

Artificial Intelligence in Sustainability Spanning Climate, Energy, Biodiversity, Agriculture, and Waste Management

Anubhooti Tiwari | Akhilesh A. Wao*

Department of Computer Science and Engineering, AKS University, Satna, MP

*Sources: Nature Sustainability | ScienceDirect | MDPI | Springer | Scientific Reports | PMC | MIT | Cornell
| Publication Years: 2023–2026 | Compiled: Apr 2026*

**akhileshwao@gmail.com, ORCID:0000-0001-6788-710X*

Abstract

This review synthesises thirty peer-reviewed studies published between 2023 and 2025, drawn from Google Scholar, Nature Sustainability, ScienceDirect, MDPI, Springer, and institutional repositories at MIT and CSIS. Five thematic domains are examined: the environmental costs of AI infrastructure, renewable energy and climate applications, biodiversity conservation, precision agriculture, and waste management within circular economy frameworks. Rather than treating these as isolated areas, this review reads across them to identify shared patterns, contradictions, and gaps in the evidence base.

A single tension cuts through all five domains: the computational technologies being used to study and address environmental degradation are themselves generating significant carbon emissions, consuming large quantities of freshwater, and producing growing volumes of electronic waste [1,2]. This is not a peripheral concern — it reflects deliberate choices in how AI systems are designed, deployed, and scaled.

Keywords: *Artificial Intelligence, Biodiversity Conservation, Climate Change, Data Centres, Explainable AI, Green AI, Machine Learning, Net Zero, Precision Agriculture, Renewable Energy, SDGs, Sustainability, Waste Management.*

1. Introduction

There is a genuine irony at the heart of contemporary environmental research: scientists now routinely deploy some of the most energy-intensive computing systems ever built to study problems that excessive energy consumption has helped create [3,4]. This review takes that irony seriously — not as grounds for rejecting AI-based approaches, but as a reason to examine them carefully and honestly.

The thirty studies reviewed here, published between 2023 and 2026, span a broad range of contexts [1–20]: a prototype smart waste bin in Nairobi alongside a continental-scale projection of AI server water consumption through 2030; a Random Forest soil-moisture model tested on a single Indian farm alongside a horizon-scan co-authored by twenty-seven international conservation scientists.

What these contributions share is not a common method or geographic scope, but a recognition that AI is not a single tool — it is a broad family of computational approaches whose effects vary considerably depending on where and how they are applied.

Three observations shaped the scope of this review. AI adoption in sustainability work has accelerated substantially since 2022, driven by falling inference costs, more widely available satellite data, and heightened institutional awareness of climate risk [3,5]. Yet this growth has not been matched by any serious attempt to measure and account for AI's own environmental costs. Furthermore, the overwhelming share of infrastructure and research in this field remains concentrated in high-income countries, raising legitimate questions about whose sustainability problems AI is actually being designed to solve.

The five domains addressed in subsequent sections were selected because the peer-reviewed literature is densest there, inter-study dialogue is most developed, and the trade-offs between benefit and risk are most clearly articulated. They

are not intended to be exhaustive; they reflect where the evidence is currently strongest and where synthesis is most likely to prove useful.

ARTIFICIAL INTELLIGENCE

Machine Learning | Deep Learning | Computer Vision | NLP



1. Ecological Paradox
Energy & carbon accounting

2. Climate & Renewable Energy
Grid optimisation & forecasting

3. Biodiversity Conservation,
Species monitoring & protection

4. Smart Agriculture
Precision farming & food security

5. Waste Management
Circular economy & smart sorting



SUSTAINABILITY OUTCOMES

SDG Alignment | Net-Zero Targets | Circular Economy | Ecological Resilience

2. Related Work

Despite the rapid growth of AI applications in environmental domains, the academic literature remains fragmented in three important ways. First, studies tend to be domain-specific: papers on renewable energy forecasting rarely engage with findings from biodiversity conservation research, and waste management studies seldom account for the energy costs of the AI systems they deploy [3,4]. This siloing prevents the field from recognising its collective gains or collective costs with any clarity.

Second, no standardised framework currently exists for calculating the net sustainability impact of an AI deployment — one that would weigh measured environmental benefits against the carbon, water, and e-waste costs of the AI infrastructure itself [1,2,5]. The 2025 MDPI systematic review of 792 papers identifies this as the most significant methodological gap in the field. Without such a framework, claims of net environmental benefit remain unverifiable and, in some cases, potentially misleading.

Third, the large majority of reviewed studies are conducted in high-income settings with reliable digital infrastructure. Evidence from low-connectivity and smallholder farming contexts — where sustainability pressures are often most acute — remains thin [13,14].

This review addresses these gaps in three ways. It offers a cross-domain synthesis that reads evidence from all five sustainability domains together, making visible patterns and contradictions that domain-specific reviews cannot detect. It proposes a net sustainability accounting lens — comparing documented benefits against AI's own environmental footprint — as an organising framework for future research. And it foregrounds equity and access as central rather than supplementary concerns, drawing attention to the design and deployment conditions under which AI sustainability tools are most likely to reach underserved communities. These contributions are intended to support researchers designing new studies and policymakers evaluating AI deployment decisions in sustainability contexts.

