance Computing-Based AI in New Power Systems with Large-Scale Renewable Energy Integration," *Processes* 13, no. 2 (2025): 595, <https://doi.org/10.3390/pr13020595>.

face environmental and energy strain. Traditional ethical theories, including utilitarianism and deontology, were not built for this. They assume a single human actor making a choice, not the distributed, multi-agent reality of how AI actually gets built and deployed. This gap is the reason why this paper adopts the Sustainable AI Responsibility Framework (SARF), built to trace accountability across the whole chain behind an AI system's footprint: developers, deployers, organizations, regulators, users, instead of forcing a single-actor theory onto a problem no single actor created. Therefore, folding sustainability into the core principles of AI development this way, instead of treating it as an afterthought, is what keeps technological growth from coming at the expense of the environment.

UNDERSTANDING AI'S CARBON FOOTPRINT

AI's carbon footprint refers to the greenhouse gases, mostly carbon dioxide (CO₂), released while AI systems are built, trained, deployed, and used. John McCarthy defined AI as "the science and engineering of making intelligent machines, brilliant computer programs."³ Thus, running such a science takes real mechanical and electrical infrastructure. As models scale up to chase Artificial General Intelligence, they burn through more energy and resources, and the environment absorbs the cost.⁴

Before delving into the stages through which AI's carbon footprint is generated, it is appropriate that we probe the meaning, nature, and effects of carbon emissions to understand this section better.

A Survey into the Meaning, Nature, and Effect of Carbon Emissions

Carbon emissions primarily refer to the release of carbon dioxide (CO₂). This is a greenhouse gas that is composed of one carbon atom bonded to two oxygen atoms. Most of these emissions come from everyday human activities, such as burning fossil fuels including coal, oil, and natural gas for energy and transportation, industrial processes like cement and steel production, and deforestation.⁵ On the one hand, these activities drive economic growth, but on the other hand, they pump enormous volumes of CO₂ into the atmosphere, which contributes to the greenhouse effect and, in some cases, ozone depletion.

At normal atmospheric concentrations, around 0.04 percent or 400 parts per million, CO₂ is not directly toxic. The danger comes from accumulation. This is when it builds up by trapping heat that drives global warming. It also disrupts weather patterns and ecosystems in ways that eventually affect human health. At much higher concentrations, this is above roughly

³ John, McCarthy, *Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. 1956. Accessed May 6, 2025. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>.

⁴ The AGI is type of advanced AI that can generalize by simulating intellectual human abilities, adapt to new situations, and transfer knowledge between domains. It is argued hypothetically to transcend human intelligence, wherein humans will depend on it, and be able to operate independently without explicit instruction on every step. Cf. Agordi Paul Oluwatosin, "Artificial Intelligence in the African Context: A Philosophical and Ethical Inquiry," *Nnadiabube Journal of Philosophy*, Vol. 8 No. 4 (2025):93, <https://acjoi.org/index.php/NJP/article/view/7767/7488>; Legg, Shane, and Marcus Hutter, "Universal Intelligence: A Definition of Machine Intelligence," *Minds and Machines* 17, no. 4 (2007): 391–444.

⁵ U.S. Environmental Protection Agency, "Sources of Greenhouse Gas Emissions," accessed July 15, 2026, <https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions>.

30,000 ppm, CO₂ turns genuinely hazardous because it displaces oxygen, which is a separate problem from its role in air pollution and respiratory illness.⁶

A case reported out of Korea in 2020 makes the risk concrete. A 66-year-old man went into cardiac arrest after a CO₂ tank exploded in an enclosed space. The gas is heavier than air, so once released, it pooled and pushed the oxygen out of the room. He lost consciousness almost immediately and stopped breathing. Responders got his heart going again, but the damage was already done. He ended up in intensive care with hypoxic brain injury, the kind of outcome that happens when the brain goes too long without enough oxygen.⁷

Having stated some of the effects of carbon on humans and the environment, we shall now examine the stages (nodes) through which AI's carbon footprint is generated.

Stages through which AI Carbon Footprint Is Generated

1. The Manufacturing Stage

This stage involves the production of hardware for AI systems. This includes GPUs, CPUs, servers, storage devices, and networking equipment. In this phase, a carbon footprint is generated during the extraction of raw materials, the fabrication of chips, and the transportation of components from one location to another. All these activities consume large amounts of fossil fuels and release greenhouse gases.⁸

2. Model Training

In training a modern deep-learning model, it is made to run over massive datasets in several passes, across clusters of GPUs or TPUs. This sometimes takes hours at a stretch or weeks on end. This back-and-forth process, also known as forward and backward propagation, consumes electricity at a scale most people do not picture when they think about AI.⁹ Much of that electricity still comes from data centers running on fossil fuels, and the result is tens to hundreds of metric tons of CO₂-equivalent emissions, sometimes more. For the largest models, training alone can outweigh every other stage of the model's lifecycle combined. Giovanni Cocca Guardia et al. push back on the idea that renewable energy solves this. For them, even with partial renewable sourcing, cooling systems, redundancy protocols, and the hardware manufacturing behind it all still drive the footprint up considerably.¹⁰

