



## OPEN ACCESS

### EDITED BY

Brian Fath,  
Towson University, United States

### REVIEWED BY

Narcis Eduard Mitu,  
University of Craiova, Romania  
Ghifari Munawar,  
Bandung Institute of Technology,  
Indonesia

### \*CORRESPONDENCE

Esther O. Esho  
✉ ester4peace@gmail.com  
Andronicus A. Akinyelu  
✉ akinyelua1@ukzn.ac.za

RECEIVED 16 October 2025

REVISED 20 January 2026

ACCEPTED 27 January 2026

PUBLISHED 25 February 2026

### CITATION

Esho EO, Akinyelu AA and  
Dinis MAP (2026) Sustainable generative  
AI and quantum computing: review  
assessment on the environmental  
impact of generative AI and quantum  
technologies.  
*Front. Sustain.* 7:1726832.  
doi: 10.3389/frsus.2026.1726832

### COPYRIGHT

© 2026 Esho, Akinyelu and Dinis. This is  
an open-access article distributed under  
the terms of the [Creative Commons  
Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use,  
distribution or reproduction in other  
forums is permitted, provided the  
original author(s) and the copyright  
owner(s) are credited and that the  
original publication in this journal is  
cited, in accordance with accepted  
academic practice. No use, distribution  
or reproduction is permitted which does  
not comply with these terms.

# Sustainable generative AI and quantum computing: review assessment on the environmental impact of generative AI and quantum technologies

Esther O. Esho<sup>1\*</sup>, Andronicus A. Akinyelu<sup>2\*</sup> and  
Maria Alzira Pimenta Dinis<sup>3,4</sup>

<sup>1</sup>The Australian New Zealand Society for Ecological Economics (ANZSEE), Brisbane, QLD, Australia, <sup>2</sup>School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Pietermaritzburg, South Africa, <sup>3</sup>Fernando Pessoa Research, Innovation and Development Institute (FP-I3ID), University Fernando Pessoa (UFP), Porto, Portugal, <sup>4</sup>Marine and Environmental Sciences Centre (MARE), University of Coimbra, Coimbra, Portugal

The rapid advancement of Generative Artificial Intelligence (GenAI) and Quantum Computing (QC) presents transformative opportunities, yet their high computational requirements raise concerns about their environmental sustainability. This comprehensive review examines the ecological footprint of both technologies, focusing on key metrics like energy consumption, carbon emissions, and resource depletion. Findings from existing studies consistently indicate that the impact of GenAI is mostly driven by the immense energy demands of large-scale model training and inference. Moreover, findings from the review reveal that the footprint of QC largely stems from the energy-intensive cryogenic cooling and rare material requirements of its specialized hardware. This paper benchmarks current approaches to environmental assessment, highlighting the important role of Life Cycle Assessment (LCA) in providing a holistic view of the classification of environmental impacts across the entire supply chain, from manufacturing to disposal. This study proposes a range of domain-specific mitigation strategies, including algorithmic optimizations like pruning and distillation for AI, and cryogenic and material sourcing improvements for quantum systems. This study also proposes a framework for proactive, responsible innovation and identifies some gaps in the literature, such as the lack of standardized metrics and transparent reporting. There is a need to embed eco-conscious principles in the design of future technologies and highlight opportunities where these technologies can be used to handle broader climate challenges. The findings in this study can be used by policymakers, researchers, and industry stakeholders in aligning technological progress with global climate and sustainability goals.

### KEYWORDS

carbon footprint assessment, environmental sustainability, Generative Artificial Intelligence, green computing, quantum computing

## 1 Introduction

Generative Artificial Intelligence (GenAI) and Quantum Computing (QC) are two big technologies that are driving transformation in many industries today. Tools like ChatGPT and image generators have made it easy for anyone to create content, sped up scientific discoveries, and improved productivity in many industries, from medicine to finance (Dauner and Socher, 2025). Likewise, QC, though still at its infant stage, promises to solve problems that are currently impossible, such as designing new drugs, cracking complex codes, and optimizing huge systems (Kakarla et al., 2025; Nikolopoulou, 2025; Vasantrao and Saxena, 2025). These technologies are more than just economic drivers; they have the potential to genuinely improve society, helping us tackle climate change and disease. However, all these amazing benefits come with a significant and often ignored cost to the environment. The energy consumption of modern AI is enormous. For instance, training a single large language model such as GPT-3 has been estimated to consume hundreds of megawatt-hours of electricity, an energy demand comparable to the annual consumption of dozens of average households, highlighting the significant environmental footprint of contemporary AI systems (Hinov, 2025). This large amount of energy directly translates into carbon emissions, especially when it comes from unclean sources like coal. The environmental impact of QC is also serious. Current quantum machines need to be kept extremely cold; the quantum chips inside the machines must reach temperatures that are colder than anything in the world system (Meier and Yamasaki, 2025). This process consumes dozens of kilowatts of electricity per machine. Beyond the energy concern, GenAI and QC rely on complex global supply chains that use rare, toxic minerals and create a mountain of electronic waste, because the hardware becomes outdated very fast (Wang et al., 2024). Despite these serious issues, the discussion about improving the sustainability of GenAI and QC has been fragmented and treated as unimportant, secondary to performance and scale. A few isolated studies (Berthelot et al., 2024; Cordier et al., 2025; Ding et al., 2025) have looked at the carbon footprint of an AI model or the energy use of quantum hardware; an integrated analysis comparing the two has been missing. Most of the existing reviews (Bolón-Canedo et al., 2024; Hasan et al., 2025; Hidayat and Surendro, 2025) focused only on AI or only on QC, failing to provide a bigger picture that could reveal shared solutions.

The primary objective of this paper is to provide a comparative review of the environmental impacts of GenAI and QC. As these technologies move toward mass industrial adoption, it is critical to evaluate how we can break the link between rapid computational growth and environmental damage. To achieve this, the study is organized around three guiding pillars: (i) Multidimensional Footprinting, utilizing Life Cycle Assessment (LCA) to move beyond simple carbon accounting; (ii) Mitigation Paradigms, evaluating software-hardware co-design and algorithmic efficiency; and (iii) Strategic Harmonization, identifying the metrics required for transparent, cross-paradigm sustainability reporting. Guided by these themes, the paper addresses the following three interconnected research questions:

- RQ1: How do the environmental profiles of GenAI and QC diverge in terms of energy intensity, material scarcity, and “hidden” lifecycle impacts such as water usage and abiotic depletion?
- RQ2: To what extent can technical interventions, such as model pruning and room-temperature qubits, and policy frameworks

like the “SEA-framework,” effectively mitigate the projected ecological burdens?

- RQ3: What are the critical barriers to establishing standardized, end-to-end LCA benchmarks that allow for transparent comparisons between classical and quantum paradigms?

By addressing these questions, this review makes four distinct contributions to the field of sustainable computing. First, it provides a Cross-Paradigm Lifecycle Synthesis that consolidates disparate data, specifically highlighting the “Inference-to-Training” gap in GenAI and the “Cryogenic-to-Processor” energy dichotomy in QC. Second, it establishes a Green Quantum Advantage Assessment, offering a comparative framework for identifying the pivotal threshold at which these technologies transition from environmental liabilities to net-positive sustainability catalysts. Third, the study provides a Comprehensive Mitigation Taxonomy, evaluating a spectrum of interventions ranging from software-level optimizations, such as AVATAR and Retrieval-Augmented Generation (RAG), to structural hardware shifts, such as stratospheric platforms. Finally, the paper culminates in a Strategic Research Roadmap and Policy Framework, identifying systemic barriers to transparency and proposing a standardized reporting roadmap that integrates Compute Carbon Intensity (CCI) and Abiotic Depletion Potential (ADP) metrics into future regulatory governance.

## 2 Foundational context: generative AI and quantum computing

This section establishes a foundational understanding of the two principal technologies discussed in this study: GenAI and QC. By exploring their core principles, historical trajectories, and contemporary applications, we provide the necessary context for a detailed analysis of their environmental impacts in subsequent sections.

### 2.1 Introduction to Generative Artificial Intelligence

GenAI is a type of computer program that can be used to create completely new, realistic-looking data. It is different from the older kind of AI (discriminative AI), which mainly sorts, labels, or identifies things that already exist. GenAI learns the rules and structures hidden inside a large-scale set of existing data, and then uses those rules to create new content that looks or sounds very similar to the original (Zheng et al., 2023). This means GenAI can generate all sorts of things, like new text, images, audio, and even computer code. The technology behind GenAI has changed a lot. Early GenAI used things like GANs (Generative Adversarial Networks). Think of a GAN as two AI programs fighting: one (the generator) tries to create a fake image, while the other (the discriminator) tries to spot the fake. This competition makes the output better over time. The biggest game-changer was the Transformer architecture (Jamil et al., 2023). This technology was first used to understand language, but it worked so well that it was used to create very large models that could understand and create long, sensible text. This led to the creation of Large Language Models (LLMs), like GPT-4. These are the programs that have impressed everyone because they can perform complex tasks like writing essays, summarizing long documents, and chatting in a very human-like way. For creating beautiful pictures and videos, diffusion models are now

popular. These models work by capturing a clear image, scrambling it by adding noise, and then learning how to reverse that process. This allows them to start with random noise and “de-scramble” it into a brand-new, photorealistic image. However, there is a catch. These powerful new generative models are huge and incredibly demanding on resources. They need a massive amount of computing power and energy just to be trained and to run (inference). Some of the biggest LLMs consist of trillions of parameters (Kumar et al., 2025). This huge size means their need for electricity and specialized hardware is substantial, making their environmental footprint a serious issue that requires urgent attention.

## 2.2 Introduction to quantum computing

QC is a brand-new way to process information that is completely different from classical computers (i.e., computers we use every day). QC gets its power from quantum mechanics, which is the set of laws that govern the subatomic world. Instead of using simple bits (which are always either a 0 or a 1), a quantum computer uses qubits (Vasanthrao and Saxena, 2025). A qubit can be in a state called superposition, meaning it can act as both a 0 and a 1 at the same time. This, combined with entanglement (where qubits become linked and instantly affect each other no matter how far apart they are), allows a quantum computer to handle certain calculations exponentially faster than any traditional computer (Columbus Chinnappan et al., 2025). Today, scientists are exploring a few different ways to build these delicate machines. The leading technologies include superconducting qubits (the type used by major companies like Google and IBM) (Kakarla et al., 2025) and ion traps (where tiny, charged atoms are suspended and controlled by magnets and lasers) (Kolachevsky et al., 2025). A major challenge for all these methods is that they require environments that are extremely stable and controlled. Many quantum computers must operate at temperatures near absolute zero to keep the qubits stable. This requires complex, energy-intensive cooling systems. Although still in its early phase, QC is expected to eventually solve problems that are impossible for even the world’s most powerful supercomputers. Its applications are expected to transform areas like finding new drugs, creating new materials, breaking or securing codes (cryptography), and solving complex logistical and optimization challenges (Kavitha et al., 2025).

## 2.3 Methodology

The literature search for this study was conducted using Google Scholar and ScienceDirect. Google Scholar served as a broad academic search engine to identify relevant publications, while ScienceDirect provided access to peer-reviewed journals and conference proceedings in computer science, physics, and environmental science.

### 2.3.1 Search strategy and database selection

A specific search strategy was used to find relevant articles for this research. The keywords were chosen carefully. Two main groups of search terms were created, one for each technology. For Sustainable GenAI, the search included phrases like “environmental impact of Generative AI,” “LLM,” “carbon footprint,” and “Sustainable AI” with “Life Cycle Assessment.” For Sustainable Quantum Computing, the search phrases include “environmental impact of Quantum


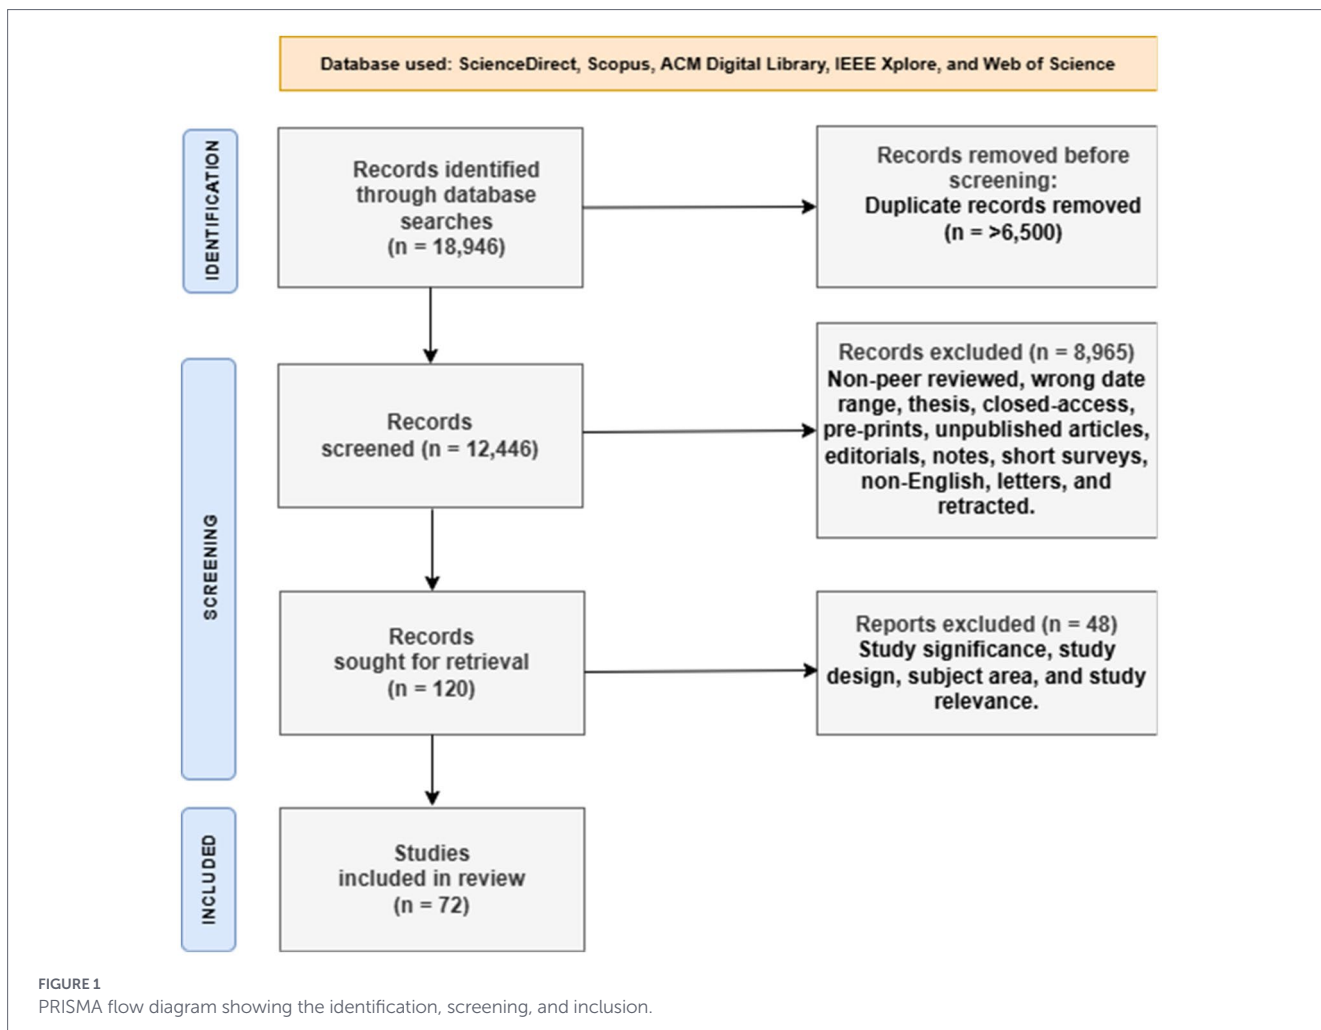
Computing,” “Quantum technology,” “sustainability,” “Quantum computing,” and “carbon footprint.” This organized search method ensured that most of the research was thoroughly covered on both topics.