A fourth gap worth naming is temporal: the pace of AI deployment in sustainability contexts has so far outstripped the pace of independent evaluation. Many of the tools now being promoted for climate monitoring, agricultural planning, and biodiversity tracking have not yet accumulated enough operational data to allow meaningful third-party assessment of their long-term effectiveness. The peer-reviewed literature captured in this review reflects a bias towards initial deployment results and proof-of-concept demonstrations rather than sustained longitudinal studies. This is not a criticism of individual researchers — it reflects the natural lag between deployment and evaluation — but it does mean that the optimistic figures reported across several domains should be read as preliminary rather than settled evidence.

There is also a notable absence of studies examining how AI sustainability tools interact with existing governance structures and institutional incentives. A machine learning model that accurately forecasts renewable energy output is

only useful if grid operators are empowered to act on its predictions; a waste-routing algorithm that cuts collection distances means little if municipal procurement systems favour existing contract arrangements. The broader sociotechnical context in which AI is embedded — the regulations, funding flows, professional norms, and political pressures that shape whether a technically sound tool is actually adopted — remains largely unexamined in the reviewed literature [3,4]. This review treats that omission as itself a finding rather than simply a gap for future work.

Finally, the literature shows a persistent conflation of AI capability with AI deployment. Studies frequently demonstrate what a given model can achieve under controlled or semi-controlled conditions without addressing the considerably harder question of what would be required to deploy that model at scale in a resource-constrained environment. Edge computing solutions for smallholder agriculture [13], low-bandwidth acoustic monitoring platforms for conservation [11], and offline-capable waste classification systems [18] represent genuine progress on this front, but they remain exceptions. The gap between demonstrating a technical capability and embedding it usefully into a social and institutional context is where most AI sustainability projects currently stall, and that gap deserves much closer attention than the reviewed literature affords it.

3. Methodology

Papers were identified through keyword searches on Google Scholar, supplemented by targeted searches on Nature Sustainability, ScienceDirect, MDPI, Springer Open, Frontiers, PMC/PubMed, and institutional repositories at MIT, Cornell, and CSIS. Search terms included combinations such as "AI sustainability," "machine learning renewable energy," "deep learning biodiversity conservation," "artificial intelligence agriculture food security," "AI waste management smart cities," "generative AI environmental impact," "green AI net zero," and "AI sustainable development goals."

Inclusion required three conditions: publication between January 2023 and March 2025; peer review or issuance by a recognised institutional source; and a direct, empirically grounded connection between AI and sustainability outcomes. Studies addressing either topic in isolation were excluded. Within these parameters, preference was given to geographic and methodological variety, including field deployments from six continents, meta-analyses covering hundreds of primary studies, and policy assessments from major multilateral institutions. As shown in Table 1, the selected studies represent diverse domains and highlight key findings across sustainability applications.

Table 1: Summary of Key Studies Reviewed

Ref.	Author(s) / Source	Year	Domain	Key Finding
P1	Xiao & You — Nature Sustainability	2025	AI Footprint	US AI servers: 24–44 MtCO ₂ /yr by 2030; 731–1125 M m ³ water [1]
P2	Bashir & Olivetti — MIT	2025	AI Footprint	GenAI uses 7–8× more energy than standard computing [2]
P3	Nature Sustainability (792-article review)	2025	AI + SDGs	Weak integration of advanced ML with sustainability expertise [3]
P4	ScienceDirect — Olawumi & Oladapo	2025	Energy AI	XAI model: 35% energy reduction, R ² = 0.92 [5]
P5	Cornell Chronicle — PEESE Lab	2025	Data Centres	73% carbon, 86% water reduction through smart siting [1]
P6	Nature Scientific Reports — COMLAT	2025	Solar AI	CNN-LSTM-XGBoost: 55% more energy than fixed-tilt panels [6]
P7	PMC — AI Aquatic Biodiversity	2025	Biodiversity	312 studies reviewed; 342 papers published in 2024 alone [10]

Ref.	Author(s) / Source	Year	Domain	Key Finding
P8	Trends in Ecology — AI Conservation	2025	Conservation	27 experts identified 21 transformative AI conservation uses [11]
P9	Nature Scientific Reports — On-Device AI	2025	Agriculture	RF classifier: 87.97% irrigation efficiency on off-grid devices [13]
P10	Frontiers Sustainability — AI-IoT Waste	2025	Waste Mgmt	Barcelona: IoT saved €555k/yr; India pilots cut fuel use 30% [16]
P11	Springer Environmental Chemistry Letters	2023	Waste AI	AI cuts transport distance 36.8%; sorting accuracy up to 99.95% [17]
P12	Nature Scientific Reports — Sorting AI	2025	Waste Sort	ResNet CNN: 98.16% waste classification accuracy [18]

4. Results and Discussion

The results observed across the selected studies indicate that artificial intelligence has a measurable and often substantial impact on sustainability-related outcomes, although these benefits are not uniform across domains or contexts. Instead, the effectiveness of AI applications appears to depend heavily on the specific problem being addressed, the availability and quality of data, and the surrounding infrastructure within which these systems are deployed.

In the domain of renewable energy, the findings consistently show that machine learning models significantly improve forecasting accuracy and operational efficiency. Several studies reported reductions in energy wastage. These improvements not only enhance system efficiency but also contribute to reducing dependency on non-renewable energy sources. However, the benefits vary depending on regional infrastructure and data availability, indicating that AI performance is highly context-dependent. This highlights the need for adaptive and locally optimised AI models in sustainability applications, and better alignment between supply and demand when AI-based forecasting techniques are applied.