3. Inference (Deployment)

Inference is what happens when a trained model actually runs, generating outputs from new inputs, whether it is a chatbot answering a prompt, a classifier reading a scan, or an embedded

⁶ Ugo Bardi, Phil Bierwirth, Kuo-Wei Huang, and John McIntyre, "Carbon Dioxide as a Pollutant: The Risks to Human Health and the Stability of the Biosphere," *Environmental Science: Advances* 4 (2025): 1364–1372, <https://pubs.rsc.org/en/content/articlehtml/2025/va/d5va00017c>.

⁷ Sang-O Park, Ju-Won Kim, and Sun-Young Kim, "Case of Cardiac Arrest Due to Carbon Dioxide Poisoning Following an Explosion of a Carbon Dioxide Tank," *Clinical and Experimental Emergency Medicine* 7, no. 4 (2020): 336–339, <https://doi.org/10.15441/ceem.19.015>.

⁸ Green AI Institute, "White Paper on Global Artificial Intelligence Environmental Impact," accessed July 15, 2026, <https://www.greenai.institute/whitepaper/white-paper-on-global-artificial-intelligence-environmental-impact>.

⁹ Xing Chen et al., "Self-Contrastive Forward-Forward Algorithm," *Computational Science* 25, no. 6 (2025): 112–130, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12217723/>.

¹⁰ Giovanni Cocca Guardia et al., "Assessing the Energy Impact and Carbon Footprint of AI Model Training: A Case Study Using GPU Servers," *IEEE Conference Paper*, accessed November 4, 2025, <https://clagtee.fi.mdp.edu.ar/full-papers-search-engine/papers/ID071.pdf>.

model running inside a self-driving car. Training happens once. Inference happens every time anyone uses the model, and that gap adds up; a model serving a billion queries a day can burn through roughly 300 MWh daily, enough to exceed its training energy cost in under five days, especially where data centers still run on coal, gas, or oil.¹¹ Any single query looks negligible, but O'Donnell and Crownhart argue that the real cost only shows up once you count what the industry is not tracking and where the usage is headed.¹²

4. **Data-Centre Infrastructure:**

Data centers need serious infrastructure beyond just servers, such as networking, storage, lighting, and especially cooling, since cooling efficiency and Power Usage Effectiveness (PUE) largely determine how much CO₂ comes out per unit of electricity used. To this end, the International Energy Agency IEA points to AI as a major driver of rising data center electricity demand in the coming years, and as AI workloads grow, that demand for power and cooling grows with it, pushing emissions up unless energy sources and cooling methods get cleaner.¹³

At this point, an ethical question presents itself: given everything AI makes possible, should we shut these systems down over their carbon footprint? It is due to this concern that we shall look into some traditional ethical frameworks that would help to keep AI innovations in check.

TRADITIONAL ETHICAL FRAMEWORKS

In the scope of Utilitarianism, it weighs AI's benefits such as medicine, climate modeling, education, and productivity against its environmental costs.¹⁴ The problem associated with this framework is that it leaves out the minority communities harmed by pollution and data centers, which rarely factor into the calculation, and the whole method assumes we can predict consequences AI often makes unpredictable.

Deontology treats environmental stewardship as a duty owed regardless of outcome, and as a result of this submission, reckless AI expansion breaks a moral rule even if it produces good elsewhere.¹⁵ The theory holds up better than utilitarianism here, but duties are only as strong as someone's willingness to enforce them, and no single body is positioned to do that across the whole AI industry. Thilo Hagendorff's critique of AI ethics guidelines backs this up by arguing that deontological approaches to AI ethics often amount to a "tick-box exercise," principles stated on paper with nothing forcing anyone to follow them.¹⁶

¹¹ James O'Donnell and Casey Crownhart, "We Did the Math on AI's Energy Footprint. Here's the Story You Haven't Heard," *MIT Technology Review*, May 20, 2025, <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>.

¹² O'Donnell and Crownhart, "We Did the Math on AI's Energy Footprint."

¹³ International Energy Agency, *Electricity 2024: Analysis and Forecast to 2026*, accessed November 4, 2025, <https://www.iea.org/reports/electricity-2024>.

¹⁴ Henry R. West, "Utilitarianism," *Encyclopædia Britannica*, last modified 2025, <https://www.britannica.com/topic/utilitarianism-philosophy>.

¹⁵ The Editors of Encyclopædia Britannica, "Deontological Ethics," *Encyclopædia Britannica*, October 22, 2025, <https://www.britannica.com/topic/deontological-ethics>.

¹⁶ Thilo Hagendorff, "The Ethics of AI Ethics: An Evaluation of Guidelines," *Minds and Machines* 30 (2020): 99–120.

