

Generative AI's Carbon Paradox: Balancing Emissions and Climate Mitigation Through Lifecycle Optimization

Gabriel Silva-Atencio, PhD¹ 

¹ Universidad Latinoamericana de Ciencia y Tecnología (ULACIT), San José, Costa Rica, gsilvaa468@ulacit.ed.cr

Abstract – It is still not clear how to solve the carbon problem of generative artificial intelligence (GenAI), which is when practical emissions decreases fight with life-cycle carbon debt from training, hardware, and return effects. Integrating process-based accounting (ISO 14040/44) with real-time operating data and physics-informed neural networks (PINNs) is what this study suggests as a mixed life-cycle assessment (LCA) approach. The empirical study looks at 12,540 GPU-hours from 36 different sector-specific applications in areas like industry, transportation, farmland, and buildings. The results show that domain-specific designs lower carbon intensity by 2.3 times compared to general-purpose models (95% CI: 1.9 to 2.7). Scheduling tasks with carbon awareness cuts operating emissions by 28% (± 3.2 pp). 18.7 to 38.2% of life-cycle pollution come from the making of gear. Behavioral rebound effects lower the expected efficiency gains by 12.1% (± 2.3 percentage points), ranging from 3.2% in the energy sector to 12.1% in the factory sector. The mean absolute percentage mistake for the combination method is 6.7%, which is 58% better than the standard LCA. Net balance is mostly determined by grid carbon intensity (sensitivity index = 0.72), which makes area benefits vary by a factor of 0.45 to 3.3. Carbon-aware computer standards and geospatially focused green rewards are two policy tools that could increase net environmental gain by a factor of 2.1 to 3.3.

Keywords-- Carbon footprint, Generative AI, Life cycle assessment, Policy optimization, Sustainable computing, Workload scheduling.

I. INTRODUCTION

One of the most important problems of the time is the climate issue that modern society have made. Reference [1] say that the state-of-the-art need to use new technologies in ways that have never been seen before to cut down on carbon pollution and stop catastrophic warming. There are two sides to generative artificial intelligence (GenAI): it is praised for making systems more efficient, but it is also blamed for using a lot of computer power without being properly accounted for.

One big "carbon paradox" is that GenAI applications in areas like energy, manufacturing, and transportation promise to cut emissions by a lot, but they also create a lot of carbon debt from the time they are made to the time they are used [2, 3]. Although promising concepts exist, they often lack by thorough real-world studies of how it will affect the environment as a whole [4, 5].

A close examination of prior work reveals three gaps in methodology and analysis that this research fills. A lot of the lifecycle assessments (LCAs) that are already out there do not mix working, absorbed, and saved emissions into a single study approach [6, 7]. As you can see in Fig. 1, this leads to a

lot of doubt. Standard LCA methods make mistakes of up to 18.7%, mostly when it comes to time and place details.

Second, there is a clear gap between technical study and models that are ready for policymaking. Foundational studies have measured AI's energy use [8], but less than 9% of artificial intelligence (AI) climate studies look at how well policies can be put in place, and AI companies still don't disclose their carbon emissions voluntarily [9, 10].

Third, a major flaw in the research is that it does not look enough at global equity aspects. The amount of carbon that AI infrastructure uses is not the same everywhere. Data centres in areas that depend on fossil fuels, like those in the Global South, can have three times bigger carbon footprints than those in grid-rich, renewable-powered areas of the Global North [11, 12] but this difference in geography and ethics is rarely taken into account when figuring out net-impact.




FIG. 1. LCA UNCERTAINTY AND ERROR REDUCTION META-ANALYSIS UTILIZING THE HYBRID TECHNIQUE

To fill in these gaps, this study produces and answers a main research question: *Under what specific design, practical, and geographic conditions can GenAI be used to reduce carbon emissions in a way that is net positive over its entire lifecycle?* A new hybrid LCA system is used to run the study. It combines process-based International Standard Organization (ISO) 14040/44-compliant accounting with high-frequency, real-time operating data. The methodological architecture shown in Fig. 2 is unique because it has a three-tier system boundary and uses physics-informed neural networks (PINNs) to add thermodynamic constraints directly to the optimization process. This reduces mean absolute error by 58% relative to conventional LCA methods.




FIG. 2. HYBRID LCA BOUNDARY CONCEPTUAL FRAMEWORK FOR GENERATIVE AI'S CARBON FOOTPRINT LIFECYCLE

A full analysis of 12,540 Graphic Processor Unit (GPU)-hours across thirty-six sector-specific methods forms the basis of this study. Table 1 shows the scope of this analysis. This large dataset allows for a level of confirmation that has never been seen before. Monte Carlo runs ($n=10,000$) are used to measure uncertainty, and a multi-objective Pareto frontier analysis is used to settle the natural trade-offs between model performance, economic cost, and environmental effect.

TABLE 1
EMPIRICAL BASIS: CASE STUDY SUMMARY (GPU-HRS)

Sector	Application	AI Model Type	Scale (GPU-hrs.)	Data Source	Deployment Type
Energy	Renewable Forecasting	Transformer	2,850	Grid Op. Data	Real-world
Manufacturing	Generative Design	Generative adversarial Networks (GAN)	3,100	Sensor Logs	Simulation
Transportation	Dynamic Routing	Reinforce Learning (RL) Agent	2,150	Fleet Telemetry	Real-world
Agriculture	Precision Irrigation	Vision Transformer (ViT)	1,450	IoT Sensors	Real-world
Buildings	Material Design	Diffusion Model	2,990	Building Management Systems	Simulation

This work makes three kinds of contributions that are meant to set a new standard for the field. It also has a mixed AI-LCA structure that clearly lowers error to $\pm 6.7\%$, which is 58% better than the old ways of doing things. It is the first set of guidelines for GenAI's net carbon balance in a certain field that include measures of doubt. You can also see a steady J-curve of emissions, which means that smart implementation can reduce original carbon debt.

There are steps that can be taken right away to make computers more aware of carbon emissions, gadgets more open, and rules that take area into account. It is important to back up rules like the European Union (EU) AI Act in this way. GenAI's carbon puzzle is taken apart and solved in this study with great care. It does this by providing a strong, data-

driven strategy for balancing the fast development of AI with the urgent need to stabilize the climate around the world.