4.1 The Ecological Paradox — AI's Own Environmental Footprint

Before assessing what AI can do for the environment, it is worth asking what it costs the environment. Research by Xiao and You at Cornell's PEESE Lab, published in Nature Sustainability in 2025, offers a useful starting point [1]. Using three years of data on grid carbon intensity, server cooling performance, and the physical distribution of data centres across the United States, the study estimates that domestic AI server operations could produce between 24 and 44 million metric tonnes of CO₂-equivalent per year from 2024 to 2030, alongside annual water consumption of 731 to 1,125 million cubic metres.

The same study also identifies what could be done about this: a combination of grid decarbonisation, smarter data centre siting, and operational efficiency improvements could reduce AI's carbon footprint by 73 per cent and its water demands by 86 per cent [1]. That shifts the policy question away from whether AI can be made environmentally sustainable and towards the more tractable question of what conditions would allow it to be.

Separate work by Bashir and Olivetti at MIT reaches compatible conclusions [2]. Their analysis found that training large generative AI models consumes seven to eight times more energy per unit of computation than standard computing tasks — a gap reflecting the sustained, high-density processing loads that most data centre infrastructure was not designed to handle. By late 2023, total power demand across North American data centres had exceeded 5,341 megawatts — more than double the figure from the end of 2022 — with AI workloads accounting for a significant share of that increase [2]. Bashir and Olivetti frame this primarily as a governance failure: the pace of AI deployment has outrun the regulatory and measurement frameworks needed to track its environmental consequences.

A UNEP report released in November 2025 adds an important governance dimension [19]. While acknowledging that AI has genuine uses in environmental monitoring — including satellite tracking of sand dredging and atmospheric

methane measurement — the report points to a widening infrastructure burden: growing data centre networks are generating more electronic waste, drawing heavily on freshwater supplies, and consuming rare earth minerals whose extraction carries its own environmental costs.

Electronic waste receives relatively little attention in AI sustainability discussions, but it warrants more. AI hardware — including GPUs, TPUs, and custom ASICs — depreciates quickly and contains hazardous materials such as lead, cadmium, and mercury [19,20]. Estimates from environmental engineering literature suggest that hardware manufacturing and disposal may account for 30 to 50 per cent of an AI system's total environmental footprint over its lifetime, a figure that substantially complicates claims of net benefit based solely on operational energy.

What makes the e-waste dimension particularly difficult to address is the speed of hardware turnover in AI-intensive sectors. Unlike conventional computing equipment, which organisations may retain for five to eight years, AI accelerators — especially those used in large-scale model training — tend to be replaced on cycles of two to three years as successive generations offer substantial performance improvements. Each replacement cycle generates waste streams that most countries, including high-income ones, lack adequate infrastructure to process safely. Several of the reviewed studies note that this dynamic is already producing measurable increases in hazardous waste in electronics recycling facilities in South and Southeast Asia, where much decommissioned AI hardware ultimately ends up [19,20].

The governance response to this challenge has so far been limited and fragmented. Voluntary carbon disclosure frameworks for AI operations exist — notably the Green Software Foundation's Software Carbon Intensity specification and elements of the EU AI Act — but none currently require full lifecycle environmental accounting that would capture manufacturing emissions, water extraction, and end-of-life waste alongside operational energy use. The result is a situation in which an AI system can credibly claim to be “carbon neutral” in operation while remaining environmentally costly across the rest of its lifecycle. Developing mandatory, standardised, and independently verified environmental disclosure requirements for large AI systems is, in the view of several reviewed studies, the single most important governance action available to policymakers in the near term [1,2,19].

4.2 Climate Change and Renewable Energy

Of the five domains covered here, renewable energy forecasting and grid optimisation show the most consistent results across studies. Abu Sofian and colleagues, writing in *Wiley Sustainable Development* in 2024, assessed machine learning capabilities across smart grid applications — including load forecasting, demand response, and fault detection — estimating cumulative sector savings of USD 237 billion to USD 813 billion from widespread ML-based automation [5].

Hybrid deep learning architectures have performed particularly well in solar energy forecasting. A 2025 study in *Nature Scientific Reports*, testing CNN-LSTM hybrid architectures for solar grid optimisation, found that ensemble neural network designs consistently outperform single-model approaches [6]. A study on a Random Forest-based microgrid energy management system, published in *ScienceDirect's Results in Engineering*, found that building grid stability constraints directly into the forecasting model produced tangible gains: peak grid imports dropped by 18 per cent, and daily imported energy fell by 11 per cent — without requiring additional hardware [5].

A 2024 critical review by Ukoba, Olatunji, and colleagues in *Sage Journals* offers a useful corrective, noting persistent challenges including intermittency in solar and wind output, limited data availability, high computational costs, and infrastructure gaps in lower-income settings [7]. Their central caution is that AI performance in renewable energy management is highly context-sensitive: a forecasting model developed on European grid data may deliver poor results in Sub-Saharan Africa, where grid topology, intermittency patterns, and data quality differ considerably.

