Evaluating and Comparing LLM Models for CO_2 Footprint and Sustainable Green AI

Gaurav Sharma

AI Test Automation Solution Architect
https://orcid.org/0009-0008-6311-9531
IEEE Sr. Member, Atlanta, GA, USA
Email: *gauravbox[at]ieee.org*

Abstract: ***Large Language Models (LLMs) have significantly advanced artificial intelligence applications across various domains. However, these developments come at environmental cost in terms of carbon emissions, energy consumption, and water usage for cooling data centers. This paper evaluates and compares popular LLMs based on their sustainability and carbon footprint, analyzing key factors such as training and computing CO2 emissions, inference emissions, energy consumption, and water usage. We also discuss concerns and strategies for optimizing AI systems to reduce their environmental impact, emphasizing the need for responsible and sustainable AI development.***

Keywords: Artificial Intelligence, Green AI, AI Co2 footprint, Sustainable AI, LLM (Large Language Models), Global Impact.

1. Introduction

The rapid adoption of Large Language Models (LLMs) has transformed industries, driving innovation in natural language processing, automation, and decision-making systems. However, these computationally intensive models require vast amounts of resources, raising concerns regarding sustainability. The environmental cost comes with training and operating these models includes high electricity consumption, significant carbon emissions, and extensive water usage for cooling data centers. Identifying and addressing these challenges requires a deeper understanding of the trade-offs between model performance and environmental impact. This paper aims to provide a comparative analysis of leading LLMs and explore strategies to mitigate their environmental footprint.[1]

2. Sustainability in AI: Why It Matters

AI-driven technologies are projected to expand exponentially, making it imperative to develop energy-efficient models. Traditional AI development focuses on improving accuracy and scalability, but sustainability factors are often overlooked. Key sustainability concerns include:

- **Carbon Footprint:** The greenhouse gas emissions resulting from energy consumption in model training and inference.
- **Energy Efficiency:** The amount of power needed to run those powerful high-performance computing clusters.
- **Water Consumption:** The cooling requirements for AI servers, which can reach millions of liters per year.
- **Ethical Responsibility:** The need for AI developers to adopt environmentally friendly practices and optimize energy utilization.

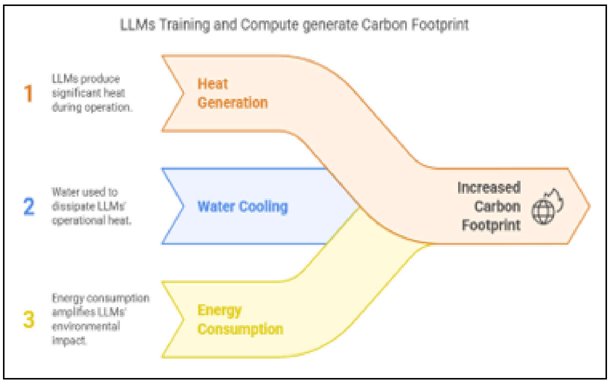

Figure 1: LLMs operations generate carbon footprints

3. Methodology: Evaluation LLM Models

To assess the environmental impact of LLMs, we use the following sustainability metrics:

- **Training CO2 Emissions (kg CO2e):** The carbon footprint incurred during model training.
- **Inference CO2 Emissions (kg CO2e/query):** The emissions generated when processing a single query.
- **Energy Consumption (MWh):** The total electricity usage in megawatt-hours.
- **Water Usage (Liters):** The estimated water consumption required to cool data centers hosting these models.
- **Sustainability Initiatives:** The measures taken by AI providers to reduce environmental impact.

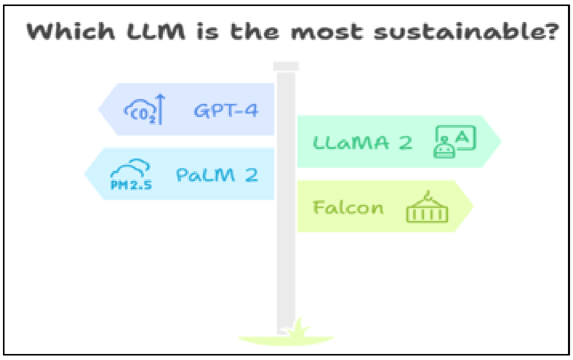

Figure 2: LLMs can be evaluated on sustainable metrices

4. Comparative Analysis of LLMs:

The following table presents a comparative assessment of major LLMs based on the aforementioned sustainability factors.[4]

This comparison metrices and publication from big companies building LLM model should be transparent to measure real carbon Footprint of models from training and computational process.