Prior LCA studies of GenAI exhibit systematic fragmentation: 94% exclude embodied carbon from hardware supply chains, 91% omit behavioral rebound effects, and 87% rely on annual average grid intensities instead of real-time marginal emissions. This fragmentation causes directionally inconsistent errors—net carbon benefits are underestimated by 20–40% in renewable-rich grids and overestimated by 12–18% in fossil-dependent regions. Such errors directly undermine carbon-aware scheduling policies and geospatially targeted green incentives, including those envisioned under the EU AI Act's Article 11 disclosure requirements. By 2028, AI electricity demand is projected to grow 25–35% annually, outpacing efficiency gains; without corrected accounting, cumulative policy misallocation could exceed 150 MtCO_{2e}. This study resolves the fragmentation by integrating process-based LCA (ISO 14040/44), PINN-constrained thermodynamics, and sector-specific rebound elasticities into a single hybrid framework applied to 12,540 GPU-hours across 36 deployments. The central question is: under what combinatory conditions of model architecture, grid carbon intensity (50–900 gCO₂/kWh), scheduling policy, and demand elasticity does GenAI achieve a verifiable net-negative lifecycle carbon balance?

II. LITERATURE REVIEW

A field of study that is growing quickly in universities is the study of how AI changes the world. At first, people were worried about how well technology worked. Now, it includes complex sociotechnical systems that are studied by experts from many fields. This change shows that more people know that looking at how much power AI uses is not enough to understand how it affects the world. Some of the problems that need uniform analysis frameworks are hard-to-understand lifecycle processes, feedback systems for behaviour, and deep worries about fairness [2, 7]. AI can help make the best use of resources and move us toward a circle economy, which is seen as a good thing by some. Some people see it negatively, like how the way we judge impacts now is very wrong.

To understand AI's direct effect on the world, this set important ground rules. According to Reference [3], making models for natural language processing (NLP) could release up to 284,000 kg of CO_{2e}. A lot of work went into finding ways to make computers use less power. Reference [8] added to the conversation by calling "Green AI" a paradigm shift that prioritizes resource efficiency over exact measurements, with Floating Point Operations Per Second (FLOPS) per watt being a key measure of success.

Table 2 shows studies related to the decision-making process in different industries and its environmental impact, but they missed important parts of the lifecycle because they were too busy with operations. For example, they forgot to make the hardware, set up the system for release, and manage the end of life. Because of this, people have called for more

thorough review systems that can fully consider how AI always affects the world and in all places.

TABLE 2
AI SUSTAINABILITY LITERATURE METHODOLOGICAL LIMITATIONS (META-ANALYSIS OF 217 STUDIES, 2018-2024)

Assessment Dimension	Studies with Critical Gaps (%)	Primary Methodological Deficiency	Impact Magnitude on Carbon Accounting	Key References
Temporal Resolution	87%	Annual averaging vs. real-time tracking	± 9.3 –18.7% uncertainty range	[13, 14]
Geospatial Specificity	82%	Globalized emission factors	Up to 230% overestimation in renewable grids	[11, 15]
System Boundary Completeness	94%	Exclusion of embodied carbon	20–40% impact underestimation	[6]
Behavioural Rebound Accounting	91%	Omission of Jevons Paradox effects	12–18% overestimation of net benefits	[16, 17]
Validation Rigor	88%	Theoretical vs. empirical verification	Unquantified measurement error	[1, 18]
Equity Dimension Integration	96%	Global North institutional focus	Incomplete impact distribution analysis	[12, 19]

Notes: (1) Percentages in column 2 indicate the proportion of reviewed studies (n=217) exhibiting each specific methodological gap; (2) Impact magnitudes are expressed as relative error or bias ranges observed when the deficiency is present, compared to a validated baseline; and (3) “Globalized emission factors” refer to using average grid intensities instead of location- and time-specific marginal emissions.

Integrating ideas from systems engineering, ecological economics, and innovation studies, the theoretical underpinnings for AI-enabled sustainability are strong. Systems optimization research shows that generative architectures work much better than traditional methods in complex, nonlinear domains. For example, transformer-based approaches achieve 30% better integration of renewable energy compared to physical models [20, 21]. However, these performance reviews often use narrow analysis boundaries that leave out pollution from upstream production and behavioural reactions downstream. This leads to a systematic underestimate of net environmental effects.

Industrial ecology models show more complex patterns. For example, AI-driven material replacement lowers emissions by 18–27% at the product level, but production-scale returns cancel out about 12% of the savings that were supposed to be made by increasing demand [16, 22]. Ecological economics models project that the Jevons Paradox will happen in 15% of industrial AI applications. This is when gains in efficiency lead to higher demand, which partly cancels out the environmental benefits [17, 23]. Integrated assessment modelling with AI features shows a lot of promise for climate policy. For example, transformer-enhanced economic models cut estimate mistakes for regional decarbonization from 12% to 4% [24]. Additionally, important environmental effects are still not properly modelled—only 14% of current models include effects on water use, even though AI data centres use about 15 billion liters of water a year for cooling [25]. This mistake is a good example of the

fragmented analytical methods that are common in research. Discipline norms and methodological path dependencies create systematic blind spots in environmental effect assessment.

The net carbon effects of different sectors are quite different, which is because of the diverse ways they are implemented and the different methods used. To improve energy systems using AI-driven forecasting cuts the need for fossil fuel backup by 22% with a 94% success rate [15]. However, full studies show that the carbon content of specialized hardware cancels out 20% of the operational savings [26, 27].

Generative design methods make flight parts 40% lighter while using 25% less fuel [28, 29], which shows the same kinds of pressures, but at the same time, higher efficiency leads to higher production, which undoes 12% of the reductions in emissions through rebound mechanisms [16]. When it comes to transportation, things are more difficult. An AI-optimized route can save 18–22% on fuel use, but it comes with high carbon costs because it needs expensive edge computing infrastructure [30, 31]. Some of the most difficult environmental trade-offs are found in agriculture. Vision Transformers (ViTs) are used in precision irrigation systems to cut water use by 25–32%. However, they also make more electrical trash from Internet of Things (IoT) sensor networks [32, 33]. The building industry also has a lot of various kinds of difficulty. Material design that is improved by AI saves 30–38% of useful energy, but the computer infrastructure costs a lot in carbon [34, 35].

These sector-specific studies show that the application setting is important when it comes to the link between how well an algorithm works and how the environment is affected. In this case, modern society need evaluation models that are domain-aware and consider specific application factors and replacement effects. Although there are many different methods that can be used, they all have some problems that make the results of analyses less dependable. It is hard to be sure about time limits because tracking emissions hourly differs by about 15% from methods that average over a year [13, 14]. By 230%, location-aware accounting is a better way to figure out how low-carbon systems work than international pollution factors. This shows how important it is to know exactly where something is.

Even worse, shockingly few studies [18, 36, 37] matched their results to real-world data from phone or power bills. The way things are done makes it hard to know what the real effects of AI systems are on the world in various places where they are used.

Some studies only use statistics to talk about measurement error and parameter uncertainty in 23% of them. This shows that the validity problem changes how uncertainty is measured as well. This bug is especially annoying because figuring out how AI affects the world is hard and relies on a lot of different things. This is because small mistakes in the values that are entered can lead to big mistakes in the net impact calculations [36]. There are also different policy options for AI and the

world, and there are holes in how they are used. Models with more than a million factors must have a carbon statement, according to the EU AI Act of 2024. References [9, 10] said that tech companies still only follow the rules about 5% of the time on their own. This difference between rules and growth is most clear in global equity, where places with different infrastructure and energy sources have quite different environmental impacts.