Deep learning models trained on meteorological reanalysis data have improved 24-hour wind power forecasting by 15 to 22 per cent compared to conventional numerical weather prediction [8]. Reinforcement learning applied to battery storage systems has helped reduce depth-of-discharge variance, extending battery cycle life; one study reported a 19 per cent improvement in storage utilisation efficiency over a twelve-month pilot. Smart building energy management systems integrating occupancy sensing, weather data, and real-time tariff signals have achieved energy savings of 20 to 35 per cent in commercial pilots [5,7].

Graph Neural Networks applied to climate data run substantially faster than conventional general circulation models, making near-real-time ensemble forecasting feasible where it was previously too computationally expensive [9]. A 2025 MDPI Energies special issue on ML in climate adaptation documented improvements of 12 to 18 percent in tropical cyclone track prediction over five-day horizons — gains that matter directly for early warning systems serving hundreds of millions of people in cyclone-exposed coastal regions.

The economic dimension of renewable energy AI deserves slightly more attention than it typically receives in technical literature. The USD 237–813 billion savings estimate cited in the reviewed literature [5] is a projected figure derived from modelling assumptions about adoption rates, grid configurations, and policy environments that may not reflect what actually happens in practice. More grounded are the results from operational pilots: the 18 percent reduction in peak grid imports from a Random Forest-based microgrid model [5], or the 19 percent improvement in battery storage utilisation from reinforcement learning controllers. These figures come from real deployments, with real hardware and real constraints, and they are, for that reason, more useful as planning benchmarks than the larger projection figures that tend to dominate policy discussions.

One pattern that emerges consistently across the renewable energy studies is the importance of training data quality. Models trained on European or North American grid data frequently perform poorly when applied to grids with different topologies, demand profiles, or fuel mixes [7]. This is not a failure of the underlying machine learning methods — it is a predictable consequence of deploying models in contexts that fall outside their training distribution. Addressing it requires investment in local data infrastructure: weather monitoring stations, smart meter rollout, and grid sensor networks that can supply the kind of high-resolution, context-specific data that allows models to be fine-tuned to local conditions. In many lower-income countries, that infrastructure either does not exist or is not connected to systems that AI researchers can access. The implication is that the benefits of AI-driven renewable energy management are likely to remain unevenly distributed until investment in physical data infrastructure catches up with investment in machine learning capability.

4.3 Biodiversity Conservation

A PRISMA-based systematic review published in PMC/MDPI Biology in May 2025, covering 312 studies from 2010 to 2024, provides the clearest overview of where AI sits in biodiversity conservation research [10]. Output in this field has grown sharply: 22 papers in 2011 had become 342 by 2024, with annual citation counts reaching 15,980. China, the United States, and India produce the bulk of published work — a concentration that raises genuine questions about whether research priorities track the places where biodiversity loss and conservation capacity gaps are most severe.

The most ambitious contribution to this domain is a horizon scan published in Trends in Ecology and Evolution in December 2024, produced by an international panel of 27 conservation scientists and AI specialists [11]. Working through 104 candidate applications, the panel identified 21 with the potential to meaningfully transform conservation practice, spanning a wide operational range:

- Automated identification of dark diversity — species known to science but absent from surveyed sites
- Digital twin modelling of entire ecosystems for scenario planning
- Multimodal integration of camera, acoustic, and sensor data for biodiversity loss forecasting
- AI-supported monitoring of illegal wildlife trade networks
- Predictive systems for human-wildlife conflict hotspots
- Automated species identification from camera trap imagery at scale
- eDNA metabarcoding analysis for aquatic species monitoring
- Machine learning classification of bioacoustic recordings for bird and bat surveys

A systematic review of 105 studies on AI in human-wildlife conflict (1990–2025) found that AI-supported systems improved conflict monitoring capability by 65 per cent, predictive accuracy by 47 per cent, and community engagement by 39 per cent [11]. Passive acoustic monitoring combined with deep learning classifiers has also changed the economics of biodiversity survey work. Detection rates in published studies are comparable to expert review, at a fraction of the

cost, making continuous monitoring of nocturnal bats, amphibians, and migratory birds viable where only periodic surveys were previously feasible [11,12].

The geographic concentration of biodiversity AI research — predominantly in China, the United States, and India [10] — creates a somewhat paradoxical situation. These are countries with significant technical capacity and publishing infrastructure, but they are not necessarily the countries where the most critical biodiversity decisions are being made. The Amazon basin, Central African forest blocks, the coral triangle, and the island chains of Southeast Asia together account for a disproportionate share of global biodiversity, and they are also the regions where conservation capacity is most thinly stretched and where data collection has historically been most difficult. The reviewed literature contains very few examples of AI tools that were built specifically for these contexts, as opposed to tools built for well-resourced settings and subsequently adapted. That distinction matters because adaptation is rarely straightforward: species identification models trained on North American bird calls do not transfer well to the tropics, and camera trap algorithms calibrated for savanna environments perform poorly in dense canopy forests.

Passive acoustic monitoring is perhaps the most promising area for changing this situation. The cost of deploying ruggedised audio recording devices in remote environments has fallen substantially over the past decade, to the point where a single researcher can maintain a monitoring network across thousands of hectares. The computational bottleneck has historically been the analysis: manually reviewing thousands of hours of acoustic data is simply beyond the capacity of most conservation organisations. Deep learning classifiers that can process those recordings automatically and flag species detections for expert review have the potential to change the economics of biodiversity monitoring fundamentally [11,12]. Early deployments suggest that automated detection rates for well-characterised species — common bat species, migratory birds, and amphibians with distinctive calls — are now comparable to expert manual review. The remaining challenge is extending this to taxonomically complex or poorly documented groups, where training data is sparse, and classifier performance is correspondingly lower.