4.1 Table to compare Co2 foot print

The table below provides a comparative evaluation of major LLMs with respect to sustainability factors.[4].

Let's look at the carbon footprint and water usage that comes with training and running large language models (LLMs). It's good to remember that AI is still in its infancy, developing quickly, and starting to be used widely in businesses and public sectors. With the growing demand for these advanced models, we're seeing an increase in the use of resources like water, coal, and renewables to keep up with the needs of this exciting and rapidly evolving technology.

Table 1: Comparative assessment of major LLMs based on sustainability factors.

Model	Training CO2 Emissions (kg CO2e)	Inference CO2 Emissions (kg CO2e/query)	Energy Consumption (MWh)	Water Usage (Liters)	Sustainability Initiatives
GPT-4	~502,000	0.02	~1,287	~700M	Optimized architecture, Microsoft's carbon-neutral data centers
LLaMA 2	~350,000	0.015	~900	~500M	Meta's efficiency improvements in model training
PaLM 2	~440,000	0.018	~1,100	~600M	Google's carbon-neutral cloud infrastructure
Falcon	~180,000	0.012	~600	~250M	Open-weight model, trained with efficiency in mind
Mistral 7B	~120,000	0.009	~400	~150M	Optimized for lightweight deployment
Claude 2	~320,000	0.014	~800	~450M	Anthropic's research into green AI
Grok	~280,000	0.013	~750	~350M	X's AI division focusing on efficiency
Deep Seek	~200,000	0.01	~500	~200M	Trained with sustainability in mind, China-based optimization

5. Environmental Impact of LLM CO2 Emissions

The widespread use of LLMs has a notable environmental impact, primarily driven by CO2 emissions, water consumption, and energy demands. Below are the key effects of these emissions:

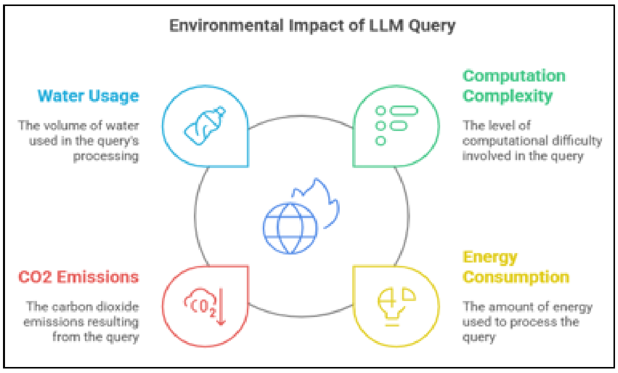

Figure: LLM Co2 footprint impact on Environment

1) Climate Change Contribution
The CO2 emissions generated by training and inference processes add to global greenhouse gas (GHG) levels, exacerbating climate change. For instance, training a single large-scale LLM like GPT-4 can emit as much carbon as several hundred transatlantic flights.

2) Increased Energy Demand
As LLMs become more complex, their energy requirements grow significantly. Data centers hosting AI models consume vast amounts of electricity, often sourced from fossil fuels, which further intensifies the carbon footprint of AI applications.

3) Water Scarcity and Thermal Pollution
Data centers require large quantities of water for cooling, which can contribute to water scarcity in regions with limited freshwater resources. Additionally, the release of heated water into natural water bodies can disrupt local ecosystems, harming aquatic life.

4) Electronic Waste (E-Waste) Production
Artificial intelligence relies on powerful hardware, leading to regular upgrades and the disposal of old components. When we toss outdated GPUs, TPUs, and other devices, we create electronic waste that can release harmful toxins into the environment if not handled correctly.

5) Air Pollution and Resource Extraction
The production and maintenance of AI hardware involve mining rare-earth metals, leading to habitat destruction, deforestation, and increased air pollution due to industrial activities.

6. Carbon Offsets: A Real Solution or a Band-Aid Fix?[3]

Some companies attempt to offset their carbon footprint by purchasing carbon credits and essentially paying to have a "neutralize" their emissions. However, this is not the same as actually reducing energy consumption. While companies like Google Cloud and Microsoft Azure buy carbon credits to offset their emissions, Amazon AWS—the largest cloud computing provider—only covers 50% of its energy use with renewable sources.