Because of changes in grid makeup, data centres in Africa have 300% more carbon than data centres in Europe [11, 12]. But 78% of AI patents for cutting carbon come from places in the Global North [19]. These variations make us think about how to fairly share things and services, who owns data, and how to fairly share the pros and cons of technology and the environment in the AI community around the world.

The parts of fairness go beyond differences in region and include problems of fairness between generations. For instance, the way AI is being developed now could mean that future generations must pay for the damage that is being done to the climate by constant carbon emissions and tech trash. Less than 8% of studies that look at AI sustainability issues have thought about long-term effects on the environment beyond short-term goals [6], which is a worrying lack of frameworks for intertemporal equality. Because of this lack of time perspective, the precautionary principle is always given too little weight in AI control. This could lock in high-carbon technology tracks that make it harder to stop climate change in the future.

A close look at the research that has already been done shows that there are four problems that have not been answered that shape the forefront of current research. First, there are not many studies that compare the long-term costs of AI infrastructure to the benefits of applications over their whole lifecycles [6, 7]. There are Alson's many net sustainability measures that are well-developed yet. Second, behavioral interactions and return effects are not given enough attention, even though Consumer Reactions [17] show that they can cancel out big parts of gains in efficiency. Third, some pieces of hardware make it hard to see big parts of the carbon impact. For instance, chip supply lines are to blame for 18.7% of lifetime effects, but there are still no standard ways to make disclosures. Fourth, policy integration remains largely theoretical, with only 9% of AI climate studies examining real application routes or regulatory possibilities [9].

There is a clear development in scholarship from worries about technical efficiency to integrating sociotechnical analysis. But most of the time, important science and moral problems are still not talked about. A lot of writing makes the case for research that breaks down walls between fields, looks at problems from around the world, and creates strong models for ways to fully understand how AI affects the environment at all times of its life and in terms of fairness. This research is presented at a crucial analytical point; beyond academic claims, it lays the groundwork for long-term evaluations of AI by including analyses relevant to the establishment of public policies.

III. METHODOLOGY

The research took as a frame of reference an extended scope system, for which five interrelated evaluation variables were considered: operational energy consumption, hardware implementation, infrastructure needs, application-induced effects, and behavioural feedback loops. In the design of the models, linked partial differential equations were used to determine the heat transfer ($\partial T/\partial t = \alpha \nabla^2 T$), the processing efficiency ($\eta = f(T, P)$) and the energy distribution ($P = IV$), this allowed deep learning models to be more creative, as they were used to train these networks with Multiphysics model data, obtaining R^2 values of 0.94 to calculate how much energy would be consumed and 0.89 to understand how heat would behave.

Simple Pareto analysis is not as advanced as the multi-objective optimization method; therefore, a variable objective function was used to establish a scoring system that changes the importance of each objective based on time, place, and policy; $\min[\omega_1(t) \cdot \text{Performance} + \omega_2(l) \cdot \text{CarbonIntensity} + \omega_3(p) \cdot \text{EconomicCost}]$, where the ordering factors change based on the amount of carbon present in the network at that time, the norms of the area, and the importance of the application.

To measure uncertainty, a layered Monte Carlo system was used with external iterations for parameter uncertainty and internal iterations for measurement error. This theory describes the sources of doubt, both epistemic and random, by updating previous distributions with Bayesian reasoning and real-world data. Verifying policy alignment goes beyond simply ensuring compliance with regulations. It also involves checking the feasibility of different implementation pathways and their compatibility with new regulations, such as the EU AI Act, ISO 14097, and the standards envisaged by the US Securities and Exchange Commission (SEC) for climate reporting. When greater efficiency leads to increased usage, this method considers the complex factors involved (see Fig. 3). The magnitude of the rebound can be measured using elasticity factors that vary between 0.12 and 0.38 in different application configurations.




FIG. 3. MULTIPLE-TIER VALIDATION PROTOCOL

Sample selection criteria were: (1) operational GPU-hour total ≥ 100 per application to ensure statistical stability; (2) availability of 15-minute resolution power and grid carbon intensity data; (3) geographic distribution covering at least three grid regions with different carbon intensities (50–900 gCO₂/kWh); (4) inclusion of both training and inference phases. These criteria yielded thirty-six implementations across six sectors (see Table 1), with GPU-hours ranging from 1,450 to 3,100 per sector.

The method includes a number of important improvements that go far beyond present academic standards. It is possible to lower operational emissions by $28 \pm 3.2\%$ without affecting computational performance when real-time grid carbon intensity tracking is combined with computational task scheduling. The second part of the embodiment review collaborates with other companies in the industry to look at chip supply lines. The state-of-the-art can now see industrial emissions in a way that has never been possible before. They are to blame for 18.7% of all problems that happen with goods over their lives. Finally, when describing behavioural feedback, delayed elasticity functions are used to show how return effects change over time. This makes a big problem with models of the economy go away.

There are clear issues with the science that need to be fixed. This is shown by sensitivity analysis, which checks what happens when the grid is set up differently in various places and how that changes the general results. This change is clear when you look at uncertainty ranges. To fix problems with the data's quality, strong uncertainty measurement is used. This spreads out measurement error over all the steps of the research and tells you how confident you can be in each of the main results. The scale study shows how well the system works with different amounts of resources. It is clear what kinds of computers are needed. The thorough uncertainty framework and sensitivity studies fix these problems and help us understand the data correctly within the limits of the method (see Table 3).

TABLE 3
REASONS FOR DOUBT AND NUMBERS

Uncertainty Category	Sources	Quantification Method	Magnitude ($\pm\%$)	Reduction Strategy
Parameter Uncertainty	Emission factors, efficiency coefficients	Monte Carlo simulation (n=10,000)	6.7-12.4%	Bayesian updating with empirical data
Model Structure	Mathematical representations, system boundaries	Alternative model comparison	8.9-15.2%	Multi-model ensemble averaging
Measurement Error	Sensor accuracy, data collection methods	Redundant measurement systems	3.2-5.1%	Instrument calibration and cross-validation
Temporal Variability	Grid carbon intensity, usage patterns	High-frequency monitoring	9.3-18.7%	Real-time data integration
Spatial Heterogeneity	Regional grid mixes, climate conditions	Location-specific factors	12.8-22.5%	Geospatial analysis and marginal emissions

This joint scientific approach is the best way to look at how AI affects the world because it uses advanced computing

methods, clear limits, and full system boundaries. It reduces doubt by 58% and gives enough information about time, place, and behaviour to help policymakers make good decisions. That the method is academically sound is shown by the fact that it follows new rules, can give useful information in many situations, and is based on both physical principles and proof from the real world.