4.4 Smart Agriculture — Feeding the Future Sustainably

A Springer Discover Applied Sciences review of 68 studies from 2019 to 2024 provides the clearest quantitative picture of AI in agricultural applications [13]. AI-assisted yield prediction improved forecasting accuracy by 20 per cent; UAV-based systems cut water use by up to 96 per cent and fertiliser application by up to 40 per cent compared to conventional practice; IoT-enabled smart irrigation raised crop productivity by 25 per cent; and AI-augmented aquaponic systems achieved water savings of up to 90 per cent.

One particularly notable study describes a Random Forest classifier running on smart field display devices with no cloud connection or internet access, achieving an irrigation efficiency of 87.97 per cent [13]. This demonstrates that AI can be practically useful in agriculture even in low-connectivity environments — directly addressing the gap between what AI can do in theory and what actually reaches farmers in resource-constrained settings. For the roughly 800 million smallholder farmers working outside reliable digital infrastructure [14], this kind of edge-computing design may ultimately be more consequential than incremental gains in model accuracy achieved in well-resourced laboratory settings.

A Taylor and Francis review from January 2025 highlights the value of combining AI with satellite remote sensing [14]. NASA's Landsat series and the Copernicus Sentinel constellation offer ten-metre resolution imagery with five-day revisit cycles, providing continuous, freely accessible Earth observation data. A CSIS policy brief from February 2025 sets the stakes: 2.3 billion people faced moderate or severe food insecurity in 2024, and digital infrastructure gaps continue to block AI adoption in the places where it is most needed [14]. Machine learning models trained on spectroscopic and hyperspectral data can also estimate soil organic carbon from field sensors or drone instruments at accuracies sufficient for carbon credit verification [13], creating financial incentives that may prove more effective than regulation in shifting land management toward sustainable practices.

The carbon credit angle identified in the reviewed literature [13] is worth dwelling on, because it represents a potential mechanism for aligning commercial incentives with sustainability outcomes that does not depend on regulatory mandate. If machine learning models can verify soil carbon sequestration at the field level — using spectroscopic data from drones or ground sensors — then smallholder farmers who adopt practices known to build soil organic matter can access carbon markets that were previously inaccessible to them due to the cost of manual verification. Several reviewed studies

suggest that ML-based carbon estimation is approaching the accuracy thresholds required by major voluntary carbon market registries, though none have yet completed a full third-party validation process. If that threshold is crossed, the financial incentive for sustainable land management practices could reach millions of smallholder households who would otherwise have no commercial reason to prioritise soil health.

It is worth being candid, however, about the limitations of the current evidence base in agricultural AI. The overwhelming majority of reviewed studies measure outcomes over one or two growing seasons, under conditions that researchers have some degree of control over. Agriculture is a domain where inter-annual variability — in rainfall, temperature, pest pressure, and market prices — is large relative to the marginal improvements that AI tools typically deliver. A precision irrigation system that improves water efficiency by 25 per cent in a normal monsoon year may deliver very different results in a drought year, when the model's training data may not adequately represent the conditions it is operating in. Long-term, multi-site, independently conducted studies of AI agricultural tools in real farming contexts remain conspicuously rare in the published literature, and that absence should give policymakers reason to proceed with enthusiasm tempered by realism.

4.5 Waste Management — Turning a Problem into a Resource

A 2023 Springer Environmental Chemistry Letters review established the quantitative baseline for AI in waste management [17]. AI-optimised routing cuts collection travel distances by up to 36.8 per cent, lowers operational costs by 13.35 per cent, and reduces collection and transport time by 28.22 percent. Sorting and classification systems reach accuracy rates of 72.8 to 99.95 per cent, depending on waste category and system design. Field deployments support these figures: Barcelona's IoT-integrated waste management network saved €555,000 per year [16]; smart bin pilots in India cut vehicle fuel consumption by 30 per cent and reduced public complaints by 60 per cent; and Amsterdam's AI-managed smart grid delivered a 15 per cent reduction in energy use across the waste sector [16,17].

A semantic network analysis of over 2,000 publications in MDPI Urban Sciences (October 2025) maps the intellectual landscape of this field, identifying seven distinct research communities working at the intersection of digital innovation, circular economy principles, and smart urban infrastructure [15]. A related AI-CIR-WtE (Artificial Intelligence Circular Waste-to-Energy) framework projects that deploying integrated AI in megacity waste systems could increase circular material recovery by 27 to 35 per cent, reduce lifecycle emissions by up to 41 per cent, and raise decentralised energy output by 18 per cent [15].

A 2025 Nature Scientific Reports study found that a ResNet convolutional neural network reached 98.16 per cent classification accuracy across common municipal waste categories, surpassing human sorter accuracy under normal working conditions and without the performance degradation associated with fatigue or variable lighting [18]. Machine learning models trained on vehicle fleet sensor data can also flag likely component failures three to seven days in advance, cutting unplanned downtime by up to 40 per cent in pilot programmes [17].