### 2.3.2 Inclusion and exclusion criteria

The search was primarily focused on articles published between 2023 and 2026 to capture the most recent advancements and research findings. A small number of important and highly relevant articles published before 2023 were also included to provide historical context. The inclusion criteria were restricted to peer-reviewed journal articles and a few conference papers that provided empirical data, theoretical analysis, or a comprehensive overview of the environmental footprint of the specified technologies. To ensure the rigor and academic quality of the review, several types of publications were excluded. Specifically, any unpublished articles or preprints were not included in the review. Also, academic theses or dissertations, as well as news articles, blog posts, or other web-based content that were not peer-reviewed, were excluded. This selective process ensured that the analysis is based on a high-quality, relevant set of academic literature. Figure 1 presents the PRISMA flow diagram, illustrating the identification, screening, and inclusion processes. Figure 2 shows the breakdown of the total number of articles reviewed in this study.

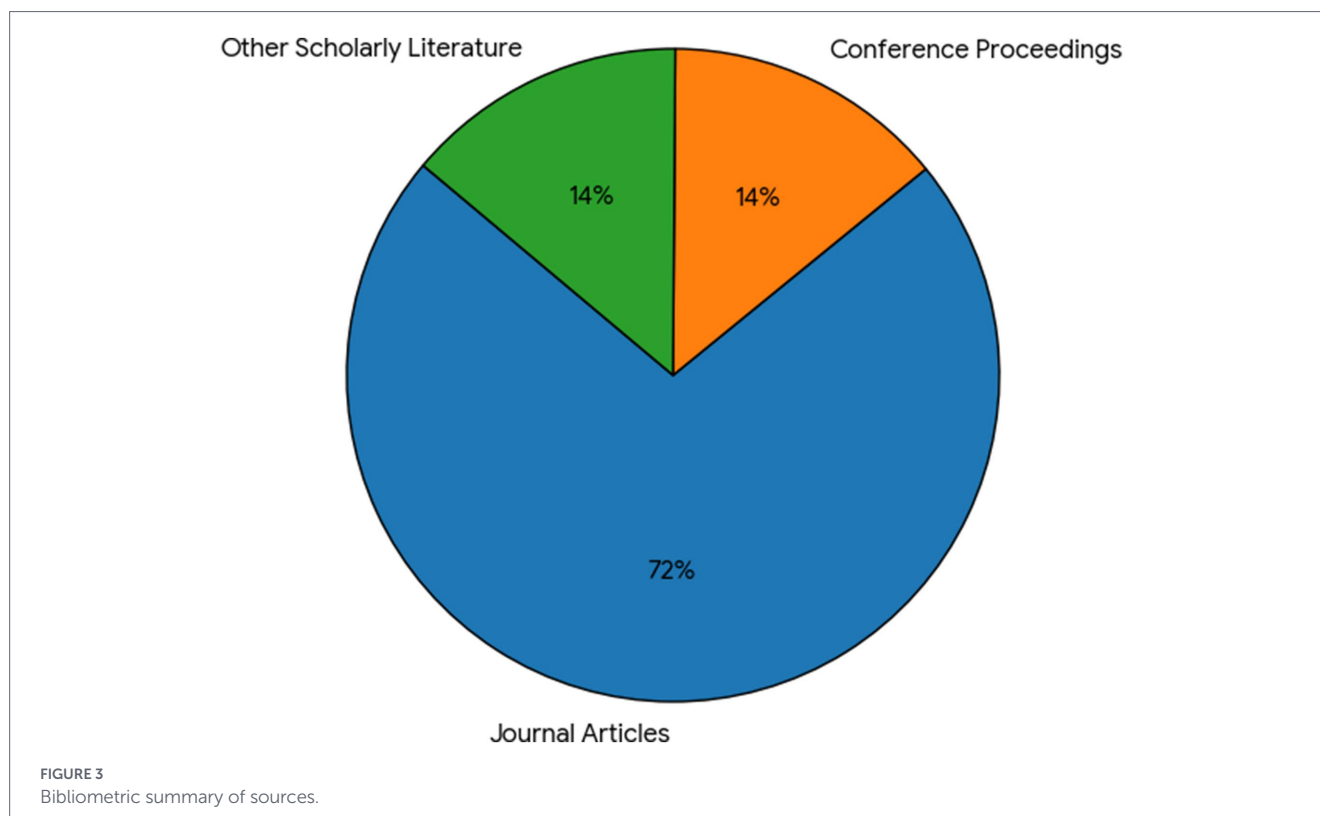
### 2.3.3 Bibliometric analysis

The bibliometric analysis of the selected literature demonstrates a high level of academic rigor and a focus on current, formally vetted research. As shown in Figure 3, peer-reviewed journal publications form the cornerstone of this study, accounting for 52 out of the 72 sources (approximately 72%). This high proportion reflects the scholarly depth and established scientific consensus integrated into the domains of AI sustainability and quantum hardware. Conference proceedings constitute 10 sources (approximately 14%). In the rapidly evolving fields of computer science and engineering, proceedings from top-tier venues are considered primary literature, such as the International Conference on Machine Learning (ICML) and IEEE High Performance Extreme Computing (HPEC). These sources employ rigorous double-blind peer review and are essential for capturing state-of-the-art breakthroughs and algorithmic optimizations that may not yet appear in traditional journal cycles. The remaining 10 sources (approximately 14%) are classified as “Other Scholarly Literature,” including book chapters from premier academic presses like Oxford University Press and Springer, as well as verified university-hosted research. A temporal analysis of the bibliography reveals a heavy emphasis on the most recent scientific developments. Articles published in 2025 and 2026 account for 48 of the 72 sources (66%), ensuring the inclusion of the latest empirical data on generative AI lifecycle assessments and fault-tolerant quantum hardware. Specifically, 12 articles (17%) are dated 2026, providing a forward-looking perspective on emerging “Green AI” techniques and sustainable quantum dot synthesis. Articles from 2024 (15 sources), 2023 (7 sources), and pre-2023 (2 sources) provide the necessary longitudinal context and baseline theoretical frameworks. In terms of publication venues, the literature is predominantly drawn from high-impact global publishers. Elsevier (ScienceDirect) contributes the largest share, representing 33% of the total references (24 sources). This is followed by Springer Nature at 17% (12 sources), IEEE at 14% (10 sources), and the ACM Digital Library at 11% (8 sources). Additionally, MDPI accounts for



8% of the sources, while the remaining 17% is comprised of specialized publishers and academic presses such as Oxford, IGI Global, and Frontiers. This diverse distribution underscores the interdisciplinary

breadth of the study, spanning environmental science, hardware engineering, and computational theory, while ensuring all data is anchored in academic credibility and rigorous peer-review standards.



### 3 Literature review

GenAI and QC present both opportunities and environmental challenges. These two technologies have the potential to transform different industries; however, their emergence is accompanied by a substantial ecological footprint. This section presents a comprehensive review of the environmental impact of GenAI, specifically focusing on important issues such as energy consumption, carbon emissions, and resource depletion. In the review, the methodologies used to assess these impacts are analyzed, including LCA. The review also explores the sustainability trade-offs among different models and deployment strategies. This literature review aims to synthesize existing research, identify key gaps, and outline innovative techniques to mitigate the ecological costs of GenAI and QC.

#### 3.1 Literature review on generative AI

This section provides a structured review of the current literature on GenAI sustainability, focusing on the precision of environmental impact measurements, the efficacy of technical mitigation strategies, and the evolving role of governance in ensuring responsible innovation.

##### 3.1.1 Environmental impact of generative AI

This section synthesizes recent studies categorized into three critical themes. It first examines the scope of what is measured in life cycle assessments, then analyzes the timing of energy consumption from development to daily use, and finally evaluates how specific design and infrastructure choices determine the technology's total footprint.

##### 3.1.1.1 Life-cycle perspectives and hidden environmental costs of generative AI

Current research marks a critical transition from simplistic carbon accounting to comprehensive LCAs that capture the “hidden” ecological costs of Generative AI. While early estimates focused primarily on the staggering electricity consumption of model training, such as the 1,287 MWh required for GPT-3 (Hosseini et al., 2025), contemporary frameworks by Berthelot et al. (2024), Berthelot et al. (2025) and d’Orgeval et al. (2026) argue that carbon is merely one variable in a broader crisis of resource depletion. This “operational-centric” bias often ignores embodied carbon (Nguyen et al., 2024) and the abiotic depletion of rare materials required for hardware manufacturing. The literature suggests that the total environmental impact of a model is highly sensitive to infrastructure variables; for instance, the disparity between air-cooled and liquid-cooled data centers can result in a 70% difference in Global Warming Potential per eFLOP (d’Orgeval et al., 2026). However, a significant gap remains in the transparency of Small Language Models (SLMs) and data center overheads, which are frequently omitted from empirical benchmarks, leading to an underestimation of the total ecological burden (Jeanquartier et al., 2026; Ji and Jiang, 2026). Identifying these lifecycle variables provides a general picture of AI’s footprint. However, looking at the different phases of energy use, from initial development to daily use, shows that the environmental impact is much more unbalanced than it first appears.

##### 3.1.1.2 The shift from training intensity to inference-phase sustainability

The environmental gravity of GenAI has shifted from the intensive event of training to the continuous process of inference. While training large-scale models like Llama 3.1-405B remains a massive

financial and energetic undertaking (d'Orgeval et al., 2026), the cumulative cost of billions of daily user queries creates a dominant footprint that can exceed training emissions by a factor of 25 annually (Boulais et al., 2026; Chien et al., 2023; Samsi et al., 2023). This realization has catalyzed a surge in “carbon-aware” computing and algorithmic efficiency research. Strategies ranging from the AVATAR optimization method (Shi et al., 2024) to the CarbonMin scheduling algorithm (Chien et al., 2023) can substantially reduce a model's projected emissions. Furthermore, the emergence of Mixture of Experts (MoE) architectures, as seen in DeepSeekV3, proves that computational efficiency, activating only a fraction of parameters per token, can reduce training costs by an order of magnitude compared to traditional dense models (d'Orgeval et al., 2026). Ultimately, the literature underscores that sustainable AI requires a decoupling of model utility from energy demand through geographic relocation to low-carbon grids and the adoption of more efficient architectural paradigms (Ding et al., 2025; Li et al., 2025).

### 3.1.1.3 Architectural innovation, infrastructure dependency, and the emerging inference gap

Table 1 consolidates metrics across the AI lifecycle, highlighting the variance in environmental costs based on architecture, infrastructure, and phase of use. The synthesized data reveals a divergence between the escalating computational demands of state-of-the-art models and the emerging efficiency of specialized architectures. The comparison between GPT-4o and DeepSeek V3 is particularly telling; by using a Mixture of Experts (MoE) approach and FP8 precision, DeepSeek achieved a 90% reduction in training costs and significantly lower FLOP requirements than its dense counterparts (d'Orgeval et al., 2026). This suggests that the “brute force” scaling laws of the past are being challenged by architectural innovation. Furthermore, recent comprehensive life-cycle assessments indicate that the environmental impact is not an inherent property of the model alone but is dictated by the supporting infrastructure. While hardware-level interventions like liquid-cooling can reduce Global Warming Potential per eFLOP by over 70% (d'Orgeval et al., 2026), new findings emphasize that embodied emissions from the manufacturing and disposal of AI accelerators are a massive, often overlooked contributor to the total footprint. The introduction of the Compute Carbon Intensity (CCI) metric reveals that while hardware efficiency improved 3x from Tensor Processing Unit (TPU) v4i to TPU v6e, manufacturing still represents a significant portion of a chip's cradle-to-grave emissions (Schneider et al., 2025). However, the most significant critical take-away remains the “Inference Gap”: with inference potentially outstripping training emissions by 25 times annually (Chien et al., 2023), the industry's historical focus on training-phase metrics is increasingly obsolete. Future sustainability frameworks must prioritize real-time, inference-phase telemetry, embodied carbon accounting, and edge-deployment strategies to mitigate the carbon-intensive reality of daily AI use.

### 3.1.2 Mitigation strategies for the environmental impact of GenAI

Mitigating the environmental burden of GenAI begins at the architectural level, where recent research highlights a shift from “brute-force” scaling to high-efficiency design. Algorithmic optimization techniques, such as the AVATAR method, have demonstrated that

TABLE 1 Comparative synthesis of quantitative findings in GenAI sustainability.

Category	Specific metric/ finding	Model(s)/ framework	Primary source(s)
Energy consumption (Training)	1,287 MWh (552 tCO <sub>2</sub> e)	GPT-3	Hosseini et al. (2025)
	456 MWh (40 tCO <sub>2</sub> e)	Google GLAM	Ji and Jiang (2026)
	1.5 MWh	BERT	Ji and Jiang (2026)
Computational intensity	$3.81 \times 10^{25}$ FLOP	GPT-4o / Llama 3.1-405B	d'Orgeval et al. (2026)
	$3.41 \times 10^{24}$ FLOP	DeepSeek V3 (MoE)	d'Orgeval et al. (2026)
Hardware lifecycle	3x Efficiency Improvement	TPU v4i to TPU v6e	Schneider et al. (2025)
	Embodied Carbon Analysis	Life-Cycle Assessment	Schneider et al. (2025)
Economic cost	\$62 Million vs. \$6 Million	Llama 3.1 vs. DeepSeek V3	d'Orgeval et al. (2026)
Infrastructure efficiency	59.59 vs. 199 gCO <sub>2</sub> eq/eFLOP	Liquid- vs. Air-cooled	d'Orgeval et al. (2026)
Operational impact	Inference = 25x Training/yr	General Workload Model	Chien et al. (2023)
Optimization gains	184x Energy Efficiency	AVATAR Method	Shi et al. (2024)
	56% Emission Reduction	CarbonMin Algorithm	Chien et al. (2023)

software-level interventions can significantly reduce energy consumption and carbon emissions during large language model inference without substantially degrading performance (Shi et al., 2024). Beyond compression, architectural shifts like the Mixture of Experts (MoE), exemplified by models like DeepSeekV3 have proven disruptive, allowing for large-scale performance at a fraction of the computational and financial cost of traditional dense models (d'Orgeval et al., 2026). These technical solutions are further bolstered by user-driven strategies, where smarter prompting and selecting task-appropriate model sizes directly curtails unnecessary energy expenditure during the inference phase (Jeanquartier et al., 2026; Ren et al., 2024). While algorithmic gains are substantial, the literature emphasizes that true sustainability requires systemic changes in how and where AI is deployed. Infrastructural interventions, such as transitioning to high-density liquid-cooled data centers, can reduce the carbon footprint per eFLOP by over two-thirds compared to traditional air-cooled systems (d'Orgeval et al., 2026). Strategic workload management also plays a pivotal role; for example, the CarbonMin algorithm enables carbon-aware computing by shifting requests to regions with cleaner energy grids, projecting emission reductions of up to 56% by 2035 (Chien et al., 2023). Additionally, shifting workloads from cloud to edge devices offers a localized solution that achieves over 90% energy savings by eliminating the need for massive, centralized cooling and transmission overheads (Li et al., 2025). Ultimately, researchers argue that these technical fixes must be integrated into broader policy frameworks that treat environmental, social, and economic sustainability as

interconnected pillars, ensuring that AI efficiency gains are not offset by the rapid growth in global data center demand (Bashir et al., 2024; Waelen and van Wynsberghe, 2025). Table 2 categorizes the primary methods discussed in this sub-section for reducing the environmental impact of GenAI.

### 3.1.3 Governance and responsible innovation for generative AI

Based on the review of studies, several key findings related to Governance & Responsible Innovation emerge. Scholars argue that the societal and environmental implications of these technologies necessitate a move beyond mere efficiency toward “responsible innovation” frameworks that integrate ethical, legal, and social considerations into the R&D process (Lee et al., 2025; Waelen and van Wynsberghe, 2025). For instance, the proposed “SEA-framework” (Safeguarding, Engaging, and Advancing) for Responsible Quantum Technology (Root, 2025) underscores the importance of prioritizing Environmental, Social, and Governance (ESG) criteria in research and investment (Root, 2025). Similarly, current literature advocates for ‘Carbon-aware Quantum Computing’ principles, highlighting the need for standardized benchmarking of total life-cycle carbon footprints (Arora and Kumar, 2026; Cordier et al., 2025). Studies also emphasize that the governance of such emerging technologies should involve policy interventions, stakeholder engagement workshops, and a commitment to “technological frugality” and the use of certified materials to minimize ecological harm (Bashir et al., 2024; Wang et al., 2024). Ultimately, these findings suggest that effective governance requires a multi-faceted approach where technical optimizations are supported by robust regulatory and ethical guidelines to ensure long-term sustainability (Christensen et al., 2021).