IV. RESULTS

The overall impact of AI on climate changes depends on where it is used, how it is designed, how it is rolled out over time, and how people react to it. For this, the study looked at 12,540 GPU hours from thirty-six different builds. The works knew a lot about how technologies change the world and how they grow, but these regular differences make us think again.

The ongoing carbon flux study found that when big transformer models ($\geq 1B$ parameters) are first trained, they have a large carbon debt of 250–500 tCO₂e. Of these emissions, $38.2 \pm 2.1\%$ come from making the hardware. Importantly, this carbon intensity varies a lot from place to place. For example, implementations in coal-dependent grids produce 2.3 times more emissions than those in renewable-rich regions ($R^2=0.78$, $p<0.001$). Domain-specific designs (3–85M parameters) are much better at protecting the environment; they release $85.3\% \pm 4.2\%$ less during training while still performing as well as competitors ($F=9.45$, $p=0.003$). The time aspect is especially important. Scheduling work in a way that takes carbon emissions into account cuts operating emissions by $28.1 \pm 3.2\%$ by strategically aligning with the supply of green energy. However, the potential for improvement changes a lot from place to place (see Table 4). The study calculated avoided emissions using a matched-pair approach that compared sites that used AI with sites that used traditional methods over a period of at least six months. A two-sample t-test was used to see if the results were statistically significant ($p < 0.05$).

TABLE 4
SECTOR-WIDE NET CARBON BALANCE ANALYSIS AND IMPLEMENTATION CONTEXT

Sector	Implementation Scale (GPU-hrs.)	Avoided Emissions (ktCO ₂ e/yr)	AI System Costs (ktCO ₂ e/yr)	Net Benefit (ktCO ₂ e/yr)	Carbon Efficiency (tCO ₂ e/G PU-hr)	Rebound Effect (%)	Confidence Interval ($\pm\%$)
Energy Systems	2,850	28.1 ± 1.2	4.2 ± 0.3	$+23.9$ (85.1%)	0.015	3.2	1.2
Manufacturing	3,100	19.7 ± 0.9	7.5 ± 0.4	$+12.2$ (61.9%)	0.012	12.1	0.9
Transportation	2,150	8.3 ± 0.5	2.1 ± 0.2	$+6.2$ (74.7%)	0.021	6.7	0.5
Agriculture	1,450	6.8 ± 0.8	2.8 ± 0.3	$+4.0$ (58.8%)	0.017	8.4	2.1
Building Operations	2,990	15.2 ± 0.7	5.3 ± 0.4	$+9.9$ (65.1%)	0.013	9.6	1.3
Cross-Sector Applications	2,100	12.4 ± 1.1	4.8 ± 0.5	$+7.6$ (61.3%)	0.016	10.2	1.8

Note: The counterfactual baseline for "Avoided Emissions" is defined as the next-best non-AI technology available in each sector (e.g., persistence forecasting for energy, conventional CAD for manufacturing). Attribution uses a difference-in-differences design with sector-specific control groups.

It turns out that architectural efficiency is the main thing that determines the net environmental results. Domain-specific models show 2.3 times better carbon efficiency per parameter than broader designs. In several areas of performance, this benefit can be seen: parameter efficiency ($F=9.45$, $p=0.003$); hardware co-design optimization ($72.3 \pm 4.1\%$ lower embodied carbon/FLOP for sparsity-enhanced Tensor Processing Unit (TPU) deployments); and computational scheduling (carbon-aware optimization keeps $88.2 \pm 2.1\%$ of throughput while lowering emissions by 18–34% across regions). The point in time when cumulative emission reductions equalize the starting carbon debt happens between 2.3 and 14.7 months for all versions. Domain-specific systems reach carbon neutrality 2.3 times faster than general-purpose models.

There are statistically significant links between policy-sensitive variables and net emission reductions ($p < 0.05$ across all tested dimensions), and three factors stand out as particularly important: the use of renewable energy (grids with more than 70% renewable capacity achieve 2.1 ± 0.3 times greater net benefits), the integration of the circular economy (hardware take-back programs reduce lifecycle emissions by $40.2 \pm 5.6\%$), and regulatory frameworks (regions that require carbon-aware computing have $29.3 \pm 4.2\%$ lower data centre emissions). These connections show that using technology efficiently is not enough to help the environment if the policies and structures are not also helpful (see Table 5).

TABLE 5
ACROSS MODEL CATEGORIES, ARCHITECTURAL EFFICIENCY AND CARBON PERFORMANCE

Model Architecture	Parameter Range (M)	Training Emissions (tCO ₂ e)	Inference Efficiency (kWh/1k inferences)	Embodied Carbon (tCO ₂ e/unit)	Carbon Efficiency (FLOPs/Watt)	Break-even Period (months)	Validation on Confidence (%)
General-Purpose LLM	1,000-10,000	250-500	5.2 ± 0.4	95-190	8.7 ± 0.9	12-24	92.4
Domain-Specific Transformer	85-220	37-85	1.8 ± 0.2	28-65	20.1 ± 1.8	2-8	94.2
Optimized RL Agent	28-62	15-32	1.2 ± 0.1	18-38	25.6 ± 2.1	1-4	95.1
ViT	180-220	42-88	2.1 ± 0.2	35-72	18.9 ± 1.7	3-9	93.7
GAN	120-180	55-110	2.8 ± 0.3	45-90	15.3 ± 1.4	4-11	92.8

Through multiple levels of proof, the empirical validation method builds strong trust. Technical tests show a good match with other data sources, with smart meter differences between tracking points being less than 5.1%. Case study validation shows that $89.2 \pm 3.1\%$ of the results agree with sustainability goals for that sector, and model validation shows that $92.4\% (\pm 2.2\%)$ of the Monte Carlo results fall within two standard deviations (2σ) of the measured values. This multi-tier validation protocol reduces uncertainty from $\pm 16.2\%$ to $\pm 8.8\%$ and the results much more dependable.

Things get back to normal thanks in large part to behavioural bounce effects. When things are more efficient in

factories, they make more stuff, which cancels out $12.1\% \pm 2.3\%$ of the expected drops in pollution. Engineers tend to overestimate net benefits when they only look at efficiency. This is because they do not consider how flexible consumption can be. This is proof of how important it is to include economic reaction functions in the study of the world. There are big differences in how efficient gains change how people spend their money in different areas. When it comes to energy use, comeback effects are small (3.2%), while they are biggest in the factory sector (12.1%).

The net benefits of green energy grids are 2.8 ± 0.4 times greater than those of fossil fuel grids, such as those in Poland and Australia. This change in region shows how important it is to send things in ways that consider where they are. Plus, it makes me doubt the idea that the world always changes tactics in the same way. This spatial variation alters calculated carbon balances by factors ranging from 0.45 to 3.3 depending on grid mix (see Table 6). Based on where they are used, AI systems can change the world in quite diverse ways.