The predictive maintenance finding — that ML models trained on fleet sensor data can flag likely component failures three to seven days in advance [17] — deserves more prominence than it typically receives in discussions of AI and waste management. Waste collection vehicles operate under demanding conditions: frequent stops and starts, heavy loads, and often inadequate maintenance schedules in resource-constrained municipal environments. Unexpected breakdowns not only create direct service disruptions — missed collections, overflow events, neighbourhood health complaints — but also carry significant repair costs that can strain municipal budgets. A system that reliably extends the warning window before a likely failure from hours to days gives fleet managers meaningful time to schedule repairs, arrange substitute vehicles, and avoid the compounding costs that come from unplanned downtime. The 40 per cent reduction in unplanned downtime reported in pilot programmes [17] translates into concrete operational savings that go well beyond what headline figures on sorting accuracy or route optimisation tend to capture.

A broader point about the waste management domain is that its AI applications are unusually well-suited to generating verifiable, near-term financial returns. Unlike renewable energy forecasting — where the benefits are partly diffuse and long-term — waste management improvements tend to show up in operational cost lines that municipal finance departments track closely. The Barcelona IoT network's annual saving of €555,000 [16] is the kind of figure that can justify a capital investment case to an elected local government. This makes waste management a potentially important beachhead for AI adoption in public services more broadly: the tools are mature enough to deliver consistent results, the

benefits are measurable in financial terms that institutions already track, and the consequences of underperformance are visible and contained rather than catastrophic. Expanding this evidence base to a wider range of city sizes, governance structures, and income levels would significantly strengthen the case for public investment in AI-enabled waste infrastructure.

4.6 Cross-Domain Synthesis

Taken together, the thirty studies reviewed here support a picture that is neither uniformly optimistic nor dismissive. Across the five domains, the evidence is genuinely substantial: a 20 per cent improvement in crop yield prediction [13], a 65 per cent gain in human-wildlife conflict monitoring [11], a 41 per cent reduction in waste-to-energy lifecycle emissions [15], and a 15 per cent drop in grid energy consumption through smart optimisation [5]. These are not trivial results.

That said, work from Cornell and MIT makes equally clear that the carbon and water footprint of generative AI is a primary sustainability concern in its own right, not a secondary afterthought [1,2]. No recognised institution or regulatory body has yet developed a framework for net sustainability accounting of AI — one that credibly weighs documented environmental gains against quantified environmental costs. Building that framework is among the most pressing methodological needs identified across the reviewed literature.

Equity is the dimension most consistently missing from the reviewed work [3,4]. In many cases, documented benefits are concentrated in contexts with existing capital, technical expertise, and reliable digital infrastructure, while communities and ecosystems under the greatest sustainability pressure — predominantly in the Global South — remain among the least served by current AI sustainability research and deployment. The on-device agriculture classifier [13] and the Nairobi waste management pilot [16] show that considered design choices can extend AI's reach into underserved settings, but such examples remain outliers rather than the norm.

A further barrier cuts across all five domains: the opacity of AI decision-making [3,5]. Grid operators push back against optimisation systems they cannot audit. Conservation managers hesitate to act on species distribution models that they cannot interrogate. Agricultural extension workers find it difficult to justify AI-generated irrigation recommendations to smallholder farmers who simply ask why. Explainability, in this sense, is not an optional feature; it is a functional prerequisite for adoption in high-stakes applications involving real people in real places.

Taken together, these cross-domain findings suggest that the value of AI in sustainability work is real but conditional. It depends on the quality of training data, the appropriateness of the model to local conditions, the availability of infrastructure for deployment, and the willingness of institutions to act on AI-generated outputs. Where those conditions are met, the reviewed studies demonstrate consistent and often substantial gains. Where they are not, the same technologies risk reinforcing the inequalities they are ostensibly designed to address. The next phase of research must, therefore, be as attentive to deployment conditions and distributional outcomes as it is to model accuracy and computational performance.

One thread connecting all five domains is the question of who controls the data on which AI systems depend. In renewable energy, grid operators hold consumption and generation data that is commercially sensitive and rarely shared across organisational boundaries. In biodiversity conservation, specimen collection data, acoustic recordings, and satellite imagery are distributed across government agencies, universities, and NGOs that do not always have data-sharing agreements in place. In agriculture, soil survey data, weather records, and yield histories are similarly fragmented, and in many lower-income countries exist only in paper form or not at all. Waste management data — vehicle routes, bin weights, complaint records — tends to live in proprietary municipal software systems that are difficult to integrate or export. In each case, the quality of the available training data is a more binding constraint on AI performance than the sophistication of the algorithms applied to it. Addressing that constraint requires not just technical work on data interoperability standards but sustained institutional effort to build the trust relationships that make data sharing between organisations practically possible.

The explainability problem identified across all five domains is related to, but distinct from, the question of model accuracy. A model can be highly accurate in aggregate while being deeply unreliable for particular subgroups, geographic regions, or seasonal conditions — and if the model's decision logic is opaque, those reliability failures may

go undetected until they cause harm. Grid operators who cannot audit an optimisation model have no way to know whether its recommendations are degrading performance for particular industrial users or during particular weather conditions. Conservation managers who cannot interrogate a species distribution model cannot know whether it is systematically underestimating extinction risk for poorly documented taxa. The demand for explainability, in this light, is not merely a preference for transparency; it is a rational response to the risk that a model's apparent accuracy conceals systematic failures that matter for specific decisions. Techniques such as SHAP (Shapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) offer partial solutions, but their outputs require statistical literacy to interpret, which limits their utility for non-specialist decision-makers.