Paper ID: SR25304105212 DOI: https://dx.doi.org/10.21275/SR25304105212

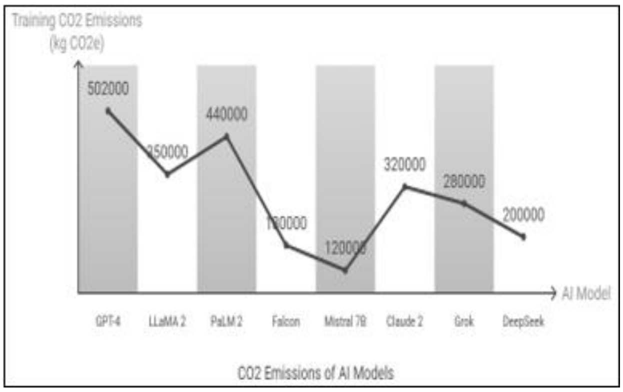

Figure 3: Graph trend of co2 emission by LLM models

An important question arises: Should AI companies prioritize reducing energy consumption instead of just offsetting emissions? Improving the efficiency of AI models could reduce the need for carbon credits.

Additionally, the effectiveness and transparency of carbon offsets are questioned. Some doubts that companies use offsets to appear environmentally friendly without making meaningful changes. There's also the issue of double counting, where the same carbon reduction is claimed by multiple parties, undermining the credibility of the offset market.

Verifying and certifying carbon offsets can be complex and inconsistent, leading to exaggeration in the actual environmental benefits. Heavy reliance on carbon offsets might hamper real progress towards sustainability.

Despite these concerns, some offset projects have made positive impacts. Projects focusing on reforestation, renewable energy, and methane capture from landfills can offer environmental and social benefits. But these projects need careful selection, monitoring, and management to deliver their promised outcomes.

Ultimately, while carbon offsets can help mitigate the environmental impact of AI, they are not a cure-all. A holistic approach that includes improving energy efficiency, investing in renewable energy, and developing sustainable practices is essential for the long-term sustainability of AI technologies.

7. Transparency in AI Research: The Need for Better Reporting [3]

A major challenge in AI research is the lack of transparency in reporting computational costs. Many studies only report the final best result without revealing how many failed experiments were run in the background. This makes it difficult for future researchers to understand how much time, money, and energy were spent to achieve that final breakthrough.

If researchers fully reported all experiments (including failed ones), it would help the AI community better understand the real cost of progress. This would also encourage more efficient research practices, saving both energy and money.

7.1 Example Queries and Environmental Cost

To illustrate the real-world impact of LLM queries on resource consumption, consider the following examples:

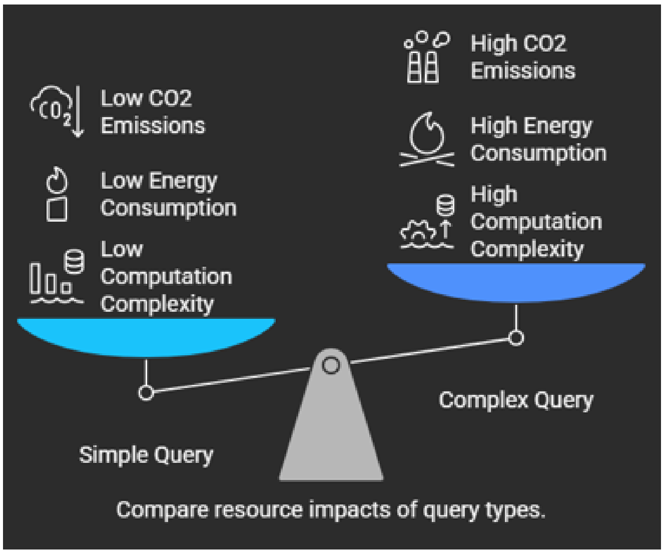

Figure 4: Compare Resources used between Simple and complex query

Example 1: Simple Query
Query: "How many planets have water in the Milky Way galaxy?"
- **Computation Complexity:** Low
- **Energy Consumption:** ~0.02 kWh
- **CO2 Emissions:** ~0.02 kg CO2e
- **Water Usage:** ~1 liter (equivalent to a small water bottle)