Virtue ethics wants technologists and policymakers to act with prudence and restraint instead of chasing growth.¹⁷ That is a fine standard for a single person. However, it is a weak one for an institution, since personal virtue does not scale to corporate or government accountability, and most of the time there is no rulebook behind it.

Furthermore, environmental justice preoccupies itself with naming who actually pays.¹⁸ Wealthy nations host the data centers; poorer ones absorb the mining and e-waste, cobalt from the DRC being the clearest case. Thus, this framework is sharp on fairness, but blind in its design. It tells us an injustice happened after the fact instead of shaping how AI gets built before it does. Due to this challenge, Schlosberg argues for the expansion of environmental justice to include recognition and participation, stressing fairness in global environmental governance.¹⁹

Why Traditional Ethical Frameworks Fall Short of AI Innovations

In the above frameworks discussed, it becomes clear that despite the efforts of the frameworks to attend to AI sustainable development, they remain inadequate. This is because these traditional ethical theories assume a moral agent with intentions, foresight, and the capacity to be held accountable. They focus on someone who can be praised, blamed, or punished for a choice they made.

AI systems do not fit in this scope. Utilitarianism needs someone who can weigh consequences in advance. Deontology needs someone capable of willing a duty. Virtue ethics needs a character built up over a lifetime. An AI model has none of that. It has no intentions of its own, and once deployed, it can act in ways its own developers never predicted. Deborah Johnson makes this exact case. She argues that computer systems fail the traditional test of moral agency because, no matter how sophisticated they get, they lack what she calls “intendings to act”; this is the free, deliberate willing that only a conscious agent has. Hence, whatever a system does, it does because it was built and deployed that way, not because it chose to.²⁰

Some philosophers call this the *responsibility gap*. Andreas Matthias was the first to name such a gap; for him, as learning systems grow more autonomous, the gap widens between what a system decides and what any human can foresee or control, and that gap is exactly where frameworks built for human choice stop working.²¹ Robert Sparrow pushed the same idea further with autonomous weapons. He contends that if a system was not programmed to cause harm, and no human was steering it at the time, then the harm was not foreseeable, and moral responsibility has nowhere to land.²²

¹⁷ Aristotle, *Nicomachean Ethics*, trans. Terence Irwin (Indianapolis: Hackett Publishing Company, 1999), Book II, Chapters 6–9.

¹⁸ Robert D. Bullard, *Dumping in Dixie: Race, Class, and Environmental Quality* (Boulder: Westview Press, 1990), 15–20.

¹⁹ David Schlosberg, *Defining Environmental Justice: Theories, Movements, and Nature* (Oxford: Oxford University Press, 2007), 45–50.

²⁰ Deborah G. Johnson, "Computer Systems: Moral Entities but Not Moral Agents," *Ethics and Information Technology* 8, no. 4 (2006): 195–204.

²¹ Andreas Matthias, "The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata," *Ethics and Information Technology* 6, no. 3 (2004): 175–83.

²² Robert Sparrow, "Killer Robots," *Journal of Applied Philosophy* 24, no. 1 (2007): 62–77.

Because of this gap, some philosophers and organizations have proposed recommendations to address the concerns raised by AI advancements. Virginia Dignum argues that responsible AI comes down to governance built on accountability, transparency, and human oversight, and on that basis, existing governance mechanisms already handle much of AI's ethical terrain.²³ However, her framework becomes less robust when it comes to environmental factors. It establishes institutional accountability without ever distributing responsibility for carbon emissions across an AI system's full lifecycle, from hardware manufacturing through deployment and end use.

The EU High-Level Expert Group's *Ethics Guidelines for Trustworthy AI* (2019) also seek to address many of the ethical challenges arising from AI innovation. As a result, it identifies "societal and environmental well-being" as one of the seven key requirements for trustworthy AI, alongside accountability, transparency, and human oversight, and frames environmental sustainability as a consideration throughout an AI system's lifecycle rather than only at the deployment stage.²⁴ This effort by the EU High-Level Expert Group falls short in the area of enforcement and specificity. The guidelines are voluntary; they ask developers to "consider sustainability" without assigning that consideration to any particular actor in the chain, and the assessment list built around them treats environmental well-being as one checkbox among seven.

Premised on this concern, this paper advocates for the adoption of a Sustainable AI Responsibility Framework (SARF) that will handle the complex, global, and fast-evolving nature of AI technologies. SARF picks up that governance logic and pushes it further, making environmental sustainability a central obligation instead of a policy afterthought.