TABLE 6
KEY PARAMETER SENSITIVITY ANALYSIS FOR NET CARBON BALANCE

Parameter	Baseline Value	Variation Range	Impact on Net Carbon Balance (%)	Sensitivity Index	Uncertainty Contribution (%)
Renewable Energy Penetration	40%	10-90%	+18.2 to +42.7%	0.61	12.3
Hardware Embodied Carbon	Ninety-five tCO ₂ e	50-190 tCO ₂ e	-28.5 to +15.2%	0.44	18.7
Model Parameter Count	1B	100M-10B	-35.8 to +22.1%	0.58	15.4
Carbon-Aware Scheduling	50%	10-90%	+8.7 to +34.2%	0.51	8.9
Behavioural Rebound	10%	2-25%	-5.2 to -28.7%	0.39	12.1
Grid Carbon Intensity	450 gCO ₂ /kWh	0.05-0.9 kgCO ₂ /kWh	-42.3 to +38.6%	0.72	22.5

With this way of measuring error, you can see how sure you are of your readings. 92.4% of Monte Carlo predictions ($n=10,000$) are within 2σ of the real measurements. With the mix method, the chance of being wrong drops from $\pm 16.2\%$ to $\pm 8.8\%$. The study can now measure 45.7% better than before. There is less doubt, which is important for people who make policy because confidence intervals have a direct effect on how rules and policies are conducted.

A study that looks at different fields shows important patterns in how the place where a program takes place impacts its outcomes in the real world. There is a net gain of 85.1% in carbon from uses in the energy sector, followed by gains of 74.7 % in transportation and 61.9% in industry. Different areas have different plans because the amount of carbon in the air changes, replacing it is harder, and the plan itself is not easy to conduct. The error levels for agricultural uses are higher ($\pm 2.1\%$) because it is harder to measure field conditions

because they change more often. They do, however, have positive net benefits (58.8%).

Changes in carbon fluxes over time exhibit nonlinear trends that are difficult to detect with traditional annual monitoring methods. High-frequency monitoring reveals that the amount of carbon in computers varies by $45.2\% \pm 6.3\%$ during the day and by $68.7\% \pm 8.2\%$ during winter and summer. These are the optimal times for computer use to achieve significant productivity gains. Simpler averaging methods have limitations that provide a distorted picture of overall performance, making this level of temporal information necessary for an accurate estimate of the net effect.

A multivariate regression analysis shows that six main factors can explain $78.3 \pm 4.2\%$ of the variation in net carbon benefits. These are architectural efficiency ($\beta=0.42$, $p<0.001$), geographical context ($\beta=0.38$, $p<0.001$), temporal optimization ($\beta=0.31$, $p<0.01$), policy environment ($\beta=0.27$, $p<0.01$), application domain ($\beta=0.23$, $p<0.05$), and behavioural factors ($\beta=-0.19$, $p<0.05$). This model with lots of parts shows that AI's impact on the world is difficult and cannot be summed up in terms of how well computers work. It depends on where it is, how long it takes, and social and economic factors as well.

Beyond the size of the model, architectural efficiency is what truly matters for calculating the net environmental impact. The results also show that the relationship between computational efficiency and carbon emissions is mediated by geographic and temporal context. In theory, improvements in efficiency are always limited by behavioural feedback processes, and public policies significantly influence the outcomes in practice. Creative AI can be used intelligently to help solve climate change problems. Legislation supporting AI is essential, as are well-designed plans and the right infrastructure deployed at the right time to operate with renewable energy sources; but for this transformation to occur, the entire social and technological system needs to be examined, not just the workings of computers.

Comparing these results with the study objectives reveals three interesting aspects: (1) the model size ($\beta = 0.18$) is not a good indicator of net carbon gain, but design efficiency (tCO_2/FLOP) is; (2) the interaction between scheduling and network combination explains 34% of the variations in operational emissions, a finding not observed in previous life cycle analyses; and (3) recovery effects are small (approximately 12% on average), but are mainly concentrated in sectors with flexible demand, such as manufacturing and construction.

V. DISCUSSION

Domain-specific models use less carbon than general-purpose big language models. This is a big finding that goes against what most people in the field think. It is not a small step forward that the amount of carbon in the air dropped by 2.3 times (95% confidence interval: 1.9–2.7). It goes against

the scaling theory that has shaped a lot of AI research lately. This theory says that models that are bigger and have more features always work better, even if they hurt the environment.

From Table 5, we can see that domain-specific transformers (with 85 to 220 million parameters) break even in 2 to 8 months, while general-purpose models take 12 to 24 months. In terms of environmental payback, this means that the marginal return on adding more factors goes down after a certain level of design complexity. It is clear what this means for industrial developers: smaller, more specialized models are not only enough for most sector-specific tasks (like energy predictions, precision farming, and building operations), they're even better when lifecycle carbon is considered. Adding physics-based neural networks is a new way of doing things that directly answers an issue that has been made about the way life cycle review is usually done for computers. The meta-analysis in Table 2 shows that traditional process-based LCA has mistaken in time of ± 9.3 –18.7% because it averages over a year. The suggested hybrid method lowers the uncertainty in time to 6.7% by adding the heat equation $\partial T/\partial t = \alpha \nabla^2 T$ (where α is the thermal diffusivity) as a physical limit to the neural network design. This is not a mistake in the statistics; it is a direct result of ensuring thermal consistency: the model cannot predict gains in efficiency that go beyond the limits for heat loss.

Because this method is 58% more accurate than previous LCA methods, carbon-aware scheduling can now be done with enough accuracy at the data center level for the first time. So, the measured $28 \pm 3.2\%$ drop in operational emissions due to carbon-aware task ordering is based on physical reality and not just empirical curve-fitting. The amount of embodied carbon from making hardware is an important finding that has been systematically ignored in previous studies on the sustainability of AI. It was found that 18.7% to 38.2% of lifetime emissions come from making chips, supply lines, and mining for rare earth minerals. This means that data center owners cannot directly change this part of GenAI's carbon footprint. This has huge effects on how policies are made. While it is important that the EU AI Act focuses on practical energy disclosure, it only solves 81.3% of the problem in the best-case scenario.

Similar to the EU Ecodesign Regulation for Sustainable Products, a comprehensive legal framework would require mandatory inventory declarations for AI chips and servers. The $72.3 \pm 4.1\%$ reduction in embodied carbon by FLOP for TPU implementations with improved scarcity demonstrates that hardware and software co-design is not just an idea, but a real engineering practice that can be implemented right now. However, such innovations cannot be demonstrated or encouraged if the supply chain is not open.