## 3.2 Literature review on quantum computing

QC, while offering transformative computational power, also raises important questions about its environmental footprint and sustainability. This section explores both the potential environmental impacts of quantum technologies and their capacity to drive solutions for global sustainability through responsible innovation and lifecycle management.

### 3.2.1 Environmental impact of quantum computing

This section examines the environmental landscape of QC by categorizing current literature into three thematic areas: the scope of its total ecological footprint including rare material use, the timing and energy overhead of maintaining fault-tolerant systems, and the design and infrastructure required to bridge the gap between cryogenic costs and industrial sustainability gains.

#### 3.2.1.1 Quantum infrastructure costs and hidden material risks

Measuring the environmental impact of QC reveals a huge gap between the tiny amount of energy the processor uses and the massive amount of power needed for the equipment that supports it. Current life-cycle sustainability discussions highlight that, although superconducting quantum processors themselves consume relatively little power, the associated cryogenic cooling infrastructure represents a significant share of the overall energy footprint of quantum computing systems (Chauhan et al., 2025). This cooling challenge is further complicated by the scaling of Quantum Error Correction (QEC); as systems transition toward fault tolerance, the energy overhead of maintaining stable qubits is projected to become the dominant environmental cost (Cordier et al., 2025; De Michielis and Ferraro, 2025). Beyond energy, the quantification of impact must account for “hidden” material risks. The reliance on rare-earth metals, noble gases, and toxic nanomaterials introduces significant toxicity concerns, including DNA damage and oxidative stress resulting from heavy metal ion release (Arora and Kumar, 2026; Giroux et al., 2022). Consequently, a comprehensive LCA must move beyond simple power metrics to include the global warming potential (GWP) of the entire hardware supply chain (Gheibi et al., 2026).

#### 3.2.1.2 Green quantum advantage and industrial decarbonization potential

The quantification of QC’s environmental utility is increasingly measured against the “Green Quantum Advantage Threshold,” the pivotal point where quantum efficiency surpasses classical

TABLE 2 Summary of mitigation strategies for Generative AI sustainability.

Level of intervention	Specific strategy/technology	Environmental & operational benefit	Primary source(s)
Architectural /code	Mixture of Experts (MoE) & AVATAR method	High performance at a fraction of traditional computational costs.	d’Orgeval et al. (2026); Shi et al. (2024)
User-driven	Smart prompting & task-appropriate model selection	Directly curtails unnecessary energy waste during inference.	Jeanquartier et al. (2026)
Infrastructure	High-density liquid cooling	Reduces carbon footprint per eFLOP by over 66% compared to traditional air-cooled mixes.	d’Orgeval et al. (2026)
Systemic/Grid	CarbonMin scheduling algorithm	Shifting workloads to cleaner grids; projected 56% emission reduction by 2035.	Chien et al. (2023)
Deployment	Cloud-to-Edge computing shift	Achieves over 90% energy savings by eliminating centralized cooling overhead.	Li et al. (2025)
Policy	Integrated Sustainability Frameworks	Ensures efficiency gains are not lost to rapid industry growth (rebound effect).	Bashir et al. (2024), Waelen and van Wynsberghe (2025)

TABLE 3 Comparative quantitative synthesis of QC environmental metrics.

Category	Qualitative/quantitative finding	Context/comparison	Source(s)
Infrastructure & cooling	High supporting power demand	Consumption of one dilution refrigerator exceeds multiple average households.	Chauhan et al. (2025)
Operational scaling	Significant power per qubit	Scaling to Fault-Tolerant systems (QEC) shifts impact from manufacturing to usage.	Cordier et al. (2025); De Michielis and Ferraro (2025)
Simulation overhead	Substantial emission disparity	Simulating quantum circuits on classical hardware is far more carbon-intensive than training large ML models.	Li et al. (2023)
Algorithmic efficiency	Orders of magnitude energy saving	Native quantum search algorithms vs. classical supercomputers for large-scale operations.	Jaschke and Montangero (2023)
Carbon mitigation	Massive projected annual offsets	Long-term potential via quantum-accelerated grid and battery optimization.	Vudugula et al. (2025)
Industrial impact	Meaningful carbon reduction	Optimization of HVAC and industrial systems using quantum-inspired methods.	Hidayat and Surendro (2025)
Macro-context	Rising share of global energy	Projected increase in total ICT sector consumption by the end of the decade.	Arora and Kumar (2026)

high-performance computing (Jaschke and Montangero, 2023). Current comparative studies indicate that the “hidden cost” of the pre-quantum era is substantial; for instance, simulating quantum circuits on classical hardware can generate up to 48 times the carbon emissions of training large-scale machine learning models (Li et al., 2023). However, the literature suggests that transitioning to native quantum execution may offer energy-consumption advantages for certain computational tasks, with potential implications for reducing the carbon footprint of large-scale digital systems (Chauhan et al., 2025). While “input–output bottlenecks” in data preparation for QRAM may initially diminish these gains (Morstyn and Wang, 2024), the long-term projection identifies QC as a net-positive environmental tool. In particular, the literature highlights the potential role of quantum-accelerated modeling in supporting decarbonized power grids, thereby contributing to long-term emissions reduction pathways (Chauhan et al., 2025). Thus, the critical analysis of QC’s footprint reveals that while the immediate infrastructure is energy-intensive, its algorithmic precision serves as a vital mechanism for achieving global carbon neutrality.

### 3.2.1.3 Sustainability catalysts and the energy transition gap

Table 3 consolidates the environmental and energy data points extracted from the literature, categorizing them by the nature of their impact. The synthesized data reveals a profound decoupling between the energy efficiency of the quantum processor and the energy intensity of its supporting environment. While theoretical frameworks suggest that a “Green Quantum Advantage” can lead to exponential energy savings (Jaschke and Montangero, 2023), the current physical reality is dominated by the substantial power overhead required for cryogenic cooling (Chauhan et al., 2025). This disparity creates a sustainability transition gap, where the environmental cost of simulating and developing quantum systems on classical hardware (Li et al., 2023) currently outweighs their operational utility. However,

emerging research into quantum-accelerated catalyst design and logistics optimization indicates a significant potential for reducing carbon intensity across high-emission industrial sectors (Hidayat and Surendro, 2025). These findings suggest that the primary environmental value of quantum computing lies not in its direct energy consumption, but in its role as a “sustainability catalyst” that optimizes complex chemical and industrial cycles. Long-term projections indicate that the path to a net-positive environmental impact depends on leveraging these technologies for global grid and supply chain optimization (Vudugula et al., 2025). Ultimately, achieving this balance requires overcoming the data transfer inefficiencies and mitigating the energy costs associated with maintaining fault-tolerant error correction.

### 3.2.2 Mitigation strategies for the environmental impact of quantum computing

Mitigation strategies for the environmental impact of quantum computing focus primarily on decoupling computational power from the energy-intensive cooling infrastructure. A significant technical pathway involves the transition to room-temperature architectures, such as photonic qubits or diamond Nitrogen-Vacancy centers, which aim to eliminate the massive energy overhead required to maintain the near-absolute zero temperatures necessary for superconducting systems (Vasantrao and Saxena, 2025). Complementing these hardware shifts is the emergence of “Green Quantum Software” and NISQ-friendly algorithms. By prioritizing algorithmic efficiency and implementing unit testing protocols like assertion checks to prevent wasteful re-runs of flawed operations, developers can significantly reduce the energy overhead inherent in early-stage quantum experimentation (Dwivedi et al., 2024). Furthermore, advanced computational modeling inspired by quantum principles can refine the product design phase itself, leading to major reductions in the carbon footprint and waste generated during the manufacturing lifecycle (Hassan and Hassan, 2025; Lukyanenko, 2025). Beyond direct technical interventions, researchers advocate for the integration of

sustainability into the foundational design and governance of quantum technology. The adoption of frameworks such as responsible innovation and the “SEA-framework” (Safeguarding, Engaging, Advancing) provides a structured approach to embedding environmental and ethical considerations into the R&D phase before mass adoption occurs (Root, 2025). This proactive stance is exemplified by the application of Safe-and-Sustainable-by-Design (SSbD) principles, which utilize AI-driven predictive modeling to mitigate the toxicity and ecological risks of specialized materials like silicon nanowires and heavy metals (Gheibi et al., 2026). By establishing metrics such as the “green quantum advantage threshold,” the research community is moving toward a standard where quantum progress is measured not just by speed, but by its net-positive contribution to global decarbonization (Jaschke and Montangero, 2023; Theophilo and Oughton, 2025). Table 4 outlines the technical and framework-based solutions currently being researched to reduce the ecological footprint of QC.

### 3.2.3 Governance and responsible innovation

The governance of emerging technologies like QC is shifting toward a proactive, values-based model that integrates environmental stewardship directly into the innovation lifecycle. Central to this transition is the responsible innovation framework, which advocates for “anticipatory governance,” identifying potential ecological and social harms while the technology is still in its malleable, pre-industrial stage (Root, 2025). This perspective aligns with emerging work on Carbon-Aware Quantum Computing (CQC), which emphasizes the need to account for both embodied carbon arising from rare-earth mining and system fabrication, as well as operational carbon associated with energy-intensive cryogenic cooling infrastructures, in alignment with the United Nations Sustainable Development Goals (Arora and Kumar, 2026; Theophilo and Oughton, 2025). By establishing standardized metrics like the “green quantum advantage threshold,” policymakers can move beyond commercial hype to enforce accountability in energy consumption and material sourcing (Jaschke and Montangero, 2023). Ultimately, this integrated policy approach aims to prevent a “quantum divide,” ensuring that the transition to post-classical computing is globally equitable and environmentally regenerative rather than merely computationally superior.

TABLE 4 Summary of mitigation strategies for quantum computing sustainability.

Strategy category	Key technical intervention	Environmental benefit	Primary source(s)
Hardware innovation	Transition to room-temperature architectures (e.g., photonic qubits).	Eliminates the massive energy overhead required for extreme cryogenic cooling.	Vasantrao and Saxena (2025)
Software efficiency	Green Quantum Software & assertion checks (unit testing).	Reduces energy waste by preventing flawed operation re-runs.	Dwivedi et al. (2024)
Design optimization	Quantum-inspired modeling in product design.	Reduces carbon footprint and waste during the manufacturing phase.	Hassan and Hassan (2025)
Governance frameworks	Responsible Innovation and the SEA-framework.	Embeds ethical and environmental considerations into early R&D.	Root (2025)
Material safety	Safe-and-Sustainable-by-Design (SSbD) principles.	Mitigates toxicity risks of silicon nanowires and heavy metals using AI.	Gheibi et al. (2026)
Standardization	Green Quantum Advantage Threshold.	Measures QC progress by its net-positive contribution to decarbonization.	Jaschke and Montangero (2023)

## 4 Review discussion

This section reviews the core sustainability assessment frameworks for both GenAI and QC, synthesizing key comparative insights between the two domains. It then outlines practical mitigation strategies and proposes a roadmap for sustainable pathways, addressing both immediate technical solutions and long-term policy and research challenges. By consolidating these topics, this paper provides a comprehensive look at the current state of sustainable computing and identifies critical future directions for responsible technological development.

### 4.1 Environmental impact of generative AI

The remarkable capabilities of GenAI models, particularly LLMs and diffusion models, are built upon a foundation of immense computational power (He et al., 2025; Mohammed and Kora, 2025). This power, however, comes with a substantial and multifaceted environmental cost, extending beyond the commonly cited energy consumption to encompass water usage, hardware manufacturing, and electronic waste (d’Orgeval et al., 2026). This section delves into the ecological footprint of GenAI, examining its key environmental impacts and the metrics used to quantify them.

#### 4.1.1 Energy consumption and carbon footprint

The energy demands of GenAI are significant throughout its lifecycle, from the initial training phase to the inference phase, where models are used to answer user queries. The training of massive models like GPT-3 is famously energy-intensive, consuming an estimated 1,287 MWh of electricity and emitting approximately 552 tons of CO<sub>2</sub> (Hosseini et al., 2025). To put that in perspective, this is enough energy to power an average U. S. home for over 120 years, or to drive a car around the world 38 times. This immense carbon footprint highlights the high environmental cost of developing these cutting-edge AI technologies. Other models like LLaMA and PaLM also have considerable training footprints (Ding et al., 2025). The carbon emissions associated with this energy consumption are heavily dependent on the local grid’s electricity mix, which determines how much carbon is emitted per kilowatt-hour of electricity used (Hosseini et al., 2025). While training is a one-time event for a given model, the

environmental impact of inference, billions of daily queries, can be even greater over time (Argerich and Patiño-Martínez, 2024). As the Bar Chart in Figure 4 clearly illustrates, studies show that inference-related emissions can be 25 times greater than training emissions in a single year (Chien et al., 2023), which will be 32,175 MWh ( $25 \times 1,287$ ). This is a critical point, as the sheer scale of global user interactions with GenAI services means that the cumulative impact of each small query is immense (Kumar et al., 2025; Li et al., 2025). For a service like Stable Diffusion, the inference cost per query is not negligible, and for a large user base, it can quickly become the dominant environmental factor (Li et al., 2025).

The bar chart in Figure 5 further contextualizes these costs by comparing the training emissions of LLMs to the carbon footprint of a long-haul flight, revealing the staggering ecological debt incurred even before a system is deployed for public use. While a single passenger's return flight from London to New York generates approximately 1.7tCO<sub>2</sub>e, the initial development of OpenAI's GPT-3 in a standard United States data center produced an estimated 552tCO<sub>2</sub>e (Hosseini et al., 2025), an impact exceeding 320 such trans-Atlantic journeys [ $552 \div 1.7 \approx 324.7$ ]. This disparity intensifies significantly with the latest frontier models; training GPT-4o or Llama 3.1-405B using traditional air-cooled infrastructure can result in emissions surpassing 15,000 tCO<sub>2</sub>e, matching the carbon footprint of roughly 8,823 individual flights [ $15,000 \div 1.7 \approx 8,823.5$ ]. However, recent lifecycle assessments demonstrate that these impacts are highly sensitive to technical and geographic variables. According to d'Orgeval et al. (2026), transitioning the GPT-4o training process to high-density, liquid-cooled architectures (such as the RD109 Full AI) can reduce emissions to 4,243tCO<sub>2</sub>e, effectively "saving" the carbon equivalent of over 6,320 flights through hardware optimization alone [ $(15,000 - 4,243) \div 1.7 \approx 6,327.6$ ]. Furthermore, the model BLOOM illustrates the transformative potential of regional deployment, as its training on the low-carbon French grid resulted in only 30 tCO<sub>2</sub>e, a footprint nearly 18 times smaller than the older GPT-3 despite its vast computational requirements [ $552 \div 30 \approx 18.4$ ]. These findings show that while the convenience of GenAI comes at a profound

environmental cost, the industry's "carbon debt" can be substantially mitigated. By prioritizing energy-efficient architectures and low-carbon energy sources, developers and companies can move toward a more responsible and sustainable paradigm for artificial intelligence.

In the development and deployment of Generative AI models, there is an inherent trade-off between model size, performance, and energy efficiency. Larger, more capable models typically require more energy for both training and inference, as they have more parameters and require more computations (Argerich and Patiño-Martínez, 2024). However, the location and timing of these computations are also critical. Deploying GenAI models in regions with a cleaner electricity mix (e.g., lower carbon intensity) can significantly reduce their carbon footprint, a strategy known as "carbon-aware computing" (Chien et al., 2023; Hosseini et al., 2025). For instance, shifting workloads to data centers in countries with abundant renewable energy can lead to substantial reductions in emissions (Chien et al., 2023).