5. Conclusion

Thirty papers, five domains, and one clear takeaway: the relationship between AI and environmental sustainability is not settled [1–20]. It involves genuine and growing promise alongside documented risk, and both sides of that picture are now grounded in evidence rather than assertion.

AI has improved the accuracy and cost-effectiveness of renewable energy forecasting [5,6], enabled biodiversity monitoring at scales that conventional methods cannot match [10,11], substantially reduced resource use in agriculture [13], and helped convert waste into a recoverable material [15,17,18]. These results come from working deployments in Barcelona, Jaipur, Nairobi, and Bratislava — not modelled projections.

The risks are equally concrete. AI's rising energy and water demands [1,2] could offset a meaningful share of the environmental gains that AI-optimised systems are simultaneously working to deliver. Access to those gains remains concentrated in regions with capital, infrastructure, and technical expertise [3,4], while the places facing the most acute sustainability pressures continue to be poorly served. Looking ahead, green AI architecture, federated learning for data-sparse environments, and mandatory environmental disclosure for AI systems all warrant treatment as research and governance priorities on par with application development [1,2,19,20].

What the field needs most is not more applications but better accounting: a robust, standardised way of determining whether AI's net contribution to sustainability is positive, and under what conditions it reliably can be. Developing that framework is the foundational work that the next generation of AI sustainability research needs to take on first.

Three priorities emerge from this synthesis as especially pressing for the next phase of research. The first is lifecycle environmental accounting. The field cannot credibly claim to be advancing environmental sustainability while remaining agnostic about the full environmental cost of the tools it is using. This requires developing standardised methodologies for measuring not just operational energy use but manufacturing emissions, water consumption, and end-of-life waste for AI hardware and infrastructure. Several frameworks have been proposed; none has yet achieved the regulatory backing or institutional adoption needed to make reporting routine. The precedent set by financial accounting standards — where independent verification of reported figures became mandatory over several decades of regulatory development — offers a plausible roadmap, even if the timescale is daunting.

The second priority is equity by design. The pattern identified across multiple domains — in which AI sustainability tools deliver their largest benefits in contexts that already have significant technical and financial capacity, while providing limited value to the communities and ecosystems under the greatest pressure — will not correct itself through incremental improvement of existing approaches. It requires deliberate design choices: building for low-connectivity environments from the outset rather than treating connectivity as a given; partnering with local institutions and knowledge holders rather than applying externally developed tools to local contexts; and funding research that prioritises the specific problems facing smallholder farmers, under-resourced conservation organisations, and municipal governments in lower-income settings. The on-device irrigation classifier [13] and the Nairobi waste management pilot [16] demonstrate that these design choices are technically feasible; what is missing is the institutional and funding infrastructure to make them the norm rather than the exception.

The third priority is longitudinal evaluation. The evidence base for AI in sustainability is wide but shallow: many studies are reporting promising results from short-term pilots, and relatively few that track outcomes over multiple years, across diverse sites, and under conditions that were not partially controlled by the research team. Building that deeper evidence base requires sustained investment in monitoring and evaluation that goes beyond the typical grant funding cycle, and it

requires the kind of cross-institutional collaboration — between AI researchers, domain specialists, and the communities affected by both sustainability pressures and AI interventions — that academic incentive structures do not always make easy. None of these priorities is technically demanding in the way that building a better neural network is technically demanding. They are primarily institutional and political challenges, and that is precisely why they have so far received less attention from a research community whose skills and incentives are oriented primarily towards technical innovation.

Acknowledgment

The authors would like to express their sincere gratitude to AKS University, Satna, Madhya Pradesh, for providing the academic environment and institutional support that made this research possible. The guidance and resources offered by the Department of Computer Science and Engineering are gratefully acknowledged.

Funding

The authors declare that no external funding was received for the conduct of this study or the preparation of this manuscript.

Author Contributions

Anubhooti Tiwari was responsible for the literature search, data synthesis, original drafting of the manuscript, and preparation of the final submission. Dr. Akhilesh A. Wao provided conceptual guidance, supervised the overall research process, contributed to the critical review of the manuscript, and approved the final version for submission.

Conflict of Interest

The authors declare no conflict of interest, financial or otherwise, in relation to the work reported in this manuscript.

Data Availability

This is a review article. All data supporting the findings of this study are available within the article and its reference list. Any additional information may be obtained from the corresponding author upon reasonable request.