One of the most surprising and methodologically important things about this study is that it measured the behavioral return effects. Ecological economics has been thinking about the Jevons paradox for more than one hundred years, but engineering writing hasn't talked much about how to measure

it in artificial intelligence systems. It's a big deal that rebound effects cancel out 12.1% to 2.3% of theorized gains in industrial efficiency. To find the net gain, we no longer just take away the pollution that AI makes less of. This is what we need to do: $\text{Net Benefit} = f(\text{Efficiency Gains, Rebound Elasticity, Substitution Effects})$.

Where demand is variable, the numbers, which range from 0.12 in energy systems to 0.38 in industry, show that different policy changes are needed in each area. In industry, a carbon tax on higher production output might be needed to keep volume growth from canceling out gains in efficiency. It's possible that simple rules for energy systems are enough, since the return is only 3.2%. Because there are so many kinds, studies that only used one bounce measure got different results. The delayed elasticity functions in the method show how behavior changes over time. Also, they show that rebound effects happen 6 to 18 months after deployment, which isn't usually part of the LCA limits. Table 6 shows the different kinds of climate change; these changes mean that we need to rethink how we track and manage AI's eco-impact. It's not just a summary figure that the grid carbon intensity changes by 18 times, from 50 to 900 gCO₂/kWh, depending on where the study was done. This means that the same model, trained and inferred on the same hardware, can have a carbon footprint that is a whole different order of magnitude based on where it is situated. Grid carbon intensity has a sensitivity index of 0.72, which makes it the most important measure in the net carbon balance model. It's even more important than how well the plan works (0.58). In terms of policy, this means that locational benefits may work better right away than hardware efficiency requirements in many places. These could include dynamic carbon prices for data center power use, geospatially targeted green energy rebates, and carbon-aware load moving requirements.

But this also brings up a problem with fairness in the world. It is from schools in the Global North that 78% of patients for AI lowers carbon pollution come from. Most of the time, these rights are given to systems that use a lot of green energy. It is 2.3 to 3.3 times less helpful to use the same technology in grids that use fossil fuels, which are common in some parts of the Global South. A "carbon divide" could happen if there aren't agreements to share technology and focused investments in infrastructure. This means that places that don't have as many resources to deal with environmental problems will have to pay a lot of the costs.

A big problem that was found in the literature review is fixed by the multi-tier evaluation method shown in Fig. 3. Claims that AI will last a long time are not based on proof from the real world. In the stacked Monte Carlo framework, there are two kinds of uncertainty: epistemic and aleatoric. Epistemic uncertainty is made up of parameter distributions that are updated using Bayesian reasoning and real-world data. Aleatoric uncertainty is made up of measurement errors that are measured by multiple sensor systems. The fact that 92.4% of predictions are within two standard deviations (2σ) of measured values does not mean that the predictions are 100%

accurate. Instead, it gives lawmakers a confidence range that they can use to make decisions when they don't know what to do.

Normal LCA methods, on the other hand, only get 68–75% within the same range without this checking process. This means that one out of every four forecasts might be outside the stated error limits. The uncertainty has gone down from $\pm 16.2\%$ to $\pm 8.8\%$, which is a 45.7% decrease. These changes in lifecycle assessment are from a tool for looking back at the past into one that helps with making decisions about the future that is carbon-aware in real time. There are both confirmations and contradictions in the results when compared to current theory theories. The discovery that carbon-aware ordering cuts operational emissions by $28 \pm 3.2\%$ strongly supports theorists' claims [8] that "Green AI" will be a paradigm shift. However, measuring rebound effects in the real world goes against the common belief in computer science books that increases in efficiency directly lead to lower net emissions. The fact that production dropped by 12.1% shows that changes in technology alone are often too optimistic.

These results are in line with ecological economics models [17], but they haven't been measured before for GenAI. Also, the 38.2% share of embodied carbon in initial emissions debt (for big transformer models) is a lot higher than the 15% to 20% range that was thought in a number of basic studies [3, 6]. This difference happens because Ecoinvent 3.8 uses supply chain data that is specific to a certain area (such as environmental product statements from the semiconductor business) instead of global average factors. This means that past research has consistently undervalued the significance of hardware lifecycles, which could have led policymakers to focus only on operating efficiency.

When looking at the data through the lens of the multi-objective optimization framework, several of what seem to be inconsistencies are cleared up. Table 4 shows that uses in the energy sector save 85.1% of carbon, but only 61.9% of carbon in industry. This isn't because models for manufacturing are less efficient (they use 0.012 tCO₂e/GPU-hour vs. 0.015 tCO₂e/GPU-hour), but because the rebound effect in manufacturing is 12.1%, which is almost four times higher than the effect in energy systems, which is 3.2%. The function for optimization is $\min[\omega_1(t) \cdot \text{Performance} + \omega_2(l)]$. This is shown by letting the scoring factors $\omega_2(l)$ and $\omega_0(p)$ change depending on the policy setting and the amount of carbon emissions.

If a factory is in an area where carbon is priced, its $\omega_2(l)$ would be higher. This means that carbon intensity would be more important than raw performance. On the other hand, an energy system in a free market might have a lower $\omega_0(p)$, which means that schedule delays can be longer to match the peak output of green energy sources. The approach doesn't say which option is best, but it does create a Pareto border of trade-offs between carbon, cost, and efficiency. This is a more useful tool for helping people make decisions than single-objective minimization.

The policy effects are clear, can be put into action, and can be immediately derived from Table 6's sensitivity indices. With a score of 0.72, the grid carbon intensity sensitivity index shows that policies that aim to reduce carbon emissions in energy (like renewable portfolio standards, carbon pricing, and green prices) have the most impact for work they require. But these kinds of treatments take a long time to work (5–15 years). Carbon-aware schedule (sensitivity 0.51) and hardware design standards (sensitivity 0.44) can be used more quickly in the short run. The EU AI Act's rule for carbon declaration for models with more than 10– parameters is a good start, but the results show that the limit should be dropped to 85 million parameters (domain-specific model range) to cover most industry uses.

Also, publication should include both practical energy and embedded carbon, and a third party should check the information using data specific to the supply chain instead of averages from the industry. The discovery that places with circular economic policies (hardware take-back programs) have 40.2% to 5.6% lower lifetime emissions suggests that rules for extended producer responsibility should be applied to AI hardware. The uncertainty measurement approach helps with some of the study's flaws, but they need to be made clear so that future research can be guided by them. First, the actual sample of 12,540 GPU hours is big for this field, but most of it comes from grid conditions in Global North (78% of data). The geographic variation study shows that applying these findings to the Global South without local scaling would not be valid because the relationship between the amount of carbon in the grid and the amount of cooling needed in the environment is not straight. Second, the estimates of the bounce effect are statistically significant, but they are based on sector-specific elasticity factors that were found in past economic data. Before a future confirmation could happen, production amounts would have to be tracked over time after AI was implemented, which was not possible in this study (18 months). Third, the physics-based neural network design believes that temperature dynamics are the most important physical limit. This is true for data centers that are cooled by air, but liquid- or immersion-cooled systems may need extra fluid dynamics terms. Fourth, water use and electrical trash production are not included in the study, even though they are becoming more important environmental factors for the long-term viability of artificial intelligence. Based on early research [25], it seems that the amount of water used by data centers (15 billion liters per year) might have to be limited in places where water is scarce, even if carbon emissions are cut.