#### 4.1.2 Hardware and infrastructure demands

The environmental impact of GenAI is not limited to electricity consumption. The hardware and infrastructure that power these models have their own significant ecological footprint throughout their life cycle, from manufacturing to disposal (Berthelot et al., 2024). As the LCA Flowchart in Figure 6 shows, the environmental costs of hardware begin long before a model is trained. This end-to-end perspective reveals that impacts are not confined to the use phase but extend from raw material extraction (leading to habitat loss) and energy-intensive manufacturing to transportation and, critically, end-of-life e-waste disposal (which can release hazardous materials into the environment).

Building on this framework, Cordier et al. (2025) provide a specialized life cycle assessment for superconducting quantum computing, tracing a path through four primary stages: Production, Delivery, Use, and End-of-Life (Figure 7). The Production Phase is characterized by the intensive extraction of raw materials and the

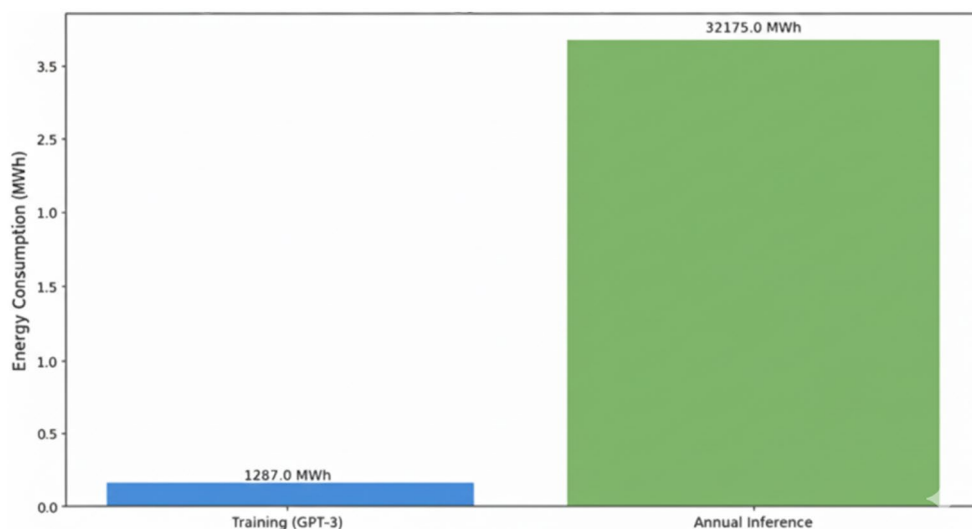
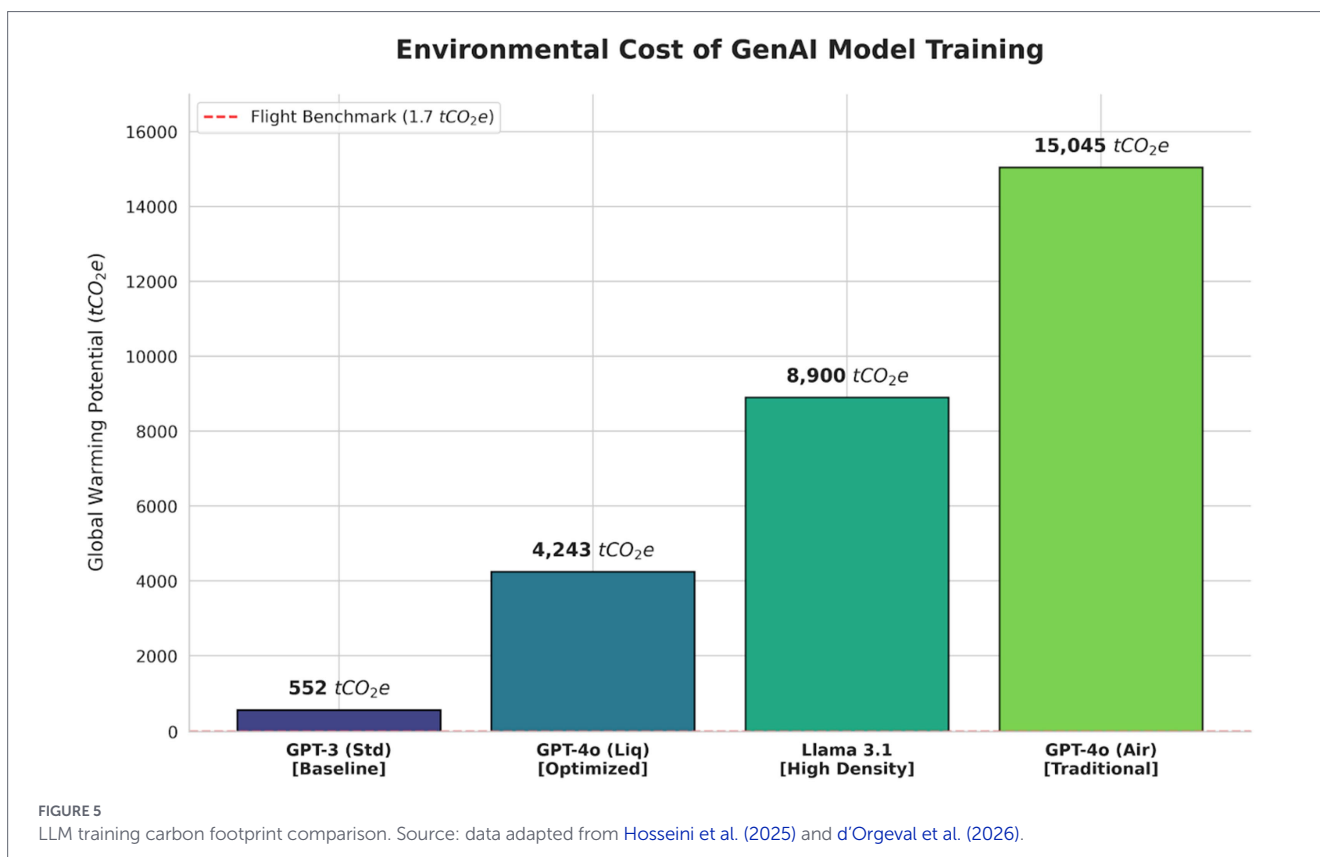


FIGURE 4

GenAI energy consumption: Training vs. annual inference. Source: calculation based on training data from Samsi et al. (2023) and the inference multiplier from Chien et al. (2023).

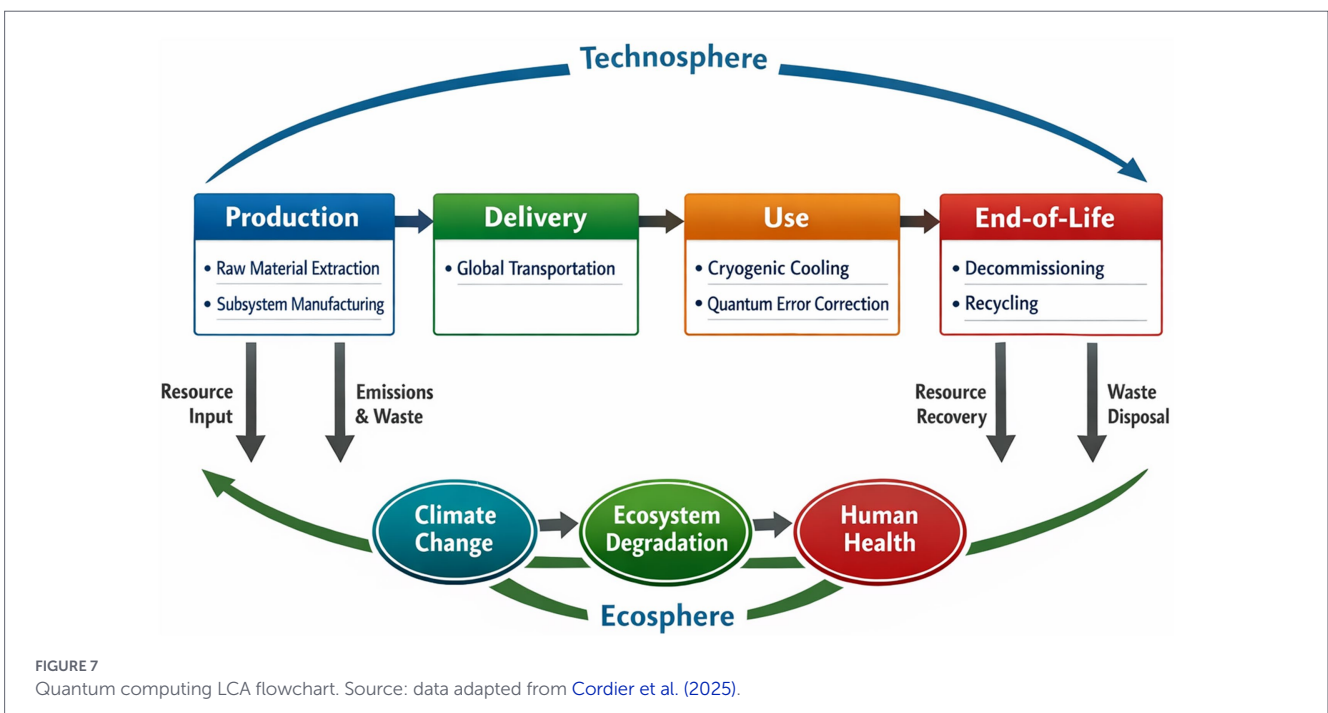
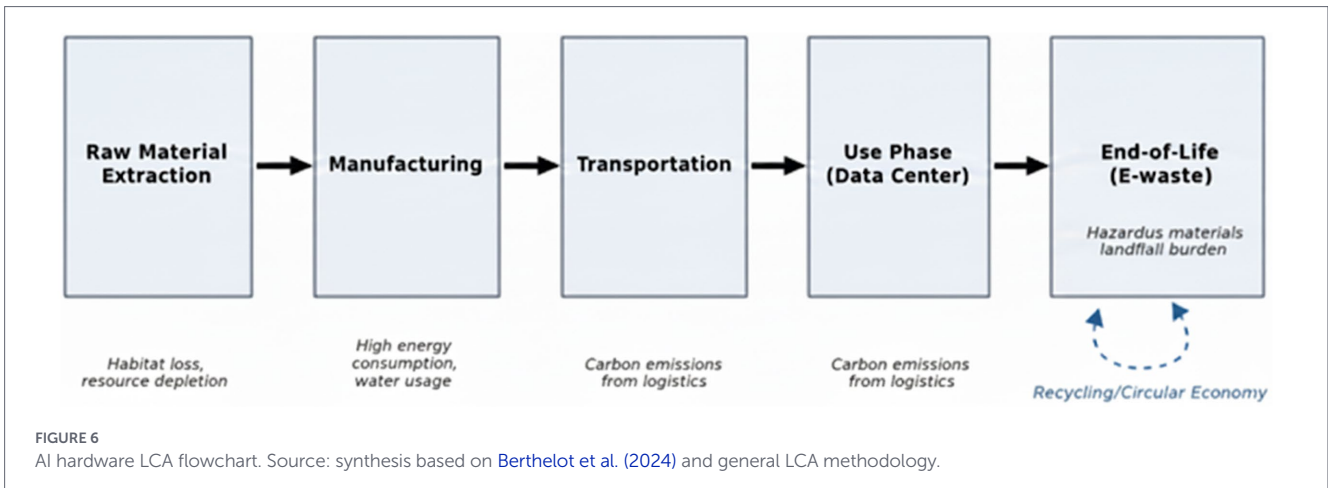


manufacturing of complex subsystems like cryostats and gas handling systems. This phase flows into the Delivery Phase, which accounts for the global logistics of transporting specialized hardware from international hubs. The sequence continues into the Use Phase, where impacts are dominated by the high energy demands of cooling systems and Quantum Error Correction (QEC) electronics. Finally, the End-of-Life Phase addresses the decommissioning of hardware, where recycling efforts can mitigate some burdens by returning materials to the start of new product cycles. The systemic relationships within this life cycle are illustrated by the flow of resources between the technosphere and ecosphere. While horizontal arrows in the flowchart trace the hardware's chronological journey, the overarching life-cycle loops define a comprehensive boundary that captures continuous environmental interactions. These inventory stages are ultimately translated into ecological consequences through an Impact Assessment using the IMPACT World+ methodology. This final step characterizes the total environmental burden into three critical areas: Climate Change (emissions), Ecosystem degradation, and Human Health (measured in disability-adjusted life years). By integrating this QC-specific data, we observe that while classical AI hardware impacts are primarily driven by chip fabrication, the quantum footprint is uniquely shaped by the extreme physical infrastructure required to maintain a stable qubit environment.

As illustrated in Figure 8, the environmental footprint of Generative AI is driven by a complex interplay between computational hardware and support infrastructure. While high-performance servers and GPUs are central to GenAI operations, research indicates that server usage alone accounts for approximately 70% of the total environmental impact during a model's operational phase

(d'Orgeval et al., 2026). The remaining impact is attributed to the "full architecture" of the data center, which includes power conversion systems and intensive cooling mechanisms. The energy required for these processes is substantial, especially for LLM inference, which carries significant energy costs per response (Samsi et al., 2023). Furthermore, Figure 8 highlights the cooling mechanism's reliance on large-scale water usage for evaporative towers; recent LCA confirm that this water consumption is a major contributor to the cloud's cumulative footprint, though its severity varies significantly by geographic region (d'Orgeval et al., 2026).

The manufacturing of GPUs and TPUs is also an energy-intensive process that relies on a complex global supply chain (Berthelot et al., 2024). These high-performance computing components require rare-earth minerals and other raw materials, the mining of which can contribute to deforestation, habitat destruction, and social impacts like child labor (Berthelot et al., 2024; Waelen and van Wynsberghe, 2025). The resulting embodied carbon (i.e., the emissions generated during manufacturing), is a crucial, but often overlooked, part of the overall environmental cost (Berthelot et al., 2024; Samsi et al., 2023). The comprehensive impact of these GPUs and TPUs is driven by their full hardware lifecycle, which contributes to global environmental challenges. The rapid pace of technological innovation in the AI sector leads to a short lifespan for these components, creating a massive e-waste problem. The disposal of servers, GPUs, and TPUs, and other components often involves hazardous materials that can contaminate soil and water (Berthelot et al., 2024). A holistic assessment of GenAI's footprint must therefore account for the entire life cycle of the hardware, from the mine to the landfill, as proposed by researchers using an LCA approach.