SUSTAINABLE AI RESPONSIBILITY FRAMEWORK (SARF)

The Sustainable AI Responsibility Framework (SARF) is not alien in its scope, nor is it a new idea. It is a framework that builds on existing ethical principles and adapts them to the specific problems AI poses, and its logic already operates in practice. It rests on four major principles already established in different philosophical literature. These include the following:

i. UNESCO's Recommendation on the Ethics of Artificial Intelligence

UNESCO's Recommendation, adopted by all 193 member states in 2021, treats environmental sustainability as one of four core values underpinning ethical AI, alongside human rights, human dignity, and diverse, flourishing societies.²⁵ It dedicates one of its policy action areas specifically to "Environment and Ecosystems," requiring member states to assess and reduce the carbon and material footprint of AI systems across their whole lifecycle, not just at the point of deployment.²⁶

²³ Virginia Dignum, *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way* (Cham: Springer Nature Switzerland, 2019), 17-29.

²⁴

²⁵ UNESCO, *Recommendation on the Ethics of Artificial Intelligence* (Paris: UNESCO, 2021), sec. III, "Values."

²⁶ UNESCO, *Recommendation on the Ethics of Artificial Intelligence*, sec. V, "Policy Action Areas: Environment and Ecosystems."

ii. **Vallor's Case for Technomoral Virtues**

Shannon Vallor argues that older ethical frameworks such as deontology, utilitarianism, etc., assume a static human condition and simply cannot keep pace with technology that changes faster than moral habits can form.²⁷ Vallor's advocacy for "technomoral virtues" is not a new rulebook; she argues for the practice of virtues like self-control, humility, and civility, meant to be practiced deliberately so people can adapt their character to tools that did not exist when the original virtue traditions were written.

iii. **Bostrom's Intergenerational Standard**

Nick Bostrom argues that AI decisions have to be judged by their long-term, intergenerational consequences, not just what they deliver this quarter.²⁸ For Bostrom, a technology this powerful can lock in outcomes for generations who never get a vote, so any ethical accounting that stops at present-day costs and benefits is measuring the wrong thing.

iv. **Floridi's Infosphere**

Luciano Floridi's idea of the infosphere treats digital systems and the natural environment as one connected space we are responsible for, not two separate ones.²⁹ In his framework, data, hardware, and ecosystems should not be considered separate categories to weigh against each other; they are part of a single environment, which means harming one through carbon-heavy computing is really harming the whole thing.

v. **Distributed Responsibility in AI Governance**

Recent work on distributed responsibility in AI governance spreads accountability across developers, regulators, and users instead of pinning it on whoever's easiest to blame.³⁰ This model treats an AI system's carbon footprint as the product of a whole chain: developers who train it, companies who deploy it, regulators who permit it, users who query it, so no single actor gets to claim it wasn't their responsibility. This aligns with Strasser's Theory of Distributed Responsibility, which argues that human-machine interactions rarely leave a single actor to blame, since the humans involved are usually both the creators of the system and the ones who set it in motion.³¹

Therefore, SARF is a synthesis of virtue ethics, distributed responsibility, principlism, and ethics of care, stitched together because each one plugs a hole the others leave open. Virtue ethics carries over the character piece traditional theories relied on. Distributed responsibility fixes the accountability problem that sinks deontology once we are dealing with a whole industry instead of one person. Principlism gives regulators something concrete to enforce. Ethics of care keeps vulnerable communities and ecosystems in the picture; the exact group

²⁷ Shannon Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (New York: Oxford University Press, 2016), 6–7.

²⁸ Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014), 209–210.

²⁹ Luciano Floridi, *The Ethics of Information* (Oxford: Oxford University Press, 2013), 9–10.

³⁰ Markus D. Dubber, Frank Pasquale, and Sunit Das, eds., *The Oxford Handbook of AI Ethics* (Oxford: Oxford University Press, 2022), 312–315.

³¹ Anna Strasser, "Distributed Responsibility in Human–Machine Interactions," *AI and Ethics* 2 (2022): 523–32.

utilitarianism tends to lose in the math. Together, these perspectives provide a coherent ethical architecture for sustainable AI governance.

Furthermore, this combination is what makes SARF a normative framework in the scope of sustainability, one that actually does something about AI's carbon footprint instead of just describing it. It hands developers a duty, asking them to act with restraint even where regulation has not caught up yet. It also spreads the cost of getting it wrong across the whole chain of actors instead of one scapegoat, and keeps the communities absorbing the damage in view alongside the ones enjoying the benefits.

Lastly, SARF echoes the mind of *Antiqua et Nova*, the Vatican's 2025 note on artificial intelligence and human intelligence, which fits its foundations closely. It insists AI must remain a tool in service of human dignity rather than a replacement for human judgment, and it flags the concentration of AI's benefits in the hands of "a few powerful companies" as a real ethical concern, the same distributional worry environmental justice raises about who hosts the data centers and who absorbs the mining and e-waste.³² The document also warns against treating AI as a stand-in for human relationships or moral discernment, which sits well next to Agordi's refusal to hand moral responsibility to the machine itself. Thus, accountability stays with the people who design, deploy, and use these systems, not with the system.³³

Operational Structure of the Sustainable AI Responsibility Framework (SARF)

SARF's implementation runs on four weighting factors spread across its five actor nodes³⁴: developers, deployers, organizations, regulators, and users, instead of treating every actor as equally liable no matter their actual position in the chain. The following are the weighting factors:

1. Foreseeability

This weighs how predictable a harm was at the point of decision, so a developer who chose an energy-hungry architecture knowing the tradeoff carries more responsibility than a user who just queried the deployed model without any idea of its footprint.

