




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 11, November 2025

Green AI: Energy-Efficient Machine Learning Model Selector

DR. Pankaj Agarkar¹, Rudrani Sonawane², Loukik Kawde³, Maheshchand Suthar⁴

Department of Computer Engineering, Ajeenkya DY Patil School of Engineering, Lohegaon Pune, India¹⁻⁴

ABSTRACT: Advancements in artificial intelligence and machine learning have taken on key roles across many different sectors. They push forward things like automated systems, medical diagnosis tools, financial operations, and even self-driving vehicles. Still, these kinds of setups demand a lot of power, both during the training periods and when they run in real time. Data volumes keep growing larger, and models get more intricate all the time. That leads to bigger environmental effects from the extra energy needed to handle all that information. In the end, it creates real risks for both the planet's health and company budgets.

Researchers have come up with a fresh approach to green artificial intelligence. This system looks at standard performance measures, such as accuracy and recall rates. It also weighs environmental aspects like how much computing power gets used and the amount of greenhouse gases produced. All this happens when picking out machine learning algorithms. The system, known as GES, checks how models do on efficiency, overall energy needs, and their impact on the environment. Software developers can use this to select the best machine learning options that fit their projects. At the same time, it encourages more eco-friendly AI practices that match up with global goals for protecting the environment.

KEYWORDS: Green AI, Machine Learning, Energy Efficiency, Model Selection, Sustainable Computing, Carbon Footprint

I. INTRODUCTION

Artificial intelligence and machine learning form the core of contemporary technological progress. These fields contribute essentially to areas like autonomous vehicles, advanced voice recognition systems, detailed medical data examination, accurate financial forecasting, and broad ecological modeling. Such developments have transformed computer-human interactions. Computers now process information more effectively, identify patterns with greater precision, and make decisions informed by analytical results. Despite the substantial progress driven by large-scale AI implementations, growing global concern emerges regarding their environmental consequences. Progress in AI has accelerated through improvements in deep learning designs and the availability of enormous datasets, both demanding intensive computational power.

Training and deploying these systems result in considerable energy consumption. This in turn adds notably to carbon dioxide emissions. Research from scholars at the University of Massachusetts Amherst reveals that developing a large natural language model such as GPT-3 produces about 284 metric tons of CO₂. That figure appears remarkably high. It matches the emissions from around five average cars across their full lifespans. As AI capabilities advance, energy requirements escalate sharply alongside expanding processing demands. This situation underscores the need for responsible, eco-aware strategies in handling data computations. Traditionally, machine learning specialists assess their models using metrics focused on accuracy, such as error rates, sensitivity, and specificity. While these measures prove crucial for gauging performance, they overlook key aspects of sustainability, including power usage, operational speed, and emission levels. Current ecological challenges call for evaluations that extend past simple accuracy. They demand holistic reviews incorporating broader impacts. An algorithm that slightly improves prediction accuracy could still demand far more energy. This creates a challenging trade-off between operational efficiency and environmental care. Such recognition has spurred the emergence of green AI. This paradigm prioritizes sustainable, nature-friendly progress in AI systems. Green AI seeks to optimize performance while minimizing resource demands. It encourages methods that harmonize computational needs with ecological considerations during algorithm design, model architecture

choices, and training procedures. Embedding environmental awareness from the outset could lessen the harmful effects of heavy computing on the planet.

The Green AI framework provides a practical path for choosing and appraising energy-thrifty machine learning approaches. Within this effort, the Green Efficiency Score emerges as a combined metric. It assesses not just how accurately algorithms predict results but also their ecological footprint via measured resource use. The GES delivers a fair way to compare models. Researchers and users can weigh both predictive strength and sustainability factors equally. Through this structure, Green AI intends to reshape evaluation standards in AI creation and deployment. It pushes for priorities centered on enduring practices over pure outcomes alone. In doing so, it aligns innovations with global environmental goals. Beyond technical shifts, the initiative fosters awareness among scientists, teachers, and industry experts about moral dimensions in AI work. Green AI progresses by boosting output without eroding potential. It represents a key step toward a balanced future where technological growth respects planetary limits.

II. MOTIVATION

Green AI came about because of some big problems in today's AI tech. Things like how much power these setups use. And there aren't really good ways yet to measure the environmental hit from machine learning projects. AI keeps spreading into all sorts of fields and industries these days. That means we have to deal with the eco costs and money costs from needing so much more computing muscle. Lately the push for compute power to build top AI has sped up like crazy. Research points to the resources for big AI efforts doubling every three years or so. That's way faster than folks expected from Moore's law patterns. That old idea predicted transistor density would quadruple each year back in the day. But now speedups happen in just four months sometimes. All this quick computing growth pulls in huge energy through giant networks. Those handle training and running complex AI models. The International Energy Agency says global data centers gulp down 200 to 250 terawatt hours of electricity yearly right now. That works out to about one percent of the world's total power use. Even with tech getting better in AI and models needing bigger data piles. People in the field see usage ramping up a lot more in the years ahead.




Fig. 2. Overview of the problem statement illustrating the relationship between high model accuracy, energy consumption, and environmental cost in AI systems.




Fig. 3. Illustration of AI's increasing carbon footprint with growing computational demands, highlighting the imbalance between accuracy and sustainability.