Example 2: Complex Query
Query: "Analyze exoplanet datasets and identify potential habitable planets with Earth-like conditions in the Milky Way. Rank them based on atmospheric composition, distance from their star, and surface water probability."
- **Computation Complexity:** High
- **Energy Consumption:** ~0.3 kWh
- **CO2 Emissions:** ~0.25 kg CO2e
- **Water Usage:** ~15 liters (equivalent to three large water bottles)

These examples highlight that even a single complex AI query can significantly contribute to energy and water consumption, reinforcing the need for optimized computation strategies.

7.2 Understanding the Co2 Footprint of Training AI Models

Training AI models is like running a giant, high-powered engine—it needs a lot of electricity. The more power it uses, the more CO2 emissions it produces, just like how a car burns fuel and releases emissions.

Scientists have found that there's a direct link between energy consumption and CO2 emissions. The more energy a model consumes, the more CO2 it emits. However, not all energy is equal. If a model is trained using clean energy (like hydroelectric power), it releases far fewer emissions than a model trained using coal or gas-based power.

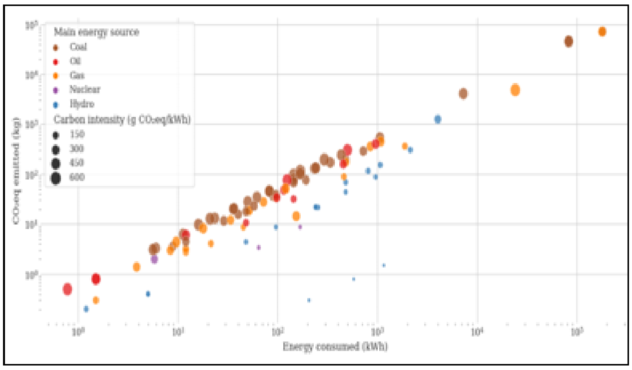

Figure 1: Estimated energy consumed (kWh) and CO2 (kg) by each model in the data set. Image Source [2]

8. How Big Are the Emissions?

AI models do not use the same amount of energy. Some smaller models only need around 10 kWh, while large-scale models like GPT-4 can consume over 10,000 kWh—a difference of 1000 times! And, naturally, the more energy they consume, the more carbon emissions they produce.

To visualize this, imagine two AI models that consume the same amount of energy:
- Model A is powered by hydroelectric energy (clean energy) → It produces very little CO2.
- Model B is powered by coal energy → It emits 100 times more CO2 for the same energy used!

In other words, where the power comes from matters as much as how much power is used.

Other Factors That Affect Emissions
Apart from the energy source, the hardware running these models also plays a role. Some high-performance chips use more power than others. However, this factor has a smaller impact compared to the choice of energy source.

9. Optimizing AI for Efficiency: Key Methods for Green and Sustainable AI [5]

To minimize computational costs and reduce environmental impact, AI can leverage innovative strategies rather than relying on extensive model training from scratch. Below are some of the most effective techniques:
- Knowledge Distillation – A compact AI model is trained by extracting knowledge from a larger model, achieving comparable accuracy while requiring significantly less energy.
  Benefit: Streamlined AI-powered applications that run more efficiently with smaller models.
- Pruning & Quantization – By eliminating redundant components (pruning) and lowering numerical precision (quantization), models become more resource-efficient while preserving accuracy.
  Benefit: These approaches can improve energy efficiency by a factor of 3 to 7, significantly lowering AI's carbon footprint.
- Fine-Tuning & Transfer Learning – Instead of building AI models from scratch, pre-trained models can be adapted for new tasks, reducing computational demands and training time.
  Benefit: Faster AI development with substantially lower energy consumption.
- Sparse Models & Mixture of Experts (MoE) – AI models activate only relevant components for specific tasks, preventing unnecessary processing.
  Benefit: MoE-based architectures can be up to 10 times more energy-efficient, making them ideal for large-scale applications.