2. Controllability

This factor weighs who actually had the power to choose differently, putting more weight on a deployer signing an energy contract or an organization approving a buildout than on a regulator working with limited authority or outdated tools.

3. Benefit

³² Dicastery for the Doctrine of the Faith and Dicastery for Culture and Education, *Antiqua et Nova: Note on the Relationship Between Artificial Intelligence and Human Intelligence* (Vatican City, January 28, 2025), para. 52.

³³ Paul Agordi, "Reassessing Moral Agency in Artificial Intelligence: A Strasserian Reconstruction of Responsibility Attribution in Machine Intelligence" (Master's Thesis, Dominican University, Samonda, Ibadan, 2026), 68-70.

³⁴ Agordi, "Reassessing Moral Agency," Chap. 4.

This weighs who profits from the activity causing the harm, so an organization capturing revenue from AI's growth carries a heavier share than a data provider who contributed information without capturing any of that upside.

4. Vulnerability

It weighs who bears the cost, tilting responsibility toward protecting communities near data centers or in resource-strained regions instead of letting their exposure pass as an acceptable externality.

Together, these four factors let SARF assign responsibility in gradients rather than blanket terms, so accountability tracks where the real decision-making power, and the real damage, actually sit.

A Look into the Operability of SARF in a Real-World Case Study

Let us consider Microsoft's FY2025 sustainability report that was released in mid-2026, which disclosed a 25 percent rise in total carbon emissions, 20 million metric tons of carbon dioxide equivalent, up from 16 million the year before. These emissions were driven mainly by data center expansion for AI workloads and a decision to stop buying certain renewable energy certificates.³⁵

When this is evaluated by SARF, the five actor nodes through which the carbon footprint is created are examined critically. Put differently, SARF reduces AI's carbon footprint by assigning responsibility to the exact point in the chain where a decision is actually made, rather than leaving the entire burden with whoever happens to be visible when emissions are reported.

1. Developers, the teams training and scaling the models that demand this infrastructure, carry responsibility weighted by foreseeability. This is because the energy cost of scaling a model is not a surprise; it is a known trade-off that was made at design time.
2. Deployers, Microsoft's cloud division and its Azure customers, carry responsibility weighted by controllability, since they are the ones deciding how aggressively to expand capacity and which energy contracts to sign, including the choice to walk away from renewable certificates in favor of new fossil-fueled capacity, as Microsoft did with its gas-fired plant deal in West Texas.³⁶
3. The organization itself carries responsibility weighted by benefit, since AI revenue growth is what justified the buildout in the first place.
4. Data providers and users carry a lighter but real share, weighted by vulnerability, since the communities near new data centers absorb strained grids and water stress without sharing in the profit.

³⁵ "Microsoft's Emissions Surged 25% in 2025 During Data Center Boom," *Bloomberg*, July 9, 2026.

³⁶ "Microsoft's Data Center Expansion Drove 25% Emissions Spike in 2025," *ESG Dive*, July 2026, <https://www.esgdive.com/news/microsofts-data-center-expansion-drove-25-emissions-spike-2025-environmental-sustainability-report/825082/>.

It is clear at this point that a single traditional framework does not hold up well against this case, as is evident in the case of traditional ethical frameworks. However, SARF does not make the tension go away. It spreads it by tying the emissions to everyone's actual role instead of leaving them as a line item that Microsoft alone gets to explain away in a report. Hence, by making the cost visible and owned at every stage instead of diffused into nobody's job, SARF helps to push firms toward the standardized disclosure, cleaner sourcing, and efficiency gains the policy section calls for.

Furthermore, SARF works best when backed by strategic sustainability policies, since that pairing does more than name who is at fault; it also points toward the fix. The policy recommendation is what the next section takes up.

POLICY RECOMMENDATIONS FOR SUSTAINABLE AI GOVERNANCE

1. Mandatory Carbon Accountability

AI firms and cloud providers should publish standardized lifecycle emissions, both location- and market-based, with full Scope 3 disclosure and independent verification.³⁷ This is to enable companies to avoid picking any number that makes them look good. Alexandra Sasha Luccioni, an AI and Climate Lead at Hugging Face, voiced this same concern. Her research has repeatedly found that tech companies are not measuring or disclosing this data at all. Thus, she called for a standardized energy rating for AI models, the same idea behind appliance energy labels, so buyers can compare models on real numbers instead of whatever a company chooses to release.³⁸

2. International Carbon Accounting

AI infrastructure is global, so national rules alone tend to push companies to relocate computing to whichever jurisdiction asks fewer questions. William Nordhaus's work on climate clubs describes the same dynamic at a larger scale. He argues that without coordinated international rules, countries and firms have a constant incentive to free-ride, letting others carry the cost of climate action while they shift their own emissions elsewhere.³⁹ A shared standard, through the UNFCCC or an OECD/ITU framework, would close that gap by harmonizing metrics like CO₂e per training run across borders, leaving no jurisdiction able to offer a discount on emissions.⁴⁰

3. Data Center Siting and Energy Rules

³⁷ Greenhouse Gas Protocol, *The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard*, rev. ed. (Washington, DC: World Resources Institute and World Business Council for Sustainable Development, 2015).