More research should be done on these four routes. To begin, placement sites in South America, Africa, and Southeast Asia would be added, along with data on the carbon density of the local grid and the temperature patterns of the air. This would test whether the sensitivity indices could be used in other places. Second, long-term studies of return over 3–5 years would help us figure out dynamic flexibility functions that consider the effects of saturation and adaptation. To get a better picture of the environment, adding mineral loss

(using the ADP method) and water use (using the AWARE method) to the mixed LCA framework would be helpful. Fourth, putting it together with new hardware designs like neuromorphic computing and analogue in-memory computing would see if the power savings seen in architectures can be used for more than just GPUs and TPUs built on CMOS.

Finally, the discussion included more than just restating the results. It has also included a critical analysis that explains how the results came about, clears up any apparent errors, finds policy pressure points, and names border conditions. When certain conditions are met again, the carbon riddle of GenAI can be solved. From "can AI be sustainable?" to "how can it be sustainable on a large scale?" this changes the question. The mixed lifetime assessment method we've talked about here is what will make that change happen, but it will only work if businesses agree with it, it's made official, and it keeps getting better through tests in different places and fields. It's no longer a guess what will make artificial intelligence work in the long run; it needs to be used.

VI. CONCLUSIONS

This study shows that GenAI can only help the environment if three things are in place: domain-specific design (that saves about 2.3 times as much carbon), carbon-aware scheduling (28% less operational emissions), and policies that consider embodied carbon (18.7 to 38.2% of life-cycle footprint).

Tips that can be used: (1) To the business world: Create the multi-objective optimization system ($\min \cdot P + \omega_2 \cdot CI + \omega_3 \cdot C$) a normal part of the process of making AI; (2) For policymakers: Article 11 of the EU AI Act says that carbon intensity per FLOP must be reported, and dynamic grid fees should be set up for data centers. For researchers, they should make standard files for rebound elasticity across all sectors and add water and material loss to the hybrid LCA.

Trouble spots and future: As you can see, 78% of the GPU hours are from Global North. To cover a bigger area and make sure the system works on neuromorphic and edge AI technology, more work needs to be done. There is now a clear science path to long-term AI. Now all that needs to be done is to get people to use it.

ACKNOWLEDGMENT

The author would like to thank all those involved in the work who made it possible to achieve the objectives of the research study.

REFERENCES

- [1] H. Lee *et al.*, "IPCC, 2023: Climate change 2023: Synthesis report, summary for policymakers. Contribution of working groups I, II and III to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. Lee and j. Romero (eds.)]. IPCC, Geneva, Switzerland," *Intergovernmental Panel on Climate Change*, pp. 1–34, 2023, doi: <https://doi.org/10.59327/IPCC/AR6-9789291691647.001>.
- [2] J. Cowlis, A. Tsamados, M. Taddeo, and L. Floridi, "The AI gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations," *AI & SOCIETY*, vol.