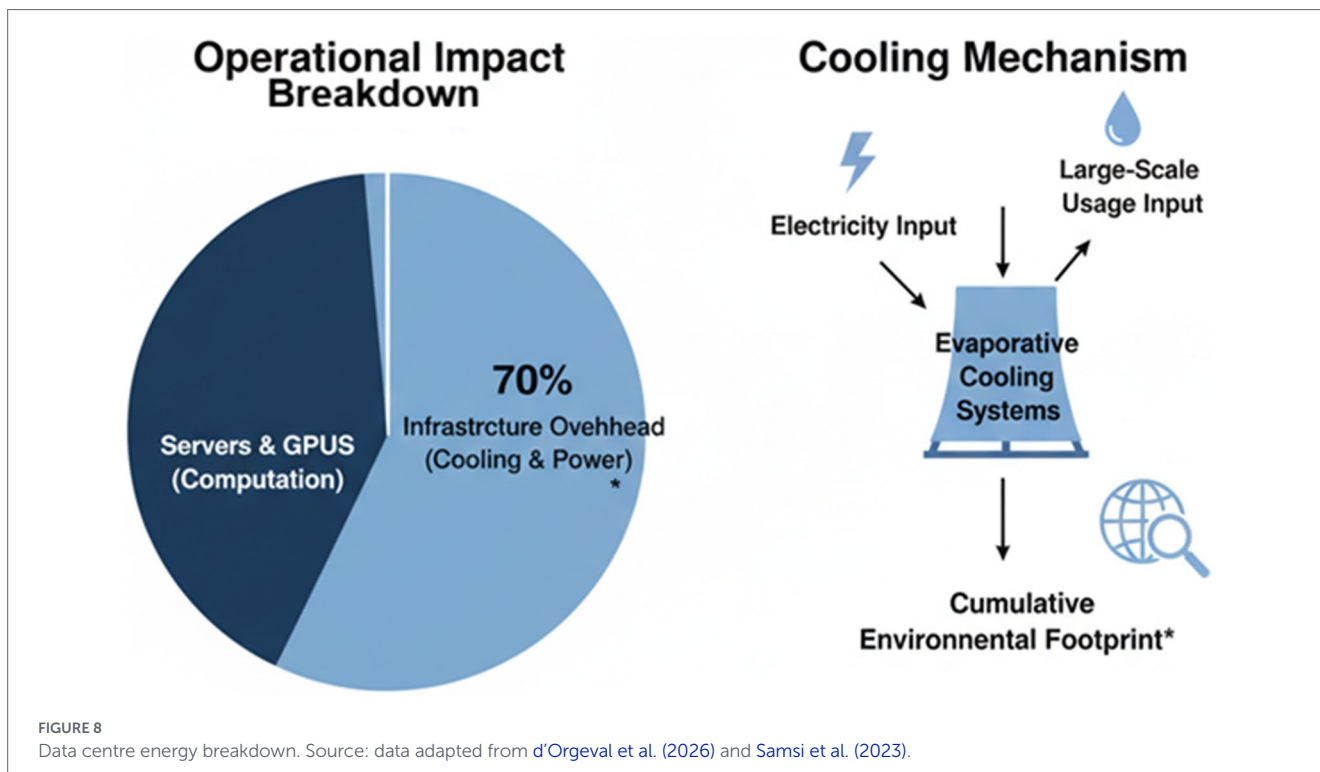
### 4.1.3 Sustainability metrics for AI

To deal with the environmental impact of GenAI effectively, there is a need for clear and standard ways to measure the damage. Researchers have developed several key metrics to achieve this:

- i Energy-per-inference and carbon-per-query: These metrics measure the energy used and the carbon released by a single action (like asking ChatGPT a question) (Li et al., 2025). This is important because the simple act of prompting AI models happens millions of times a day.
- ii FLOPs-per-watt (floating point operations per watt): This is a hardware metric, and it tells us how computationally efficient a computer system is. The metric explains how much useful work is done for every unit of energy used (Argerich and Patiño-Martínez, 2024).
- iii Abiotic depletion potential (ADP) and primary energy (PE): These metrics provide a holistic view that goes beyond just quantifying the environmental cost of carbon. Abiotic

Depletion Potential measures the reduction of non-renewable resources like rare minerals, while Primary Energy measures the total energy consumed from extraction to use (Berthelot et al., 2024).

Researchers are developing new tools to track and report these metrics, such as ML CO<sub>2</sub> Impact Calculators and energy dashboards (Argerich and Patiño-Martínez, 2024). However, one of the biggest challenges is the lack of standardization. Currently, different studies and companies use different methods to calculate their carbon footprint. Some studies focus only on the electricity used during training, while others include the full LCA of hardware manufacturing and data centers. This inconsistency makes it difficult and almost impossible to fairly compare the environmental impact of different AI models or services. This lack of clear comparison prevents policymakers and developers from making smart, informed decisions about how to build sustainable AI. This indicates that a unified and standard way of measuring and reporting these impacts is required to ensure transparency. This will guide the entire AI industry toward truly sustainable



practices. This transition toward transparency is reinforced by the development of 'Green AI' techniques; recent systematic reviews indicate that algorithmic optimizations, such as model compression, knowledge distillation, and low-precision computing (quantization), can reduce AI energy consumption by up to 50% without significant loss in performance (Khan et al., 2026).

## 4.2 Environmental impacts of quantum computing

The promise of QC lies in its ability to solve problems intractable for classical computers, yet this potential comes with a unique and substantial environmental footprint (Theophilo and Oughton, 2025). Unlike GenAI, whose sustainability challenge is dominated by operational scale and ubiquity, QC's challenge is one of concentrated, hardware-driven intensity and long-term uncertainty. This makes early governance, lifecycle-aware design, and threshold-based deployment decisions critical for preventing QC from repeating the unsustainable trajectory of classical high-performance computing. This section examines the environmental impacts of QC, including its energy-intensive infrastructure and the materials required for its components.

### 4.2.1 Energy requirements of quantum hardware

Quantum computers, particularly those based on superconducting qubits, have unique and substantial energy requirements. While a quantum computer may be more efficient for specific, complex problems, its operational energy needs are driven by two main components: cryogenic cooling and control electronics (Giroux et al., 2022). A quantum computer's primary energy demands stem not from the processing itself, but from its supporting infrastructure. The core of a quantum computer's energy consumption comes from two main sources. First, the cryogenic cooling system is a massive energy hog,

as it must maintain the qubits at temperatures just a few millikelvin above absolute zero to preserve their fragile quantum states. This is achieved using large cryostats and complex refrigeration equipment. Second, the control electronics consume a significant amount of power. These systems are responsible for sending precise signals to the qubits to manipulate and read their states. As quantum computers scale up with more qubits, the energy required for both the cooling and control systems will only increase, making them critical areas for future energy efficiency improvements.

Cryogenic cooling is essential to maintain the ultralow temperatures (near absolute zero) required for quantum states to be stable. The cooling systems, often using liquid helium or dilution refrigerators, consume a significant amount of power, with systems requiring tens of kilowatts to cool just a few dozen qubits (Li et al., 2023). The continuous energy overhead of these systems is a major challenge for scalability, as cooling larger systems will require even more power (Li et al., 2023). In addition to cooling, the control electronics and signal processing hardware that prepare and measure the qubits consume continuous power (Theophilo and Oughton, 2025). De Michielis and Ferraro (2025) estimate that when using rotated surface code QEC with silicon flip-flop qubits, the power per logical qubit and the cryogenic and thermal control overheads grows significantly with code distance and noise level, meaning large, fault-tolerant quantum computers could require large aggregate power (though not yet in the speculative 1 GW-range used in some preprints). To put that in perspective, even orders of magnitude lower demands will still be comparable to supercomputing centers when scaled. This highlights that a "cryptanalytically relevant" quantum computer, powerful enough to break modern encryption, will be a major power consumer, comparable to large-scale computing facilities. This is primarily due to the energy-intensive support systems required, such as the massive cryogenic refrigerators needed to cool the qubits to near absolute zero and the complex control electronics. The study makes it clear that while

quantum computers promise an exponential leap in computational power, this leap will come with a very significant energy cost that must be addressed for the technology to be truly scalable and sustainable. Comparisons with classical supercomputing offer crucial insights (Bashir et al., 2024). A LCA showed that while a quantum computer can achieve an environmental advantage for certain problems due to shorter runtime, the substantial carbon footprint of its hardware, particularly for quantum error correction (QEC), is a significant factor (Cordier et al., 2025; Giroux et al., 2022). This is a key trade-off: a quantum computer may solve a problem faster, but the energy required to do so might still be considerable (Desdentado et al., 2024; Theophilo and Oughton, 2025). Research is ongoing to determine the “green quantum advantage threshold,” the point at which a quantum computer becomes more energy-efficient than a classical one for a specific task (Desdentado et al., 2024; Jaschke and Montangero, 2023).

Figure 9 compares the energy use of traditional computers with that of QCs as the problem complexity increases. The graph plots two lines. It plots the energy used by a classical computer (blue curve) and a QC (purple curve) as the problem complexity increases. The plot shows that the energy consumption of a classical computer increases extremely quickly as the problem becomes more complex. It also shows that the energy consumption of a QC increases polynomially (or linearly in some contexts) while the classical energy increases exponentially. The most important feature is the red dot where the two lines cross. This point is called the Green Quantum Advantage Threshold. Before the dot, the classical computer was more energy efficient. However, beyond the dot, the quantum computer is the clear winner. For any problem harder than this threshold, using the QC not only solves it faster but also uses significantly less energy. This information shows that not all complex tasks require the use of QC; hence, developers should make informed and responsible decisions. QC

should only be used to handle difficult tasks when their complexity exceeds the Green Quantum Advantage Threshold. Doing this will ensure that the use of a QC will provide both environmental benefit and performance advantage.

#### 4.2.2 Material and manufacturing concerns

The manufacturing of quantum computers relies on a range of scarce, critical, and often toxic materials, creating significant challenges for supply-chain sustainability and future electronic waste management (Arora and Kumar, 2026). Key components such as superconducting qubits depend on elements including niobium and ytterbium (Cordier et al., 2025; Giroux et al., 2022), whose mining and processing are energy-intensive and can cause habitat destruction and environmental pollution (Cordier et al., 2025). Additional risks arise from the use of quantum dots containing heavy metals such as cadmium, which pose serious threats to terrestrial and aquatic ecosystems if released into the environment (Zhang et al., 2024). These material-related impacts highlight the importance of assessing quantum hardware through comprehensive LCA frameworks (Cordier et al., 2025). At present, manufacturing constitutes the largest environmental contributor in the quantum computing lifecycle, exceeding operational impacts due to the energy demands of fabricating highly specialized components in clean-room environments using precise chemical processes (Cordier et al., 2025; Giroux et al., 2022). In response, researchers have proposed proactive governance and design frameworks, including carbon-aware quantum computing guidelines that promote responsible material use and the early development of sustainable supply chains (Arora and Kumar, 2026). Complementing these governance efforts, recent advances in the green synthesis of inorganic halide perovskite quantum dots (IHPQDs) demonstrate that aqueous, room-temperature methods

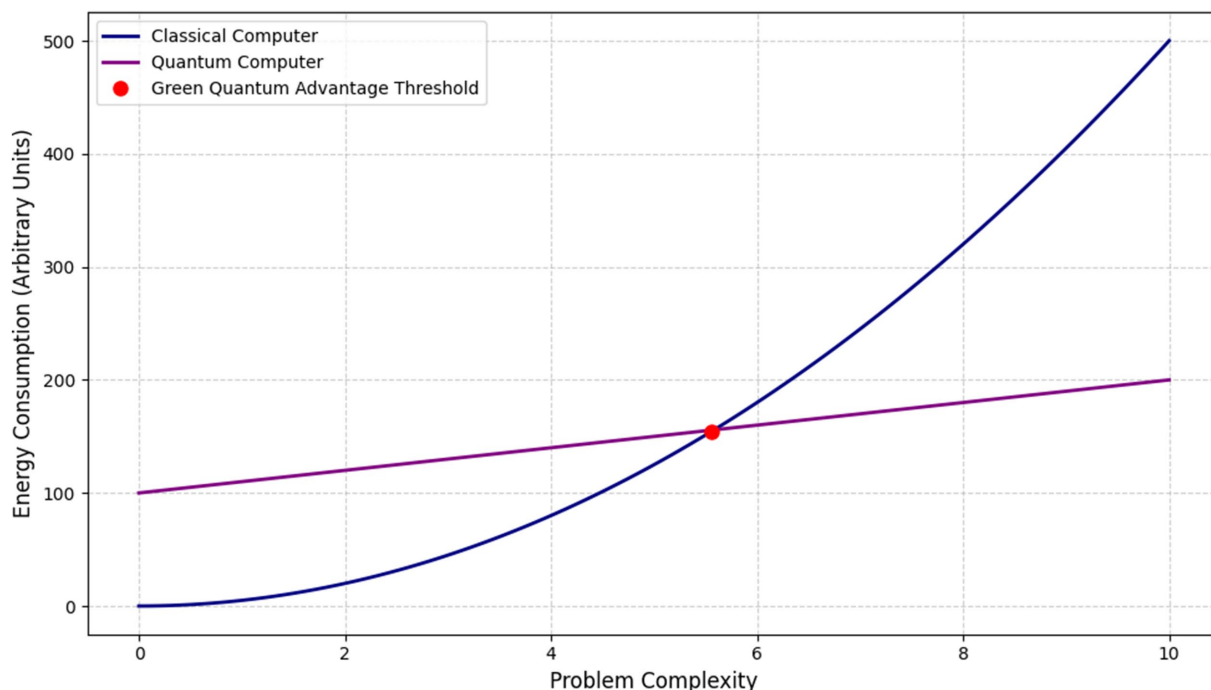


FIGURE 9

Green quantum advantage threshold. Source: synthesis based on the energy efficiency comparison models described by Jaschke and Montangero (2023) and Desdentado et al. (2024).

can substantially reduce environmental burdens, cutting hazardous solvent waste by up to 50% and lowering manufacturing energy consumption by up to 90% compared to conventional organic synthesis techniques (Abushuhel et al., 2026). These developments indicate that material innovation and process redesign can play a central role in reducing the environmental footprint of future quantum hardware.

### 4.2.3 Projected environmental trade-offs

The environmental narrative of QC is not solely one of risk. QC also holds immense promise as a powerful tool for sustainability, particularly in its ability to solve complex optimization problems that could lead to significant emissions reductions (Hassan and Hassan, 2025). Quantum for sustainability applications include:

- 1 Optimizing logistics and supply chains to reduce fuel consumption and emissions (Hassan and Hassan, 2025).
- 2 Energy modeling for more efficient power grids.
- 3 Simulating new materials for more efficient batteries, solar cells, and carbon capture technologies (Stein, 2024).

However, the potential sustainability benefits of QC are balanced against the risks of technology's own footprint. The high cooling overhead at scale and the energy demands of quantum error correction could negate the energy savings from these applications (Giroux et al., 2022). The environmental impact of classical simulation of quantum circuits is also a significant, often-overlooked, factor, with a large simulation potentially having a carbon footprint up to 48 times greater than training a large machine learning model (Li et al., 2023). There is still significant uncertainty surrounding future projections, as much of the data is based on theoretical models rather than empirical data from fully fault-tolerant systems (Desdentado et al., 2024). This underscores the need for more detailed, empirical research to understand the full environmental trade-offs of quantum technologies (Giroux et al., 2022). Ultimately, whether QC achieves a net-positive environmental impact depends on crossing the 'Green Quantum Advantage Threshold,' where the industrial-scale carbon savings from quantum applications outweigh the inherent cryogenic and error-correction overhead (Jaschke and Montangero, 2023). Ensuring this development is guided by principles of responsible innovation (Root, 2025) will be vital for aligning its powerful potential with global climate and sustainability goals.

## 4.3 Comparative sustainability frameworks and insights

Evaluating the environmental impact of GenAI and QC requires a structured approach that transcends a simple comparison of energy consumption. This section presents the key sustainability frameworks used for each technology and provides a comparative analysis, highlighting their unique challenges and potential synergies.