³⁸ Alexandra Sasha Luccioni, Boris Gamazaychikov, Sara Hooker, Régis Pierrard, Emma Strubell, Yacine Jernite, and Carole-Jean Wu, "Light Bulbs Have Energy Ratings, So Why Can't AI Chatbots?" *Nature* 632, no. 8026 (2024): 736–38.

³⁹ William Nordhaus, "Climate Clubs: Overcoming Free-Riding in International Climate Policy," *American Economic Review* 105, no. 4 (2015): 1339–70.

⁴⁰ International Organization for Standardization, *ISO 14064-1:2018: Greenhouse Gases-Part 1: Specification with Guidance at the Organization Level for Quantification and Reporting of Greenhouse Gas Emissions and Removals* (Geneva: International Organization for Standardization, 2018).

Most of AI's energy draw happens at the data center, so governments should require full environmental impact assessments before construction, regular reporting on power and water use, and incentives to reuse waste heat or generate renewable energy nearby.

4. Pricing in the Environmental Cost

Carbon pricing or surcharges on heavy computing push companies toward efficiency by making emissions cost something real, while public procurement that favors low-carbon AI and tax breaks for renewable adoption pull the same lever from the demand side, rewarding the firms already doing the right thing. Within SARF, this sharpens the organization node's accountability, wherein companies can no longer profit from AI's growth without paying for the emissions that growth produces. It also pushes developers and providers toward cleaner design and energy choices, since energy-intensive systems and fossil-fuel-powered operations now cost more than sustainable alternatives. In fact, the IMF has recommended exactly this same concern, though political resistance and measurement gaps have kept it mostly theoretical so far, which is itself the argument for pairing SARF's distributed accountability with policy that can actually enforce a price on the damage.⁴¹

5. Energy Transparency and Green AI Norms

This policy holds that developers should disclose energy use across training and deployment and treat efficiency as a real measure of progress, not an afterthought to performance, backed by funding and publication rules that actually reward it. Also, Schwartz, Dodge, Smith, and Etzioni coined the term "Green AI" to challenge a field that had come to equate progress with scale, bigger models, more parameters, more compute, while leaving the energy cost for someone else to worry about.⁴² They argue that papers should report computational cost alongside results, so a marginal accuracy gain costing ten times the energy does not pass as a clean win. And if journals and funding bodies enforced that, it would change what counts as good research, since a smaller, efficient model currently has no real shot at the recognition a record-breaking parameter count gets automatically.

CONCLUSION

A careful look at AI today shows that, alongside its real benefits to humanity, it carries serious environmental and ethical costs. Its dependence on energy-hungry computation and resource-intensive hardware ties it directly to rising carbon emissions, and left unaddressed, those emissions could undercut the very sustainability goals AI is supposed to help advance.

This paper has argued that no single traditional ethical framework (utilitarianism, deontology, virtue ethics, or environmental justice) can handle AI's environmental impact alone,

⁴¹ Shafik Hebous and Nate Vernon-Lin, "Carbon Emissions from AI and Crypto Are Surging, and Tax Policy Can Help," *IMF Blog*, August 15, 2024, <https://www.imf.org/en/Blogs/Articles/2024/08/15/carbon-emissions-from-ai-and-crypto-are-surging-and-tax-policy-can-help>.

⁴² Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni, "Green AI," *Communications of the ACM* 63, no. 12 (2020): 54–63.

given how vast, multilayered, and constantly evolving that impact is; And that is compounded by the fact that traditional ethical frameworks were designed for human agents in the first place, ones capable of intention, foresight, and accountability, not for systems that can act in ways their own developers never predicted.⁴³

However, this gap is what the Sustainable AI Responsibility Framework (SARF) seeks to close. SARF is a synthesis of virtue ethics, distributed responsibility, principlism, and ethics of care that traces accountability across the whole chain of actors, developers, deployers, organizations, regulators, and users, weighted by foreseeability, controllability, benefit, and vulnerability. SARF is not a theory floating free of practice. It rests on foundations already active in the real world, such as Vallor's technomoral virtues, Bostrom's intergenerational standard, Floridi's infosphere, the EU AI Act's tiered accountability structure, and UNESCO's global recommendation naming environmental sustainability as a core AI value. It finds an echo, too, in the Catholic Church's own concern, raised in *Antiqua et Nova*, that AI remain a tool serving human dignity rather than a technology whose benefits concentrate in a few hands while its costs fall on everyone else.