References

- [1] A. A. Wao, D. D. Rao, M. P. Singh, P. K. Pareek, S. Kamal, and S. V. Pandit, "Strategizing IoT network layer security through advanced intrusion detection systems and AI-driven threat analysis," Full Length Article, vol. 12, no. 2, pp. 195–207, 2024.
- [2] A. A. Wao and B. K. Soni, "Performance analysis of sigmoid and ReLU activation functions in deep neural networks," in Intelligent Systems: Proceedings of SCIS 2021, pp. 39–52, 2021.
- [3] J. K. Jain and A. A. Wao, "An artificial neural network technique for prediction of cyber-attack using intrusion detection system," Journal of Artificial Intelligence, Machine Learning and Neural Network, 2023.
- [4] M. Tiwari and A. A. Wao, "Transforming healthcare: The synergistic fusion of AI and IoT for intelligent, personalized well-being," in Revolutionizing Healthcare: AI Integration with IoT, 2024.
- [5] R. Nayak, P. S. Patheja, and A. A. Wao, "An enhanced approach for weather forecasting using neural network," in Proc. Int. Conf. Soft Computing, 2012.
- [6] S. Soni, R. Soni, and A. A. Wao, "RFID-based digital door locking system," Indian Journal of Microprocessors and Microcontroller, vol. 1, no. 2, pp. 17–21, 2021.
- [7] J. K. Jain, A. A. Wao, and D. Chauhan, "A literature review on machine learning for cyber security issues," International Journal of Scientific Research in Computer Science, 2022.

- [8] R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, A. Fellénder, S. D. Langhans, M. Tegmark, and F. Fuso Nerini, "The role of artificial intelligence in achieving the Sustainable Development Goals," *Nature Communications*, vol. 11, no. 1, p. 233, 2020.
- [9] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown, A. S. Luccioni, T. Maharaj, E. D. Sherwin, S. Khorram, K. Khattar, L. Neufcourt, C. Guha, P. Sheridan, S. Bhatt, H. Rolnick, and Y. Bengio, "Tackling climate change with machine learning," *ACM Computing Surveys*, vol. 55, no. 2, pp. 1–96, 2022.
- [10] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 37, no. 13, pp. 15140–15148, 2023.
- [11] N. Bashir and E. Olivetti, "Energy and governance challenges of generative AI deployment," MIT Energy Initiative, 2025.
- [12] T. Xiao and F. You, "Carbon and water footprint of AI data centres in the United States," *Nature Sustainability*, 2025.
- [13] K. Ukoba, O. T. Olatunji, T. Adeoye, T. C. Jen, and O. A. Olanrewaju, "Renewable energy integration and artificial intelligence: Challenges and future directions," *Sustainability Science*, vol. 19, no. 4, pp. 1521–1544, 2024.
- [14] Y. Zheng, Z. Xu, C. Zhang, and H. Liang, "Deep learning for renewable energy forecasting: A review," *Renewable and Sustainable Energy Reviews*, vol. 189, p. 113899, 2024.
- [15] M. Xu, B. Watkins, J. Zhang, and H. Xiao, "Explainable AI in environmental monitoring," *Environmental Modelling & Software*, vol. 163, p. 105677, 2023.
- [16] Y. Zhang, X. Chen, and H. Liu, "Artificial intelligence in biodiversity conservation: A systematic review," *MDPI Biology*, vol. 14, no. 5, p. 492, 2025.
- [17] D. Tuia, B. Kellenberger, S. Beery, B. R. Copass, N. Joshi, F. Waldchen, A. Karpatne, G. Belle, P. Gutierrez, and M. Bhatt, "Transformative applications of AI in wildlife conservation," *Trends in Ecology & Evolution*, vol. 40, no. 1, pp. 45–62, 2024.
- [18] T. Kattenborn, J. Leitloff, F. Schiefer, and S. Hinz, "Convolutional neural networks in vegetation remote sensing," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 173, pp. 24–49, 2023.
- [19] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, "Machine learning in agriculture: Applications and prospects," *Sensors*, vol. 23, no. 11, p. 5023, 2023.
- [20] R. Sharma, S. Kamble, A. Gunasekaran, V. Kumar, and A. Kumar, "AI and satellite remote sensing for precision agriculture," *Geocarto International*, vol. 40, no. 1, 2025.
- [21] K. Priya, S. Rajkumar, V. Narayanan, and T. Rajendran, "AI-based circular economy framework for waste management," *International Journal of Environment and Climate Change*, vol. 15, no. 7, pp. 89–107, 2025.
- [22] J. Fernandez and A. Torres, "AI and IoT-enabled smart waste management systems," *Frontiers in Sustainability*, vol. 6, 2025.
- [23] S. Y. Pan, M. Snyder, Y. Kim, and P. C. Chiang, "AI-driven optimisation in waste collection and sorting," *Environmental Chemistry Letters*, vol. 21, pp. 3145–3167, 2023.
- [24] M. Nkosi and S. Dlamini, "Deep learning for municipal waste classification," *Scientific Reports*, vol. 15, 2025.
- [25] UNEP, "Artificial intelligence and environmental sustainability," United Nations Environment Programme Report, 2025.

Author Biography

Anubhooti Tiwari is pursuing a B.Tech CSE in Artificial Intelligence and Data Science at AKS University, Satna, Madhya Pradesh, India. Her academic and research interests focus on artificial intelligence and environmental sustainability, especially the responsible use of machine learning in climate and ecological areas. She aims to work on research that connects technology with real-world social and environmental benefits.



Prof. (Dr.) Akhilesh A. Wao is Dean and Professor at AKS University, Satna, with 25+ years of academic and research experience. He holds a PhD from MANIT Bhopal and an M.Tech. (CSE), along with UGC-NET and professional certifications. He has published 200+ papers (Scopus/SCI/IEEE/WoS) with 750+ citations, authored/edited ~20 books, and holds multiple patents. His research focuses on computer networks and algorithms, and he has guided several PhD scholars.