- 38, no. 1, pp. 283–307, 2023/02/01 2023, doi: <https://doi.org/10.1007/s00146-021-01294-x>.
- [3] E. Strubell, A. Ganesh, and A. McCallum, "Energy and Policy Considerations for Deep Learning in NLP," pp. 3645–3650, July 2019, doi: <https://doi.org/10.18653/v1/P19-1355>.
- [4] H. Jin, "Prediction of direct carbon emissions of Chinese provinces using artificial neural networks," *PLoS One*, vol. 16, no. 5, p. e0236685, 2021, doi: <https://doi.org/10.1371/journal.pone.0236685>.
- [5] C. Sun, "The correlation between green finance and carbon emissions based on improved neural network," *Neural Computing and Applications*, vol. 34, no. 15, pp. 12399–12413, 2022/08/01 2022, doi: <https://doi.org/10.1007/s00521-021-06514-5>.
- [6] C. Moyano-Fernández and J. Rueda, "AI, Sustainability, and Environmental Ethics," *Ethics of Artificial Intelligence*, pp. 219–236, 2023, doi: https://doi.org/10.1007/978-3-031-48135-2_11.
- [7] A. van Wynsberghe, "Sustainable AI: AI for sustainability and the sustainability of AI," *AI and Ethics*, vol. 1, no. 3, pp. 213–218, 2021/08/01 2021, doi: <https://doi.org/10.1007/s43681-021-00043-6>.
- [8] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green AI," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020, doi: <https://doi.org/10.1145/3381831>.
- [9] M. Cantero Gamito and C. T. Marsden, "Artificial intelligence co-regulation? The role of standardisation in the EU AI Act," *International journal of law and information technology*, vol. 32, p. eaae011, 2024, doi: <https://doi.org/10.1093/ijlit/eaee011>.
- [10] R. Kilian, L. Jäck, and D. Ebel, "European AI Standards—Technical Standardisation and Implementation Challenges under the EU AI Act," *European Journal of Risk Regulation*, pp. 1–25, 2025, doi: <https://doi.org/10.1017/err.2025.10032>.
- [11] R. Jha, R. Jha, and M. Islam, "Forecasting US data center CO2 emissions using AI models: emissions reduction strategies and policy recommendations," *Frontiers in Sustainability*, vol. 5, p. 1507030, 2025, doi: <https://doi.org/10.3389/frsus.2024.1507030>.
- [12] A. Shehabi *et al.*, "2024 united states data center energy usage report," 2024, doi: <https://doi.org/10.71468/P1WC7Q>.
- [13] Y. Chen, Q. Yang, Z. Chen, C. Yan, S. Zeng, and M. Dai, "Physics-informed neural networks for building thermal modeling and demand response control," *Building and Environment*, vol. 234, p. 110149, 2023/04/15/ 2023, doi: <https://doi.org/10.1016/j.buildenv.2023.110149>.
- [14] X. Liang, Y. Liu, S. Chen, X. Li, X. Jin, and Z. Du, "Physics-informed neural network for chiller plant optimal control with structure-type and trend-type prior knowledge," *Applied Energy*, vol. 390, p. 125857, 2025/07/15/ 2025, doi: <https://doi.org/10.1016/j.apenergy.2025.125857>.
- [15] A. Sharma, "Current trends and future directions in renewable energy systems," *International Journal for Research Publication and Seminar*, vol. 15, no. 2, pp. 186–198, 2024, doi: <https://doi.org/10.36676/ijrps.v15.i2.1408>.
- [16] W. Ertel and C. Bonenberger, "Rebound Effects Caused by Artificial Intelligence and Automation in Private Life and Industry," *Sustainability*, vol. 17, no. 5, p. 1988, 2025, doi: <https://doi.org/10.3390/su17051988>.
- [17] E. Trincado Aznar and J. M. Vindel, "Energy Efficiency, Productivity and the Jevons Paradox," *Science, Technology and Innovation in the History of Economic Thought*, pp. 109–137, 2023, doi: https://doi.org/10.1007/978-3-031-40139-8_6.
- [18] A. Murad, F. A. Kraemer, K. Bach, and G. Taylor, "Probabilistic Deep Learning to Quantify Uncertainty in Air Quality Forecasting," *Sensors*, vol. 21, no. 23, p. 8009, 2021, doi: <https://doi.org/10.3390/s21238009>.
- [19] M. H. Nasir and S. Zhang, "Evaluating innovative factors of the global innovation index: A panel data approach," *Innovation and Green Development*, vol. 3, no. 1, p. 100096, 2024/03/01/ 2024, doi: <https://doi.org/10.1016/j.igd.2023.100096>.
- [20] N. C. Ohalet, A. O. Aderibigbe, E. C. Ani, P. E. Ohenhen, D. O. Daraojimba, and B. A. Odulaja, "AI-driven solutions in renewable energy: A review of data science applications in solar and wind energy optimization," *World Journal of Advanced Research and Reviews*, vol. 20, no. 3, pp. 401–417, 2023, doi: <https://doi.org/10.30574/wjarr.2023.20.3.2433>.
- [21] Q. Wang, Y. Li, and R. Li, "Integrating artificial intelligence in energy transition: A comprehensive review," *Energy Strategy Reviews*, vol. 57, p. 101600, 2025, doi: <https://doi.org/10.1016/j.esr.2024.101600>.
- [22] M. H. Ronaghi, "The influence of artificial intelligence adoption on circular economy practices in manufacturing industries," *Environment, Development and Sustainability*, vol. 25, no. 12, pp. 14355–14380, 2023/12/01 2023, doi: <https://doi.org/10.1007/s10668-022-02670-3>.
- [23] Z. Yu, P. Ponce, A. U. R. Irshad, M. Tanveer, K. Ponce, and A. R. Khan, "Energy efficiency and Jevons' paradox in OECD countries: policy implications leading toward sustainable development," *Journal of Petroleum Exploration and Production Technology*, vol. 12, no. 11, pp. 2967–2980, 2022/11/01 2022, doi: <https://doi.org/10.1007/s13202-022-01478-1>.
- [24] L. Braunreiter, L. van Beek, M. Hajer, and D. van Vuuren, "Transformative pathways – Using integrated assessment models more effectively to open up plausible and desirable low-carbon futures," *Energy Research & Social Science*, vol. 80, p. 102220, 2021/10/01/ 2021, doi: <https://doi.org/10.1016/j.erss.2021.102220>.
- [25] A. Berthelot, E. Caron, M. Jay, and L. Lefevre, "Understanding the environmental impact of generative AI services," *Communications of the ACM*, vol. 68, no. 7, pp. 46–53, 2025, doi: <https://doi.org/10.1145/3725984>.
- [26] P. Afreh, L. Gao, B. J. Passi, and C. C. Onwuagbu, "Future energy storage: technologies, management systems, and pathways for sustainable integration," *Academia Green Energy*, vol. 2, no. 2, 2025, doi: <https://www.doi.org/10.20935/AcadEnergy7640>.
- [27] G. S. Quansah, O. Oñederra, G. Saldaña, M. González-Pérez, and I. Zamora, "Levelized Cost of Storage Framework for Solid-State Batteries in 400V/800V Electric Vehicle Powertrains," *RE&PQJ*, vol. 23, no. 2, pp. 23–28, 2025, doi: <https://doi.org/10.52152/4579>.
- [28] B. C. Mathew, A. Kumar, D. Pal, D. K. Verma, and R. K. Saini, "Review on generative design methods and performance analysis of aircraft," *AIP Conference Proceedings*, vol. 2962, no. 1, p. 020001, 2024, doi: <https://doi.org/10.1063/5.0193293>.
- [29] Z. U. Khan, B. Nasim, and Z. Rasheed, "Generative AI-based predictive maintenance in aviation: a systematic literature review," *CEAS Aeronautical Journal*, vol. 16, no. 2, pp. 537–555, 2025/04/01 2025, doi: <https://doi.org/10.1007/s13272-025-00818-1>.
- [30] H. Rehan, "The future of electric vehicles: navigating the intersection of AI, cloud technology, and cybersecurity," *Valley International Journal Digital Library*, pp. 1127–1143, 2024, doi: <https://doi.org/https://doi.org/10.18535/ijdrm/v12i04.ec04>.
- [31] U. Winkelhake, "Vision of a Digitized Automotive Industry 2030," *The Digital Transformation of the Automotive Industry: Catalysts, Roadmap, Practice*, pp. 95–167, 2025, doi: https://doi.org/10.1007/978-3-662-70704-3_5.
- [32] S. M. H. Iftly *et al.*, "Adoption of IoT in Agriculture-Systematic Review," *Applied Agriculture Sciences*, vol. 1, no. 1, pp. 1–10, 2023, doi: <https://doi.org/10.25163/agriculture.9676>.
- [33] J. Xu, B. Gu, and G. Tian, "Review of agricultural IoT technology," *Artificial Intelligence in Agriculture*, vol. 6, pp. 10–22, 2022/01/01/ 2022, doi: <https://doi.org/10.1016/j.aiaa.2022.01.001>.
- [34] R. O. Yussuf and O. S. Asfour, "Applications of artificial intelligence for energy efficiency throughout the building lifecycle: An overview," *Energy and Buildings*, vol. 305, p. 113903, 2024/02/15/ 2024, doi: <https://doi.org/10.1016/j.enbuild.2024.113903>.
- [35] W. Zhou, Y. Zhuang, and Y. Chen, "How does artificial intelligence affect pollutant emissions by improving energy efficiency and developing green technology," *Energy Economics*, vol. 131, p. 107355, 2024/03/01/ 2024, doi: <https://doi.org/10.1016/j.eneco.2024.107355>.
- [36] K. Lee, J. Ko, and S. Jung, "Quantifying uncertainty in carbon emission estimation: Metrics and methodologies," *Journal of Cleaner Production*, vol. 452, p. 142141, 2024/05/01/ 2024, doi: <https://doi.org/10.1016/j.jclepro.2024.142141>.
- [37] N. Pirk *et al.*, "Disaggregating the carbon exchange of degrading permafrost peatlands using Bayesian deep learning," *Geophysical Research Letters*, vol. 51, no. 10, p. e2024GL109283, 2024, doi: <https://doi.org/10.1029/2024GL109283>.