None of this works without policies behind it. Standardized emissions reporting, international coordination through mechanisms like the UNFCCC, stricter siting and energy rules for data centers, carbon pricing that makes pollution cost something real, and honest disclosure of energy use. Together, these policy recommendations turn SARF's distributed accountability from a philosophical claim into something enforceable. Governments, AI companies, and researchers each carry a piece of this, and none can opt out by pointing to someone else in the chain.

In summary, SARF shows that sustainable AI governance has to align responsibility to influence, not just proximity to harm. By pulling virtue ethics, distributed responsibility, principlism, and ethics of care into one normative framework, it offers a practical way to assign responsibility across the AI lifecycle while keeping environmental sustainability central to what responsible AI innovation actually means.

***Agordi Paul Oluwatosin**

Dominican University,
Samonda, Ibadan, Oyo State, Nigeria.
agordipaul@gmail.com

BIBLIOGRAPHY

Agordi, Paul. *Reassessing Moral Agency in Artificial Intelligence: A Strasserian Reconstruction of Responsibility Attribution in Machine Intelligence*. Master's thesis, Dominican University, Samonda, Ibadan, 2026.

⁴³ Andreas Matthias, "The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata," *Ethics and Information Technology* 6, no. 3 (2004): 175–83.

- Agordi, Paul Oluwatosin. "Artificial Intelligence in the African Context: A Philosophical and Ethical Inquiry." *Nnadiesube Journal of Philosophy* 8, no. 4 (2025): 93. <https://acjoi.org/index.php/NJP/article/view/7767/7488>.
- Aristotle. *Nicomachean Ethics*. Translated by Terence Irwin. Indianapolis: Hackett Publishing Company, 1999.
- Bardi, Ugo, Phil Bierwirth, Kuo-Wei Huang, and John McIntyre. "Carbon Dioxide as a Pollutant: The Risks to Human Health and the Stability of the Biosphere." *Environmental Science: Advances* 4 (2025): 1364-1372. <https://pubs.rsc.org/en/content/articlehtml/2025/va/d5va00017c>.
- Bloomberg. "Microsoft's Emissions Surged 25% in 2025 During Data Center Boom." July 9, 2026.
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.
- Brandão, Pedro Ramos. "The Impact of Artificial Intelligence on Modern Society." *AI* 6, no. 8 (2025): 190. <https://doi.org/10.3390/ai6080190>.
- Bullard, Robert D. *Dumping in Dixie: Race, Class, and Environmental Quality*. Boulder, CO: Westview Press, 1990.
- Chen, Xing, et al. "Self-Contrastive Forward-Forward Algorithm." *Computational Science* 25, no. 6 (2025): 112-130. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12217723/>.
- Cocca Guardia, Giovanni, et al. "Assessing the Energy Impact and Carbon Footprint of AI Model Training: A Case Study Using GPU Servers." IEEE conference paper. Accessed November 4, 2025. <https://clagtee.fi.mdp.edu.ar/full-papers-search-engine/papers/ID071.pdf>.
- Dauner, Maximilian, and Gudrun Socher. "Energy Costs of Communicating with AI." *Frontiers in Communication* 10 (2025). <https://www.frontiersin.org/journals/communication/articles/10.3389/fcomm.2025.1572947/full>.
- Dicastery for the Doctrine of the Faith and Dicastery for Culture and Education. *Antiqua et Nova: Note on the Relationship Between Artificial Intelligence and Human Intelligence*. Vatican City, January 28, 2025.
- Dubber, Markus D., Frank Pasquale, and Sunit Das, eds. *The Oxford Handbook of AI Ethics*. Oxford: Oxford University Press, 2022.
- ESG Dive. "Microsoft's Data Center Expansion Drove 25% Emissions Spike in 2025." July 2026. <https://www.esgdive.com/news/microsofts-data-center-expansion-drove-25-emissions-spike-2025-environmental-sustainability-report/825082/>.