By integrating these techniques, AI can be made more efficient, sustainable, and accessible, ensuring high performance with minimal environmental impact

9.1 Discussion: Key Findings and Insights

1) **Energy-Intensive Models:** GPT-4 and PaLM 2 exhibit the highest carbon footprint, emphasizing the trade-offs between performance and sustainability.
2) **Balanced Efficiency:** LLaMA 2 and Claude 2 strike a balance between model size, performance, and ecological impact.
3) **Leading Sustainable Models:** Mistral 7B, Falcon, and DeepSeek are the most energy-efficient, consuming significantly less power and water.
4) **Water Consumption is a Hidden Cost:** AI models require millions of liters of water for cooling, highlighting the importance of eco-conscious infrastructure choices.
5) **Adoption of Carbon-Neutral Technologies:** Some AI providers leverage renewable energy sources and optimized architectures to offset emissions.

9.2 Strategies for Reducing AI's Environmental Impact

To build sustainable AI solutions, we recommend the following best practices:
- **Model Compression and Optimization:** Implement pruning, quantization, and knowledge distillation techniques to reduce model size and energy use.
- **Green Data Centers:** Deploy AI workloads in facilities powered by renewable energy.
- **Efficient Query Processing:** Minimize redundant queries and batch processes to reduce carbon emissions.
- **Water-Efficient Cooling Systems:** Explore alternative cooling technologies, such as liquid cooling, to minimize water wastage.
- **Regulatory Compliance:** Encourage policy frameworks that promote responsible AI development and sustainability initiatives.

10. Conclusion

Artificial intelligence has transformed industries, expanding the limits of machine capabilities. However, this rapid advancement comes with an environmental cost, including high energy consumption, carbon emissions, and significant water usage. While we celebrate AI's potential, we must also take responsibility for its sustainability.

This paper highlights the often-overlooked ecological impact of training and deploying large-scale AI models. If left unregulated, AI's increasing energy demands could place a heavy burden on natural resources. However, a more

Paper ID: SR25304105212 DOI: https://dx.doi.org/10.21275/SR25304105212

sustainable approach is within reach—one that prioritizes energy-efficient models, greater transparency in AI research, and responsible policies from both businesses and governments.

The progress of AI and the well-being of our planet do not have to be in conflict. By integrating environmentally friendly AI practices and leveraging renewable energy sources, we can continue to innovate without compromising sustainability. The real question is not whether AI will evolve, but whether we will guide its growth in a way that aligns with environmental responsibility.

The decisions we make today will shape the future. By designing AI systems with sustainability at their core, we can drive technological progress while preserving the planet for generations to come.

References

[1] A. Bonnaccorsi, "On the Relationship between Firm Size and Export Intensity," Journal of International Business Studies, XXIII (4), pp. 605-635, 1992. (journal style)
[2] Bannour, N., Ghannay, S., Névéol, A., & Ligozat, A. L. (2021, November). Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools. In Proceedings of the second workshop on simple and efficient natural language processing (pp. 11-21).
[3] Luccioni, A. S., & Hernandez-Garcia, A. (2023). Counting carbon: A survey of factors influencing the emissions of machine learning. arXiv preprint arXiv:2302.08476.
[4] Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green ai. Communications of the ACM, 63(12), 54-63.
[5] Strubell, E., Ganesh, A., & McCallum, A. (2020, April). Energy and policy considerations for modern deep learning research. In Proceedings of the AAAI conference on artificial intelligence (Vol. 34, No. 09, pp. 13693-13696).
[6] Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., ... & Dean, J. (2021). Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.

Author Profile

Gaurav Sharma has 20 years of extensive experience in leading and managing software engineering and quality teams. With expertise in cutting-edge technologies like AI-driven solutions, he has spearheaded innovative software solutions that have transformed industries such as Aviation, Telecom, Financial Services, CPG (Consumer Packaged Goods), Retail, and ERP systems. Gaurav has designed and implemented AI-driven automation frameworks seamlessly integrated into CI/CD pipelines, enabling a more robust software development lifecycle and accelerating time-to-market for software products and solutions. As Author, Speaker, Mentor and a passionate AI enthusiast and advocate for Responsible AI, he has multiple publications including IEEE international conferences. His thought leadership gives back to the community by volunteering for educational events at high and elementary schools, judging and mentoring AI awards. Mr. Sharma currently working on innovative AI agent to revolutionize talent industry which soon will be shared on LinkedIn and his YouTube channel **"AI – Test Hive"**

Paper ID: SR25304105212 DOI: https://dx.doi.org/10.21275/SR25304105212