### 4.3.1 Frameworks for assessment

The idea of making computers and digital technology environmentally friendly (sustainable computing) is becoming a popular field of study. To evaluate how much new technology harms the environment, researchers use several standard methods. The main method used to evaluate both GenAI and QC is called LCA. For GenAI, LCA

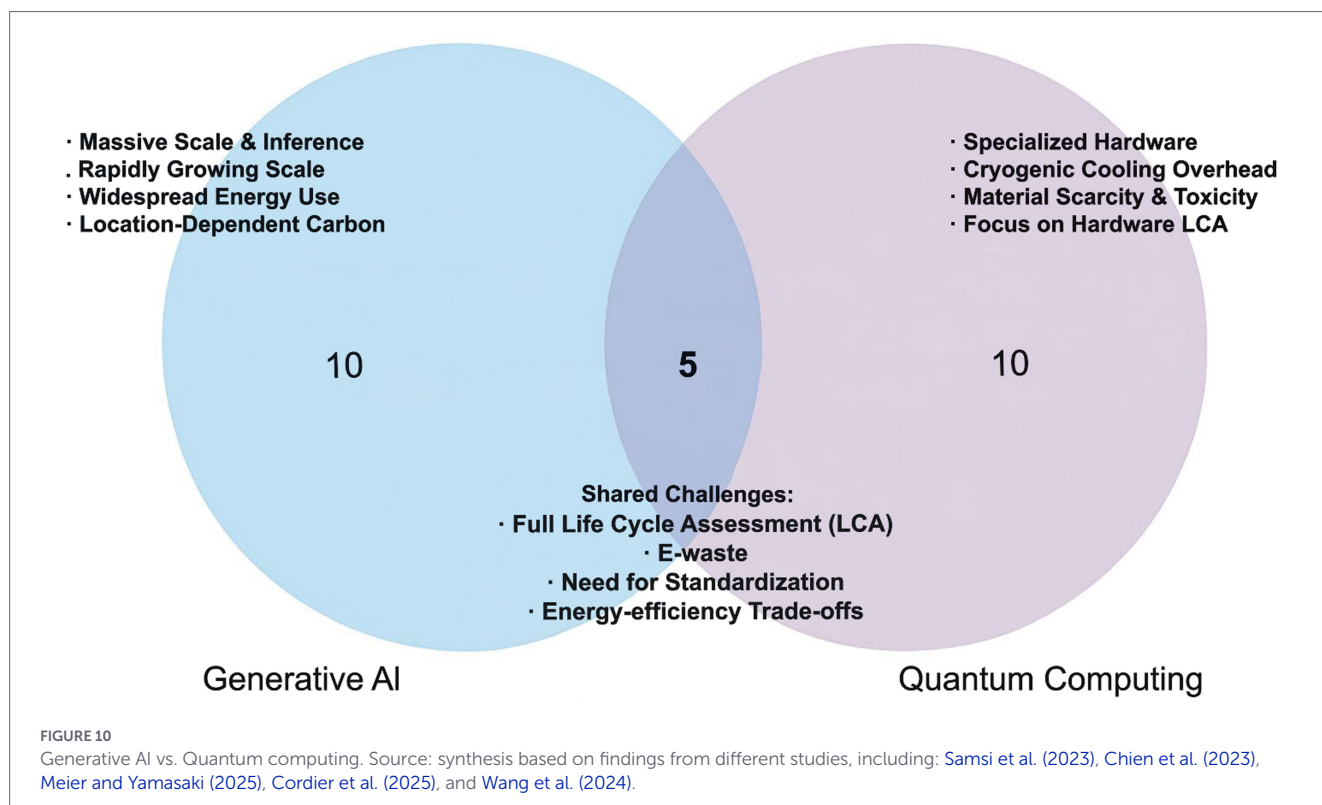
is used to evaluate the entire environmental cost, not only the energy used when training or running the AI. This includes the impact of creating the hardware and disposing of it when it is old. For example, Berthelot et al. (2024) used LCA to measure the carbon emissions, the depletion of natural resources (Abiotic Depletion Potential), and total energy use (Primary Energy). Another study (Li et al., 2025) used LCA to compare the pollution caused by running AI in a large data center (cloud) versus on a local device (edge). This also includes the pollution incurred from manufacturing the device itself. Understanding the full scope and impact of GenAI is important. LCA is also used to assess the environmental impact of QC. One study (Cordier et al., 2025) used LCA to compare a superconducting quantum computer to a traditional supercomputer. The comparison revealed that creating the quantum hardware in the first place creates a lot of damage to the environment. This concern is also seen in studies that worked on quantum dots. Giroux et al. (2022) pointed out that quantum dots pose risks to the environment because they contain toxic materials that can be released when they are discarded. While LCA is foundational, other frameworks are emerging. For AI, one study (Lee et al., 2025) uses Environmental, Social, and Governance (ESG) principles to create a comprehensive evaluation of how the AI used by an organization impacts the world and manages related risks. For QC, some researchers have proposed sustainability-oriented approaches such as Carbon-aware Quantum Computing frameworks (Lukyanenko, 2025) and approaches that embed responsibility and ethics into technological development, emphasizing sustainability and proactive governance from the earliest stages.

### 4.3.2 Generative AI vs. quantum computing: a comparative analysis

While both GenAI and QC have significant environmental footprints, the nature of their impacts differs substantially. As summarized in the Venn Diagram shown in Figure 10, each technology has unique sustainability challenges while also sharing key concerns.

The Venn diagram visually separates the unique sustainability challenges of GenAI and QC while highlighting their shared concerns. The left circle, representing GenAI, focuses on operational impacts at a massive scale. Its challenges, like "Massive Scale & Inference" and "Rapidly Growing Scale," stem from the immense and widespread energy consumption required to train and run its models. The carbon footprint of this energy is location-dependent, tied to the specific power grid of the data centers. The number 10 in this section indicates that 10 distinct, unique challenges were identified for GenAI within this framework. Conversely, the right circle for QC details challenges centered on its specialized, hardware-intensive nature. Its issues, such as "Specialized Hardware" and "Cryogenic Cooling Overhead," are related to the high energy needed for cooling systems and the sourcing of rare materials. The number 10 here signifies that those 10 unique challenges were identified for QC. The overlapping section in the middle outlines the 5 shared challenges, including the need for a Full LCA of all hardware, the growing problem of e-waste, and the lack of standardization for measuring their environmental impact.

GenAI's environmental footprint is characterized by its massive, immediate, and rapidly growing scale. The impact is primarily driven by the operational energy consumption of data centers, both for training and, more critically, for billions of daily inference queries (Argerich and Patiño-Martínez, 2024; Li et al., 2025). This creates a short-term environmental challenge defined by the need for massive computational power



and the accompanying carbon emissions from power grids. The sustainability of GenAI is heavily influenced by the electricity mix of the regions where data centers are located ([Hosseini et al., 2025](#)). In contrast, QC's environmental impact is currently defined by its specialized hardware and infrastructure ([Meier and Yamasaki, 2025](#)). The energy-intensive nature of cryogenic cooling and the reliance on rare, critical materials for fabrication are the primary concerns ([Giroux et al., 2022](#); [Li et al., 2023](#)). While the operational energy needs are high per unit, the number of operational quantum computers is still very small, making their overall environmental footprint less in the short term compared to the widespread deployment of GenAI models. However, at scale, the energy consumption of large, fault-tolerant quantum systems could become a dominant factor ([Giroux et al., 2022](#); [Li et al., 2023](#)). The sustainability trade-off between GenAI and QC is therefore a function of scale versus specialization. GenAI's problem is one of massive, distributed energy and resource consumption, while QC's problem is centered on a few highly energy-intensive and material-intensive systems. As these technologies converge into the field of Quantum Artificial Intelligence (QAI), a dual-benefit emerges: while quantum phenomena can make AI algorithms exponentially more efficient, classical AI approaches are simultaneously being used to optimize quantum circuit design and reduce hardware noise, further lowering the cumulative energy overhead of the next generation of computing ([Acampora et al., 2026](#)).

### 4.3.3 Interplay and synergies

GenAI and QC are separate technologies, each with its own set of challenges. They are both powerful tools that can be used together to make things greener and to create powerful new ways of tackling sustainability. Although GenAI has its own energy demands, it can be used to create brilliant designs that will make almost any system more energy-efficient. A good example is architecture. GenAI can be used

to quickly create and test thousands of designs at the same time, instead of an architect manually sketching and testing a few building designs. GenAI can rapidly find the best design to reduce a building's carbon footprint, optimizing for things like the materials used and how much energy the finished building will consume. However, the potential of GenAI goes far beyond construction. In finance, GenAI can be used by banks and investors to integrate the ESG (Environmental, Social, and Governance) factors in a better way. For instance, an AI can be used to analyze huge amounts of market data to find the smartest investment strategies that help a company reach its net-zero emissions goals. GenAI can also be used to create new fintech tools to make trading carbon credits easier. Moreover, GenAI is a powerful tool for resource management. GenAI can be used to analyze complex supply chains to suggest shorter, less-polluting delivery routes. GenAI can also be used to run simulations to show companies how to minimize waste and reduce energy use. AI is a creative problem-solver that can make other industries and systems much more sustainable. QC contributes to this effort by being able to solve extremely complex optimization problems that are simply too big for standard computers ([Stein, 2024](#)). Quantum algorithms are suitable for tasks like optimizing massive energy grids, perfecting logistics routes to reduce waste and emissions ([Hassan and Hassan, 2025](#)). These optimization efforts are particularly critical for the evolution of smart grids; the implementation of Wide Area Monitoring Systems (WAMS) enhanced by AI-driven observability can stabilize grids against the intermittency of renewable sources, potentially reducing technical losses and fossil fuel reliance by over 20% while managing the intense energy demands of large-scale data centers ([Al Muttaki et al., 2026](#)). QC can also be used to design new sustainable materials. Solar cells are typical examples. As illustrated in [Figure 11](#), the efficiency of a solar cell depends on its ability to receive light across the entire spectrum, from ultraviolet to infrared. A typical silicon cell only does a great job in the middle (visible light).

QC can perfectly simulate the way photons (light particles) interact with the electrons of a material. By modeling this quantum behavior in detail, scientists can use QC to design new materials with an ideal band gap, enabling solar cells to absorb a much wider range of sunlight and convert it into electricity far more effectively. This shows how QC can fast-track the discovery of breakthrough energy solutions.

The main unanswered question is the combined environmental cost of using these two high-tech technologies. This is an exciting future research possibility called quantum-accelerated AI, where quantum computers are used to supercharge AI's machine learning processes. This could be a huge net positive if the massive efficiency gains outweigh the energy cost of running the quantum hardware (Alzami, 2026). There is a need for a deeper understanding of this combined footprint, especially since even the research simulations for quantum circuits are energy-intensive (Li et al., 2023). To ensure a truly sustainable future, we must proactively develop and deploy both GenAI and QC in a unified way that keeps the global climate goals front and center.

#### 4.4 Dual perspectives on sustainability: green in IT and green by IT frameworks

This section examines the 'Green in IT' paradigm, focusing on internal strategies to enhance the inherent sustainability of computing architectures by optimizing the energy efficiency and resource management of both GenAI and QC models. Beyond internal optimization, this section explores the 'Green by IT' framework, illustrating how the advanced computational power of GenAI and QC can be harnessed as external tools to solve complex climate challenges and drive global decarbonization.

##### 4.4.1 Green in IT frameworks

The rapid scaling of GenAI has necessitated a "Green in IT" focus on reducing the computational intensity and resource

consumption of the models themselves. Current research emphasizes "efficiency-centered business model innovation" and the development of energy-aware training protocols to mitigate the staggering electricity and water demands of LLMs. By implementing techniques such as model pruning, quantization, and real-time energy monitoring, developers are striving to decouple GenAI performance from its carbon footprint. This approach is increasingly supported by regulatory alignment, which incentivizes firms to prioritize sustainable manufacturing and operational efficiency throughout the GenAI life cycle (Moravec et al., 2025). Recent peer-reviewed work also illuminates how environmental impacts vary significantly with data center architecture and operational choices. For example, a 2026 LCA of multiple data center typologies shows that optimized high-density, liquid-cooled designs and careful model selection can lower emissions per unit of computational output, highlighting energy-efficiency as a critical lever in sustainable GenAI deployment (d'Orgeval et al., 2026).

For quantum technologies, "Green in IT" involves optimizing the hardware and algorithmic efficiency to minimize the massive energy overhead required for cryogenic cooling and error correction. Since maintaining fragile quantum states at near-absolute zero is the primary driver of its environmental impact, researchers are exploring "quantum green computing" through energy-efficient hardware architectures and hybrid quantum-classical frameworks. Innovative proposals, such as hosting quantum processors in the naturally cold stratosphere to reduce active cooling needs, demonstrate a shift toward structurally sustainable quantum design (Abderrahim et al., 2025). In addition, recent work on hybrid quantum-classical models for cloud systems illustrates how integrating quantum computation can significantly reduce computational overhead and associated carbon emissions in real-world infrastructure by improving predictive efficiency and dynamic resource provisioning compared with classical

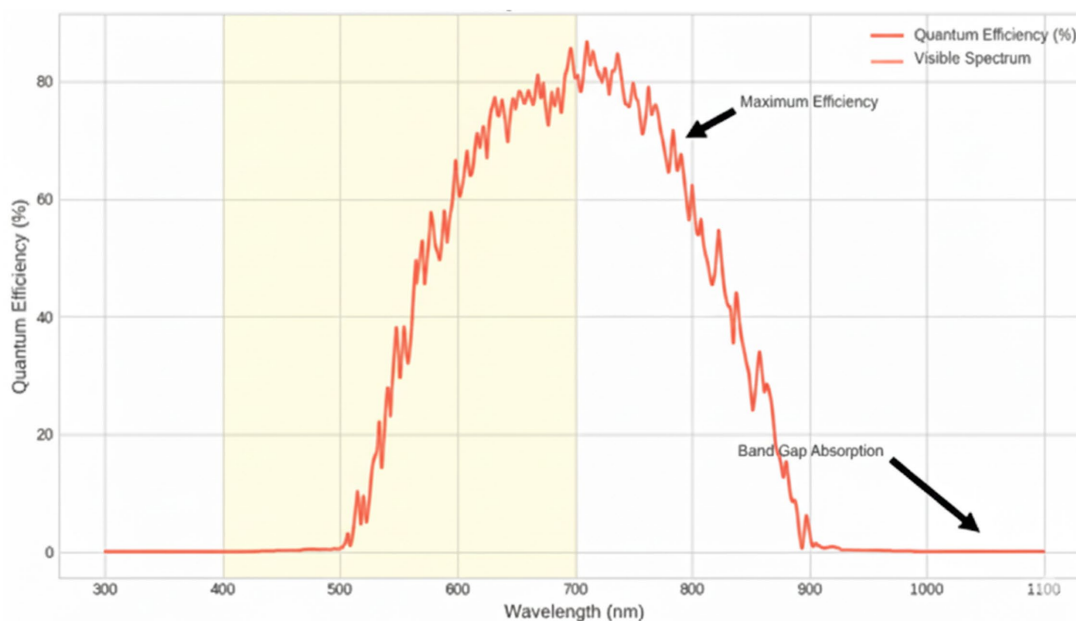


FIGURE 11

Quantum efficiency of a solar cell. Source: conceptual illustration of quantum efficiency principles; application concept based on Hassan and Hassan (2025).

approaches alone (Tripathi et al., 2025). These efforts aim to ensure that as quantum systems scale to 100 logical qubits and beyond, their physical footprint and operational energy demand do not expand proportionally with their computational power (Abderrahim et al., 2025; Tripathi et al., 2025).

#### 4.4.2 Green by IT frameworks

Green by IT leverages the predictive and generative capabilities of AI to optimize external systems, such as smart grids and urban infrastructure, for maximum sustainability. Generative models are currently being deployed to enhance renewable energy forecasting, demand-response strategies, and the design of energy-efficient buildings by simulating complex multi-objective scenarios that traditional optimization methods overlook. Recent research shows that GenAI methods, including GANs, VAEs, and transformer-based architectures, can significantly improve solar and wind forecasting accuracy. Specifically, GAN-based models have been reported to reduce root-mean-square error (RMSE) by 15–20% in solar irradiance prediction compared to conventional methods. These technologies also enhance energy storage management and dispatch planning in high-penetration renewable systems by modeling complex battery dynamics and optimizing multi-vector energy flows. Such advancements demonstrate gains in operational flexibility and system resilience that are not achievable with conventional, deterministic models (Mamat et al., 2025). GenAI thus supports more adaptive, robust grid planning, helping decision-makers manage variability and uncertainty in renewable integration. In the context of the energy transition, GenAI acts as a decision-support tool for grid operators, enabling the seamless integration of intermittent renewable sources like wind and solar. This application effectively ‘pays back’ the model’s training footprint by facilitating significant operational efficiency gains and CO<sub>2</sub> abatement across the broader economy (Mamat et al., 2025). For quantum technologies, ‘Green by IT’ focuses on the technology’s role as a computational catalyst for global decarbonization. Recent peer-reviewed analysis highlights how quantum and hybrid computational methods can enhance real-time smart grid optimization and renewable integration, improving resilience and operational efficiency while reducing energy loss and carbon emissions in next-generation power systems (Rouhani and Su, 2026). By solving optimization problems previously considered intractable, such as large-scale carbon capture simulation or hydrogen fuel cell catalysis, QC provides the ‘Green by IT’ utility necessary to offset its initial hardware-centric environmental costs.

### 4.5 Mitigation strategies and sustainable pathways

Addressing the environmental challenges of GenAI and QC requires different approaches that combine different interventions, including technological, architectural, and policy-level interventions. This section outlines some strategies that can be used to mitigate the environmental challenges of GenAI and QC.

#### 4.5.1 Mitigation strategies for generative AI

To reduce the huge amount of energy GenAI uses, the problem must be tackled from two directions. There is a need to target both the AI model and the hardware (the infrastructure) that runs it.