- Floridi, Luciano. *The Ethics of Information*. Oxford: Oxford University Press, 2013.
- Green AI Institute. "White Paper on Global Artificial Intelligence Environmental Impact." Accessed July 15, 2026. <https://www.greenai.institute/whitepaper/white-paper-on-global-artificial-intelligence-environmental-impact>.
- Greenhouse Gas Protocol. *The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard*. Rev. ed. Washington, DC: World Resources Institute and World Business Council for Sustainable Development, 2015.
- Hagendorff, Thilo. "The Ethics of AI Ethics: An Evaluation of Guidelines." *Minds and Machines* 30 (2020): 99-120.
- Hebous, Shafik, and Nate Vernon-Lin. "Carbon Emissions from AI and Crypto Are Surging, and Tax Policy Can Help." *IMF Blog*, August 15, 2024. <https://www.imf.org/en/Blogs/Articles/2024/08/15/carbon-emissions-from-ai-and-crypto-are-surging-and-tax-policy-can-help>.
- International Energy Agency. *Electricity 2024: Analysis and Forecast to 2026*. Accessed November 4, 2025. <https://www.iea.org/reports/electricity-2024>.
- International Organization for Standardization. *ISO 14064-1:2018: Greenhouse Gases-Part 1: Specification with Guidance at the Organization Level for Quantification and Reporting of Greenhouse Gas Emissions and Removals*. Geneva: International Organization for Standardization, 2018.
- Jacobson, Tyler A., et al. "Direct Human Health Risks of Increased Atmospheric Carbon Dioxide." *Nature Sustainability* 2 (2019): 691-703. <https://www.nature.com/articles/s41893-019-0323-1.pdf>.
- Johnson, Deborah G. "Computer Systems: Moral Entities but Not Moral Agents." *Ethics and Information Technology* 8, no. 4 (2006): 195-204.
- Legg, Shane, and Marcus Hutter. "Universal Intelligence: A Definition of Machine Intelligence." *Minds and Machines* 17, no. 4 (2007): 391-444.
- Liu, Haoyang, and Jiangtao Zhai. "Carbon Emission Modeling for High-Performance Computing-Based AI in New Power Systems." *Processes* 13, no. 2 (2025): 595. <https://www.mdpi.com/2227-9717/13/2/595>.
- Luccioni, Alexandra Sasha, Boris Gamazaychikov, Sara Hooker, Régis Pierrard, Emma Strubell, Yacine Jernite, and Carole-Jean Wu. "Light Bulbs Have Energy Ratings, So Why Can't AI Chatbots?" *Nature* 632, no. 8026 (2024): 736-738.
- Matthias, Andreas. "The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata." *Ethics and Information Technology* 6, no. 3 (2004): 175-183.

- McCarthy, John. *Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. 1956. Accessed May 6, 2025. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>.
- National Institute for Occupational Safety and Health. "Carbon Dioxide-IDLH." Centers for Disease Control and Prevention. Accessed July 15, 2026. <https://www.cdc.gov/niosh/idlh/124389.html>.
- National Oceanic and Atmospheric Administration. "Understanding Ocean Acidification." NOAA Fisheries. Accessed July 15, 2026. <https://www.fisheries.noaa.gov/insight/understanding-ocean-acidification>.
- Nordhaus, William. "Climate Clubs: Overcoming Free Riding in International Climate Policy." *American Economic Review* 105, no. 4 (2015): 1339-1370.
- O'Donnell, James, and Casey Crownhart. "We Did the Math on AI's Energy Footprint. Here's the Story You Haven't Heard." *MIT Technology Review*, May 20, 2025. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>.
- Park, Sang-O, Ju-Won Kim, and Sun-Young Kim. "Case of Cardiac Arrest Due to Carbon Dioxide Poisoning Following an Explosion of a Carbon Dioxide Tank." *Clinical and Experimental Emergency Medicine* 7, no. 4 (2020): 336-339. <https://doi.org/10.15441/ceem.19.015>.
- Schlosberg, David. *Defining Environmental Justice: Theories, Movements, and Nature*. Oxford: Oxford University Press, 2007.
- Schwartz, Roy, Jesse Dodge, Noah A. Smith, and Oren Etzioni. "Green AI." *Communications of the ACM* 63, no. 12 (2020): 54-63.
- Sparrow, Robert. "Killer Robots," *Journal of Applied Philosophy* 24, no. 1 (2007): 62-77.
- Stetter, Christian, et al. "Therapeutic Hypercapnia for Prevention of Secondary Ischemia After Severe Subarachnoid Hemorrhage: Physiological Responses to Continuous Hypercapnia." *Scientific Reports* 11 (2021): 1-12. <https://www.nature.com/articles/s41598-021-91007-7.pdf>.
- Strasser, Anna. "Distributed Responsibility in Human-Machine Interactions." *AI and Ethics* 2 (2022): 523-532.
- The Editors of Encyclopaedia Britannica. "Deontological Ethics." *Encyclopaedia Britannica*. Last modified October 22, 2025. <https://www.britannica.com/topic/deontological-ethics>.
- UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.
- U.S. Environmental Protection Agency. "Climate Change Indicators: Atmospheric Concentrations of Greenhouse Gases." Accessed November 4, 2025.

<https://www.epa.gov/climate-indicators/climate-change-indicators-atmospheric-concentrations-greenhouse-gases>.

U.S. Environmental Protection Agency. "Sources of Greenhouse Gas Emissions." Accessed July 15, 2026. <https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions>.

Vallor, Shannon. *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. New York: Oxford University Press, 2016.

West, Henry R. "Utilitarianism." *Encyclopædia Britannica*. Last modified 2025. <https://www.britannica.com/topic/utilitarianism-philosophy>.