##### 4.5.1.1 Algorithmic and architectural strategies

Sustainable AI development relies on intelligent techniques to design and deploy models that will require reduced computing power and get the most out of the machines they run on. Some key strategies include:

- 1 Model optimization: Pruning (trimming away the parts of the model that does not do much), distillation (training a small, efficient “student” model to learn from an experienced “teacher”), and quantization (using lower-precision numbers, like rounding decimals, to speed things up) are key techniques for creating smaller and more energy-efficient AI models, and empirical evidence shows that applying these methods to transformer-based models substantially reduces energy consumption and carbon emissions without greatly compromising performance (Paula et al., 2025). Even simple things we do as AI users, such as writing a short prompt or choosing an AI model that is just the right size for a specific task, can help save energy every time the AI generates a response (Ren et al., 2024).
- 2 Efficient architectures: Instead of depending on one gigantic, all-knowing language model (known as a monolithic LLM) that needs constant, expensive retraining, researchers are moving toward modular designs. Approaches like Retrieval-Augmented Generation (RAG) allow an AI model to search for specific, relevant information from an external, current database to answer a question (Nguyen et al., 2024). This is much more targeted and energy-efficient compared to constantly rebuilding the entire foundational model.
- 3 Edge computing: Shifting AI workloads from centralized cloud platforms to edge devices offers a significant environmental advantage. One study found that edge platforms can achieve over 90% energy savings and reduce carbon emissions by more than 80% compared to cloud platforms, while also enhancing user privacy (Li et al., 2025).

##### 4.5.1.2 Infrastructure strategies

Beyond algorithmic improvements, significant reductions in the environmental footprint of AI rely on implementing sustainable practices within the physical infrastructure that hosts these models. These infrastructure-level strategies focus on the data center environment and the hardware itself:

- 1 Green data centers: The high energy demands of data centers can be mitigated by transitioning to renewable-powered energy sources (Hosseini et al., 2025). This can be achieved through utility partnerships, direct purchasing of green energy, or deploying facilities in regions with a low-carbon electricity grid (Berthelot et al., 2025).
- 2 Cooling and hardware: Adopting more efficient cooling methods, such as liquid cooling, can reduce the substantial water and energy consumption associated with traditional air-cooling systems (d’Orgeval et al., 2026). Furthermore, designing hardware for greater efficiency is crucial, as GPU-based computing can require up to 15 times more energy than CPU-based computing (Samsi et al., 2023).

## 4.5.2 Mitigation strategies for quantum computing

Mitigation strategies for QC are focused on its unique challenges related to extreme hardware requirements and speculative future scaling.

- 1 **Cryogenics and hardware improvements:** A key area of focus is on improving the energy efficiency of the cryogenic systems that keep qubits in a stable state. [De Michielis and Ferraro \(2025\)](#) provide estimates on how the power per logical qubit and the overhead of cryogenic and thermal control scale with noise levels and code distance, thereby giving realistic benchmarks for what hardware improvements are needed. The development of room-temperature qubits is a long-term goal that, if achieved, could drastically reduce the need for energy-intensive cooling systems ([Desdentado et al., 2024](#)).
- 2 **Sustainable material sourcing:** Given the reliance on rare and critical materials, frameworks are needed to ensure responsible material sourcing and supply chains ([Giroux et al., 2022](#)). Researchers are also investigating cadmium-free alternatives for quantum dots (QDs) to mitigate environmental risks from toxic metals ([Giroux et al., 2022](#); [Zhang et al., 2024](#)).
- 3 **Green quantum algorithms:** Beyond hardware, the design of more energy-efficient algorithms, called “green quantum algorithms” and “green software,” can help reduce the energy required to solve specific problems ([Theophilo and Oughton, 2025](#)). Research on the “green quantum advantage threshold” aims to determine the point at which a quantum computer becomes more energy-efficient than a classical one for a specific task, guiding future software and hardware development ([Desdentado et al., 2024](#)).

## 4.5.3 Cross-cutting opportunities

The path to sustainable computing requires collaboration and shared frameworks that transcend the boundaries between GenAI and QC.

- 1 **Shared policies and metrics:** The lack of standardized carbon and energy reporting is a major obstacle for both fields ([Bashir et al., 2024](#); [Berthelot et al., 2024](#)). Establishing unified metrics and encouraging open carbon reporting would allow for more transparent comparisons and incentivize sustainable practices ([Berthelot et al., 2024](#)).
- 2 **Interdisciplinary collaboration:** Researchers from both domains can collaborate on a combined sustainable innovation framework that considers the interplay between the technologies ([Root, 2025](#); [Stein, 2024](#)). For example, quantum-accelerated AI could utilize quantum computers to optimize and accelerate AI algorithms, resulting in a net reduction in energy consumption if the gains outweigh the quantum hardware’s footprint ([Alzami, 2026](#)).
- 3 **Leveraging technology for sustainability:** Both GenAI and QC have the potential to be powerful tools for sustainability in other fields ([Columbus Chinnappan et al., 2025](#); [Stein, 2024](#)). They can be applied to complex problems in climate modeling, logistics, and materials science, helping to drive broader emissions reductions ([Hassan and Hassan, 2025](#); [Vudugula et al.,](#)

[2025](#)). This underscores a crucial paradigm shift: rather than just focusing on the internal footprint of these technologies, we must also focus on their external, positive impact on a global scale.

## 4.6 Open challenges and future directions

Despite a growing body of research, the assessment of the environmental impacts of GenAI and QC remains a nascent field with significant open challenges. Overcoming these hurdles is essential for guiding these technologies toward a sustainable future. This section synthesizes the key research gaps and proposes a roadmap for future work.

### 4.6.1 Cross-cutting opportunities

A primary challenge across both domains is the lack of standardized metrics and unified reporting frameworks ([Bashir et al., 2024](#); [Berthelot et al., 2024](#)). Studies on GenAI often use different methodologies, making it difficult to compare the carbon footprint of models like GPT-3, PaLM, and LLaMA directly ([Berthelot et al., 2025](#)). Similarly, in QC, the nascent stage of the technology means that data is often based on theoretical models rather than empirical measurements, which limits the robustness of comparisons with classical supercomputing ([Desdentado et al., 2024](#)). A concerted effort is needed to develop and adopt common metrics for sustainability assessment. These should go beyond simple energy consumption to include Abiotic Depletion Potential (ADP), water usage, and embodied carbon from hardware manufacturing ([Berthelot et al., 2024](#); [Cordier et al., 2025](#)). Greater transparency and public reporting from manufacturers and service providers are crucial to enable more accurate, end-to-end assessments. This transparency is a foundational requirement for all future research and policy-making.

### 4.6.2 Key trade-offs

Future research must also navigate the complex trade-offs inherent in designing sustainable computing systems. A primary challenge is the tension between accuracy and efficiency: while models can be optimized for energy efficiency through techniques like pruning and quantization, this often involves a perceived, and sometimes real, trade-off with performance. The core research challenge lies in finding the optimal balance where significant energy savings are achieved with minimal or negligible loss in accuracy, as demonstrated by approaches such as the AVATAR methodology and by deployment-level and user-interaction strategies that reduce energy use through model choice and prompt/response management ([Ren et al., 2024](#)).

Furthermore, the pursuit of energy savings introduces a trade-off between efficiency and accessibility. Although shifting AI workloads to edge devices offers clear environmental benefits, this approach may introduce new hurdles related to hardware heterogeneity and performance sacrifices, such as increased latency, which have been noted as needing further study ([Li et al., 2025](#)). Finally, a critical concern specific to QC is the balance between performance and sustainability. The pursuit of fault tolerance through Quantum Error Correction (QEC) will require a massive increase in hardware and energy consumption ([Jaschke and Montangero, 2023](#)). Consequently, a key research question for future hardware development is how the field can scale

quantum systems while effectively managing this significant environmental footprint (Li et al., 2023).

#### 4.6.3 Policy and regulatory needs

The rapid advancement of these technologies has outpaced the development of effective governance and policy frameworks. A proactive approach is needed to prevent negative environmental consequences, a lesson learned from the semiconductor revolution.

- 1 Mandatory reporting: Governments and international bodies should consider policies that mandate carbon and resource consumption reporting for large-scale computational services (Christensen et al., 2021). This would incentivize sustainable practices and provide the data needed for more accurate assessments (Berthelot et al., 2024).
- 2 Sustainable innovation frameworks: Research is needed to develop and apply frameworks that integrate sustainability into the earliest stages of technological development (Root, 2025). This includes establishing guidelines for responsible material sourcing, waste management, and the ethical use of resources throughout the supply chain (Wang et al., 2024).

#### 4.6.4 Research opportunities

Several key research areas are ripe for exploration to address the identified challenges.

- 1 Empirical studies on energy consumption: While some studies have focused on the energy costs of training, there is a significant need for more empirical data on the energy consumption of GenAI inference at scale, as inference is a far more frequent operation (Argerich and Patiño-Martínez, 2024; Li et al., 2025).
- 2 Energy-aware algorithms: Future work should focus on designing and testing algorithms for both classical and QC that are inherently energy-efficient (Theophilo and Oughton, 2025). In QC, research is needed to empirically validate the “green quantum advantage threshold” and to develop algorithms that minimize energy consumption in hybrid quantum-classical workflows (Desdentado et al., 2024).
- 3 Cross-disciplinary sustainability studies: There is an opportunity to analyze the combined environmental footprint of the intersection of these two fields, such as in quantum-accelerated AI (Alzami, 2026). This includes exploring the hidden costs, such as the energy consumed by the classical simulation of quantum circuits (Hidayat and Surendro, 2025).
- 4 Improved tech applications: Research should continue to explore how both technologies can be used as tools to accelerate global sustainability efforts, from optimizing supply chains to modeling climate systems (Columbus Chinnappan et al., 2025; Hassan and Hassan, 2025; Stein, 2024).
- 5 Lifecycle assessment (LCA) of hardware: There is a need to conduct comprehensive lifecycle assessments of hardware used in GenAI (e.g., GPUs, TPUs) and quantum computers. This includes evaluating the environmental impacts of raw material extraction, manufacturing, deployment, and disposal, which are often overlooked compared to operational energy costs (Ma and Zhou, 2024).

- 6 Sustainable data management for AI training: Training generative models requires massive datasets, yet little research has examined sustainable data management strategies, such as low-carbon data centers, green storage technologies, and federated data-sharing frameworks that reduce redundancy and carbon emissions (Dibsi and Cho, 2023; Oloruntoba, 2024).
- 7 Integration with renewable energy systems: Both GenAI and QC infrastructures could be strategically integrated with renewable energy sources. Research opportunities include developing scheduling frameworks that align computationally intensive tasks with peak (Crozier and Liska, 2025; Guo et al., 2025).
- 8 Standardized sustainability metrics and benchmarks: There is currently no universally accepted framework for measuring the environmental footprint of GenAI and QC. Research is needed to design standardized benchmarks and sustainability metrics (e.g., carbon intensity per model parameter or per quantum gate operation) that allow for fair comparisons across systems and approaches (Bolón-Canedo et al., 2024; Hasan et al., 2025).

By tackling these open challenges, researchers, policymakers, and industry stakeholders can ensure that the transformative power of GenAI and QC is harnessed in a way that aligns with, rather than undermines, global climate and sustainability goals.

## 5 Conclusions, limitations, and future prospects

The rapid evolution of GenAI and QC represents a significant technological shift, offering profound problem-solving potential alongside escalating ecological costs. This review synthesized contemporary research to evaluate the sustainability of these two paradigms, specifically addressing three core research questions. Regarding the first research question, the findings illustrate a clear distinction between the environmental profiles of these technologies. The ecological burden of GenAI is primarily driven by the extreme energy requirements of the initial development phase and the cumulative electricity consumption resulting from widespread daily use. In contrast, the environmental footprint of QC stems largely from the massive energy demands of supporting infrastructures, specifically the specialized systems required to maintain stable operating environments, and the reliance on rare materials that pose risks to natural ecosystems. In response to the second research question, this study identified several pathways for mitigation. Technical advancements, such as designing software that requires fewer computational resources and developing hardware that functions in less energy-intensive conditions, provide significant opportunities for reduction. However, these innovations must be supported by comprehensive policy frameworks. Such frameworks ensure that environmental considerations are integrated into the earliest stages of research and development, preventing ecological risks before technologies reach a global scale. Addressing the third research question, the review highlights that a lack of industry transparency remains a primary barrier to progress. The omission of overhead costs in energy reporting and the proprietary nature of manufacturing data hinder the establishment of standardized benchmarks. Without these benchmarks, it

remains difficult to conduct transparent comparisons between the environmental performance of classical and quantum computing systems. In conclusion, achieving sustainable computing requires a dual focus: managing the immense scale and distribution of AI systems while simultaneously addressing the hardware intensity of quantum technologies. Aligning these advancements with global climate goals requires a transition from isolated technical fixes to integrated strategies that prioritize environmental health as a fundamental concern.

## 5.1 Limitations

Despite the comprehensive nature of this review, some limitations should also be acknowledged. First, this analysis is limited to two technologies: GenAI and QC. The study does not address the environmental footprint of other computationally intensive technologies, such as blockchain/distributed ledger technologies (DLTs), general-purpose deep learning (e.g., classification models), or high-performance computing (HPC) for climate modeling. Second, getting precise and reliable numbers for energy use is difficult. Especially when it comes to the energy consumed by GenAI when everyone is using it (inference) and the costs of large-scale QC. The data is often inconsistent and frequently kept secret by companies that own the technology. This made it tricky to make direct comparisons; hence, this study had to rely on scientific modeling and educated estimates reported in previous studies.

## 5.2 Further prospects

If we want computing to have a sustainable future, there is an urgent need to completely change the way we think about it. Sustainability cannot be treated as simply ticking boxes; sustainability needs to be built into the very foundation of every new technology that is created. This implies that there is a need to focus our research on solving the difficult problems QC and GenAI present to us. There is also a need for better transparency by sharing data openly and using standardized reporting. Moreover, people from different fields must work together on these issues. By honestly tackling the environmental downsides while using these powerful tools to fight climate change, we can steer ourselves toward a future where computing is not just incredible and life-changing, but also genuinely sustainable. This vision, where green computing is a fundamental principle, is one of the few ways to ensure that technological progress aligns with a healthy and sustainable planet.

## Author contributions

EE: Formal analysis, Data curation, Resources, Validation, Methodology, Investigation, Writing – review & editing, Project administration, Funding acquisition, Writing – original draft, Conceptualization. AA: Data curation, Validation, Project

administration, Methodology, Conceptualization, Supervision, Writing – review & editing, Software, Resources, Investigation, Writing – original draft, Formal analysis, Visualization. MD: Investigation, Conceptualization, Funding acquisition, Writing – review & editing, Supervision.

## Funding

The author(s) declared that financial support was not received for this work and/or its publication.

## Acknowledgement

This work acknowledges the support of the Foundation for Science and Technology (FCT), within the framework of the UID/04292/2025, Marine and Environmental Sciences Centre (MARE).

## Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declared that Generative AI was used in the creation of this manuscript. Generative AI was used to refine the grammar, improve sentence clarity, and improve the overall readability of the manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## References

Abderrahim, W., Amin, O., and Shihada, B. (2025). Green quantum computing in the sky. *NPJ Wireless Technol.* 1:5. doi: 10.1038/s44459-025-00005-y

Abushuhel, M., Padmapriya, M. G., Shaker, A.-H., Ray, S., Chennakesavulu, K., Sharma, R., et al. (2026). Advancing inorganic halide perovskite quantum dots:

- breakthroughs in green synthesis, stabilization, and sustainability perspectives. *Chin. J. Anal. Chem.* 54:100679. doi: 10.1016/j.cjac.2025.100679
- Acampora, G., Chiatto, A., Schiattarella, R., and Vitiello, A. (2026). Quantum artificial intelligence: A survey. *Comput Sci Rev* 59:100807. doi: 10.1016/j.cosrev.2025.100807
- Al Muttaki, M. R., Anil, A. I. A., Afrin, S., and Hasan, S. (2026). Towards smarter grids: a systematic review of wide area monitoring enhancing stability, protection, and promoting environmental sustainability. *Energy Conversion Manage* 29:101424. doi: 10.1016/j.ecmx.2025.101424
- Alzami, I. (2026). Quantum-enhanced residual convolutional attention architecture for renewable forecasting in off-grid cloud microgrids. *Mathematics* 14, 1–37. doi: 10.3390/math14010181
- Argerich, M. F., and Patiño-Martínez, M. (2024). Measuring and improving the energy efficiency of large language models inference. *IEEE Access* 12, 80194–80207. doi: 10.1109/ACCESS.2024.3409745
- Arora, N., and Kumar, P. (2026). Sustainable quantum computing. *Commun. ACM* 69, 84–90. doi: 10.1145/3745782
- Bashir, N., Donti, P., Cuff, J., Sroka, S., Ilic, M., Sze, V., et al. (2024). “The climate and sustainability implications of generative AI” in An MIT exploration of generative AI, Cambridge, MA: MIT Press, 57–71.
- Berthelot, A., Caron, E., Jay, M., and Lefèvre, L. (2024). Estimating the environmental impact of generative-AI services using an LCA-based methodology. *Procedia CIRP* 122, 707–712. doi: 10.1016/j.procir.2024.01.098
- Berthelot, A., Caron, E., Jay, M., and Lefevre, L. (2025). Understanding the environmental impact of generative AI services. *Commun. ACM* 68, 46–53. doi: 10.1145/3725984
- Bolón-Canedo, V., Morán-Fernández, L., Cancela, B., and Alonso-Betanzos, A. (2024). A review of green artificial intelligence: towards a more sustainable future. *Neurocomputing* 599:128096. doi: 10.1016/j.neucom.2024.128096
- Boulais, I., Dauner, M., Hegenderfer, L., and Huang, I. (2026) Behind the prompt: The environmental impact of LLM inference. Available online at: [https://digitalcommons.calpoly.edu/cgi/viewcontent.cgi?filename=1&article=1117&context=ceng\\_surp&type=additional](https://digitalcommons.calpoly.edu/cgi/viewcontent.cgi?filename=1&article=1117&context=ceng_surp&type=additional)
- Chauhan, A. S., Novita, M., and Bhavna (2025). “Carbon footprint reduction in virtual worlds: innovations and implications” in *Metaverse and sustainability: Business resilience towards sustainable development goals* (Cham, Switzerland: Springer), 231–262.
- Chien, A. A., Lin, L., Nguyen, H., Rao, V., Sharma, T., and Wijayawardana, R. (2023). Reducing the carbon impact of generative AI inference (today and in 2035). In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*, 1–7.
- Christensen, H. B., Hail, L., and Leuz, C. (2021). Mandatory CSR and sustainability reporting: economic analysis and literature review. *Rev. Account. Stud.* 26, 1176–1248. doi: 10.1007/s11142-021-09609-5
- Columbus Chinnappan, C., Thanaraj Krishnan, P., Elamaran, E., Arul, R., and Sunil Kumar, T. (2025). Quantum computing: foundations, architecture and applications. *Eng. Rep.* 7:e70337. doi: 10.1002/eng.2.70337
- Cordier, S., Thibault, K., Arpin, M.-L., and Amor, B. (2025). Scaling up to problem sizes: an environmental life cycle assessment of quantum computing. *Quantum Sci. Technol.* 10:025058. doi: 10.1088/2058-9565/ad0c0a
- Crozier, C., and Liska, M. (2025). The potential of data center energy demand to provide grid flexibility. *Curr. Sustain. Renew. Energy Rep.* 12, 1–6. doi: 10.1007/s40518-025-00258-9
- d’Orgeval, A., Sheehan, S., Avenas, Q., Assoumou, E., and Sessa, V. (2026). Generative AI impact assessment through a life cycle analysis of multiple data center typologies. *Appl. Energy* 406:127288. doi: 10.1016/j.apenergy.2025.127288
- Dauner, M., and Socher, G. (2025). Energy costs of communicating with AI. *Front. Commun.* 10:1572947. doi: 10.3389/fcomm.2025.1572947
- De Michielis, M., and Ferraro, E. (2025). Energy and power scaling in quantum computers based on rotated surface codes with silicon flip-flop qubits. *EPJ Quantum Technol.* 12:64. doi: 10.1140/epjqt/s40507-025-00351-4
- Desdentado, E., Calero, C., Moraga, M. A., Serrano, M., and García, F. (2024). Exploring the trade-off between computational power and energy efficiency: an analysis of the evolution of quantum computing and its relation to classical computing. *J. Syst. Softw.* 217:112165. doi: 10.1016/j.jss.2024.112165
- Dibs, J., and Cho, J. (2023). Data management for environmentally sustainable and profitable business: evidence from the food retail industry. *Sustainability* 15:16298. doi: 10.3390/su152316298
- Ding, Z., Wang, J., Song, Y., Zheng, X., He, G., Chen, X., et al. (2025). Tracking the carbon footprint of global generative artificial intelligence. *Innovation* 6:100866. doi: 10.1016/j.xinn.2025.100866
- Dwivedi, K., Haghparast, M., and Mikkonen, T. (2024). Quantum software engineering and quantum software development lifecycle: a survey. *Clust. Comput.* 27, 7127–7145. doi: 10.1007/s10586-024-04362-1
- Gheibi, M., Canga, E., and Eldessouki, M. (2026). Life cycle assessment of nanowires for sustainable material production: a systematic review. *Results Eng.* 29:108300. doi: 10.1016/j.rineng.2025.108300
- Giroux, M. S., Zahra, Z., Salawu, O. A., Burgess, R. M., Ho, K. T., and Adeleye, A. S. (2022). Assessing the environmental effects related to quantum dot structure, function, synthesis and exposure. *Environ. Sci. Nano* 9, 867–910. doi: 10.1039/D1EN00712B
- Guo, H., Yu, H., Wang, M., Liu, C., and Li, C. (2025). Integrated management of workloads and energy system for data centers. *Energy* 327:136400. doi: 10.1016/j.energy.2025.136400
- Hasan, S. M., Islam, T., Saifuzzaman, M., Ahmed, K. R., Huang, C.-H., and Shahid, A. R. (2025). Carbon emission quantification of machine learning: a review. *IEEE Trans. Sustain. Comput.* 10, 1085–1102. doi: 10.1109/TUSC.2025.3578834
- Hassan, S., and Hassan, S. (2025). Quantum technology application in product development for sustainability. *Int. J. Sci. Res. Technol.* 2, 9–18.
- He, R., Cao, J., and Tan, T. (2025). Generative artificial intelligence: a historical perspective. *Natl. Sci. Rev.* 12:nwaf050. doi: 10.1093/nsr/nwaf050
- Hidayat, W., and Surendro, K. (2025). Quantum computing in addressing greenhouse gas emissions: a systematic literature review. *Egypt. Inform. J.* 29:100622. doi: 10.1016/j.eij.2025.100622
- Hinov, N. (2025). The energy hunger of AI: large language models as challenges and enablers for sustainable energy. *Energies* 18:4701. doi: 10.3390/en18174701
- Hosseini, M., Gao, P., and Vivas-Valencia, C. (2025). A social-environmental impact perspective of generative artificial intelligence. *Environ. Sci. Ecotechnol.* 23:100520. doi: 10.1016/j.esc.2024.100520
- Jamil, S., Jalil Piran, M., and Kwon, O.-J. (2023). A comprehensive survey of transformers for computer vision. *Drones* 7:287. doi: 10.3390/drones7050287
- Jaschke, D., and Montangero, S. (2023). Is quantum computing green? An estimate for an energy-efficiency quantum advantage. *Quantum Sci. Technol.* 8:025001. doi: 10.1088/2058-9565/acae3e
- Jeanquartier, F., Jean-Quartier, C., Rieder, P., Misirlic, V., Pasero, C., Hohensinner, R., et al. (2026). Assessing the carbon footprint of language models: towards sustainability in AI. *Resour. Conserv. Recycl.* 226:108670. doi: 10.1016/j.resconrec.2025.108670
- Ji, Z., and Jiang, M. (2026). A systematic review of electricity demand for large language models: evaluations, challenges, and solutions. *Renew. Sust. Energy Rev.* 225:116159. doi: 10.1016/j.rser.2025.116159
- Kakarla, S., Bathyal, S., Dara, M., and Ponneri, D. (2025). Transforming research with quantum computing: fundamentals, implementations, and applications. *IJSAT-Int. J. Sci. Technol.* 16, 1–8. doi: 10.71097/IJSAT.v16.i3.8356
- Kavitha, K. S., Kamalakumari, J., Krishna, G., Najma, U., and Rama, S. (2025). “AI and quantum computing for accelerating drug discovery and precision medicine” in *Modern SuperHyperSoft computing trends in science and technology* (Hershey, PA: IGI Global Scientific Publishing), 33–60.
- Khan, S., Naz, N. S., Mazhar, T., Tariq, M. U., Shahzad, T., Guizani, S., et al. (2026). Green AI techniques for reducing energy consumption in AI systems. *Array* 29:100652. doi: 10.1016/j.array.2025.100652
- Kolachevsky, N. N., Khabarova, K. Y., Zalivako, I. V., Semenin, N. V., Zhadnov, N. O., Galstyan, K. P., et al. (2025). Quantum computing with trapped ions: principles, achievements, and prospects. *Physics-Uspekhi* 68, 552–583. doi: 10.3367/UFNe.2024.12.039884
- Kumar, A., Salim, M. M., Camacho, D., and Park, J. H. (2025). A comprehensive survey on large language models for multimedia data security: challenges and solutions. *Comput. Netw.* 267:111379. doi: 10.1016/j.comnet.2025.111379
- Lee, S. U., Perera, H., Liu, Y., Xia, B., Lu, Q., Zhu, L., et al. (2025). Integrating ESG and AI: a comprehensive responsible AI assessment framework. *AI Ethics* 5, 5121–5148. doi: 10.1007/s43681-025-00741-5
- Li, J., Guan, Q., Tao, D., and Jiang, W. (2023). Carbon emissions of quantum circuit simulation: more than you would think. *Proceedings of the 14th International Green and Sustainable Computing Conference*, 11–13.
- Li, P., Islam, M. J., and Ren, S. (2025). A case study of environmental footprints for generative AI inference: cloud versus edge. *ACM SIGMETRICS Perform. Eval. Rev.* 53, 21–26. doi: 10.1145/3764944.3764950
- Lukyanenko, R. (2025). *Quantum data management to reduce environmental impact of IT. Pre-ICIS Workshop Proceedings 2024.*
- Ma, K., and Zhou, Y. (2024). A comprehensive quantitative lifecycle cost and environmental impact analysis model for computing infrastructure. *MethodsX* 13:103009. doi: 10.1016/j.mex.2024.103009
- Mamat, R., Ghazali, M. F., Basrawi, F., and Rosdi, S. M. (2025). Emerging role of generative AI in renewable energy forecasting and system optimization. *Sustain. Chem. Clim. Action* 7:100099. doi: 10.1016/j.scca.2025.100099
- Meier, F., and Yamasaki, H. (2025). Energy-consumption advantage of quantum computation. *PRX Energy* 4:023008. doi: 10.1103/PRXEnergy.4.023008
- Mohammed, A., and Kora, R. (2025). A comprehensive overview and analysis of large language models: trends and challenges. *IEEE Access* 13, 95851–95875. doi: 10.1109/ACCESS.2025.3573955
- Moravec, V., Gavurova, B., and Kovac, V. (2025). Environmental footprint of GenAI: changing technological future or planet climate? *J. Innov. Knowl.* 10:100691. doi: 10.1016/j.jik.2025.100691
- Morstyn, T., and Wang, X. (2024). Opportunities for quantum computing within net-zero power system optimization. *Joule* 8, 1619–1640. doi: 10.1016/j.joule.2024.03.020
- Nguyen, S., Zhou, B., Ding, Y., and Liu, S. (2024). Towards sustainable large language model serving. *ACM SIGENERGY Energy Inform. Rev.* 4, 134–140. doi: 10.1145/3727200.3727220

- Nikolopoulou, K. (2025). Generative artificial intelligence and sustainable higher education: mapping the potential. *J. Digit. Educ. Technol.* 5:ep2506. doi: 10.30935/jdet/15860
- Oloruntoba, O. (2024). Green cloud computing: AI for sustainable database management. *World J. Advanced Res. Rev.* 23, 3242–3257. doi: 10.30574/wjarr.2024.23.3.2611
- Paula, E., Soni, J., Upadhyay, H., and Lagos, L. (2025). Comparative analysis of model compression techniques for achieving carbon efficient AI. *Sci. Rep.* 15:23461. doi: 10.1038/s41598-025-07821-w
- Ren, S., Tomlinson, B., Black, R. W., and Torrance, A. W. (2024). Reconciling the contrasting narratives on the environmental impact of large language models. *Sci. Rep.* 14:26310. doi: 10.1038/s41598-024-76682-6
- Root, D. (2025). Quantum technologies in the context of climate change: emphasizing sustainability in a responsible innovation approach to quantum innovation. *NanoEthics* 19:4. doi: 10.1007/s11569-025-00468-x
- Rouhani, S. H., and Su, C.-L. (2026). Future cyber-resilient renewable and sustainable smart grids: a critical review from power system researchers' perspective on emerging threats and solutions. *Renew. Sust. Energ. Rev.* 226:116299. doi: 10.1016/j.rser.2025.116299
- Samsi, S., Zhao, D., McDonald, J., Li, B., Michaleas, A., Jones, M., et al., (2023). *From words to watts: benchmarking the energy costs of large language model inference. 2023 IEEE High Performance Extreme Computing Conference (HPEC)*, 1–9.
- Schneider, I., Xu, H., Benecke, S., Patterson, D., Huang, K., Ranganathan, P., et al. (2025). An introduction to life-cycle emissions of artificial intelligence hardware. *IEEE Micro* 45, 9–19. doi: 10.1109/MM.2025.3592568
- Shi, J., Yang, Z., Kang, H. J., Xu, B., He, J., and Lo, D. (2024). *Greening large language models of code. Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Society*, 142–153.
- Stein, A. L. (2024). Generative AI and sustainability. In Philipp Hacker, and Others (Eds), *The Oxford handbook of the foundations and regulation of generative AI*, Oxford, United Kingdom: Oxford Academic.
- Theophilo, K., and Oughton, N. (2025). How might quantum computing impact climate change and the wider environment? *Res. Directions Quantum Technol.* 3:e4. doi: 10.1017/qtt.2025.1
- Tripathi, S., Misra, R., and Singh, T. N. (2025). Anomaly-aware quantum convolutional neural network for carbon-efficient job arrival rate prediction in cloud computing. *Futur. Gener. Comput. Syst.*:108096. doi: 10.1016/j.future.2025.108096
- Vasanthrao, K. S., and Saxena, A. (2025). Bits to qubits: an overview of quantum computing. *2025 International Conference on Intelligent Control, Computing and Communications (IC3)*, 120–124. Doi:10.1109/IC363308.2025.10956309
- Vudugula, S., Chebrolu, S. K., Zaman, S., and Saha, R. (2025). Ai and quantum computing for carbon-neutral supply chains: a systematic review of innovations. *Am. J. Interdiscip. Stud.* 6, 40–75. doi: 10.63125/nrdx7d32
- Waelen, R., and van Wynsberghe, A. (2025). Considering the social and economic sustainability of AI. *Sci. Eng. Ethics* 31, 19–13. doi: 10.1007/s11948-025-00544-1
- Wang, P., Zhang, L.-Y., Tzachor, A., and Chen, W.-Q. (2024). E-waste challenges of generative artificial intelligence. *Nat. Computational Sci.* 4, 818–823. doi: 10.1038/s43588-024-00712-6
- Zhang, Y., Qin, C., Wang, H., Hu, X., and Gao, Y. (2024). Occurrence, behavior, and ecotoxicity of quantum dots in aquatic environments. *Environ. Sci. Nano* 11, 4044–4059. doi: 10.1039/D4EN00542B
- Zheng, C., Wu, G., Bao, F., Cao, Y., Li, C., and Zhu, J. (2023). *Revisiting discriminative vs. generative classifiers: theory and implications. International Conference on Machine Learning (ICML 2023)*, 42420–42477.