A major problem here involves the fact that much of this energy still relies on fossil fuels and other nonrenewable sources. This reliance leads to substantial carbon emissions and environmental degradation. In typical modern machine learning workflows, experts focus mainly on boosting model accuracy. They often overlook the underlying environmental and energy costs. For instance, a deep convolutional neural network might outperform a simpler model like a decision tree by only a small margin in accuracy. Yet the neural network could demand over fifty times more computational resources and energy. These performance driven evaluation practices increase the carbon footprint of AI. They do so without delivering matching gains in efficiency or sustainability. Green AI aims to tackle this imbalance. It does this by integrating sustainability into the core principles of AI research and development. The main goal centers on making energy efficiency and carbon awareness built in parts of the model evaluation and selection process. In practical terms, this approach means giving metrics like power consumption and computational load the same weight as accuracy or precision. This shift in perspective seeks to promote responsible AI practices. Such practices balance innovation with environmental responsibility. Beyond cutting environmental impact, these practices also help reduce operational costs for companies that depend heavily on high performance computing. From an economic standpoint, pursuing Green AI brings notable advantages. Large tech firms, research institutions, and AI startups together spend millions of dollars yearly on electricity, cooling systems, and hardware maintenance.

By emphasizing energy efficient algorithms and models, organizations can achieve significant cost savings. They can do this without compromising model quality. Adopting sustainable practices thus provides a dual benefit. It drives operational savings while advancing environmental protection. Green AI also aligns well with the wider global sustainability agenda. This agenda appears in the United Nations Sustainable Development Goals. It supports SDG 9 on industry, innovation, and infrastructure particularly. Green AI promotes advancements in efficient and sustainable computational technologies. These foster responsible digital industrialization. It backs SDG 13 on climate action too. This happens through proactive steps to curb climate change. Such steps involve lowering carbon emissions in AI training and deployment systems. At its core, the drive for Green AI involves reshaping the AI ecosystem. Currently, success gets measured only by predictive accuracy. Green AI pushes for a balance between performance and sustainability instead. By weaving environmental awareness into every stage of model development, Green AI helps fight climate change. It also builds the groundwork for a more efficient, ethical, and forward looking AI landscape.

III. LITERATURE REVIEW

The field of Green AI underscores the essential need to develop energy-efficient methods in machine learning. As applications of artificial intelligence continue to expand, the urgency to tackle their environmental impacts grows accordingly. Several foundational factors and efforts have shaped this emerging discipline. They stress the importance of transparency, sustainability, and efficiency in designing and deploying AI systems. An early influential work came from Strubell et al. in 2019. These researchers were among the first to systematically assess the environmental costs of natural language processing models. Their analysis showed that training one large-scale transformer model can generate about 626,000 pounds of carbon dioxide emissions. This amount roughly matches the lifetime emissions from five average cars. Such a comparison highlights the substantial carbon footprint of deep learning models. The results acted as an alert to the AI field. They sparked widespread discussions on sustainable AI practices and helped launch the Green AI movement.

Subsequent research by Henderson et al. in 2020 built on this base. It advanced the conversation by calling for standardized reporting of computational resources and energy use in AI experiments. The authors suggested that every major machine learning study should record the computational time, hardware details, and carbon impact of model development and training. This idea paved the way for clear and reproducible practices in energy reporting within AI research. It promotes accountability for resource use in experiments.

Practical tools have also emerged to aid researchers in evaluating and managing the environmental effects of machine learning models. Code Carbon, an open-source project from 2023, offers automated tracking of carbon emissions. It estimates outputs based on CPU and GPU activity, runtime length, and the regional energy mix in the computing setup. Developers can integrate sustainability tracking right into Python workflows using a simple API. In addition, pyRAPL provides a lightweight option for measuring CPU power consumption. This makes it suitable for small-scale experiments and academic projects without access to large infrastructure.

Further work by Patterson et al. in 2021, from Google Research, examined efficiency at the ecosystem level for AI computations. Their study looked at how improvements in data center design, hardware performance, and renewable energy sources could greatly lower the energy demands of large AI systems. The researchers also pointed to algorithmic approaches like mixed-precision training and model pruning. These techniques reduce power use while preserving or even improving model accuracy. The findings suggest that optimizations in both hardware and software are vital for creating truly sustainable AI systems. Schmidt et al. in 2022 added to this area by comparing standard deep learning architectures to lightweight neural networks. Their investigation found that streamlined models, such as MobileNet and EfficientNet, can deliver performance similar to typical convolutional neural networks. At the same time, they consume far less energy. This work reinforces the view that innovations in neural architecture design offer a direct path to greener AI solutions. It achieves this without compromising computational power.

In a related development, Rae et al. in 2022 introduced carbon-aware scheduling. Their method involves dynamically timing AI training tasks to coincide with peak availability of renewable energy, like solar or wind power. This strategy could substantially cut carbon emissions from large-scale model training. It presents a practical route to more environmentally responsible AI practices. Despite these advances in studies and tools, a key challenge persists. Most current systems and sustainability efforts cater to large research groups or corporate AI labs with ample technical resources. Smaller organizations, academic settings, and independent researchers often lack access to these advanced tools and facilities. The Green AI initiative aims to close this gap. It develops a comprehensive, easy-to-use platform. This allows users, including students and small firms, to measure, compare, and refine their machine learning models. The focus is on both energy use and predictive accuracy. By making sustainability assessment tools widely available, Green AI seeks to integrate environmental awareness at every stage of AI education and practice. This effort supports the global shift to sustainable computing. It also encourages a culture of responsible innovation throughout the AI community.

IV. METHODOLOGY

The Green AI system is a multi-layered, all-encompassing framework that evaluates, analyzes, and suggests energy-efficient machine learning models. To encourage computational sustainability, the approach incorporates both conventional performance indicators and environmental factors. The method ensures fairness, reproducibility, and transparency across outcomes by methodically evaluating different algorithms under identical experimental conditions. Data collection and preprocessing, model training, performance evaluation, energy consumption measurement, Green Efficiency Score computation, and visualization with model recommendation comprise the six main phases of the complete process.




Fig. 1. System architecture of the proposed *GreenAI: Energy-Efficient Machine Learning Model Selector* framework showing dataset input, model pool, energy profiling, and recommendation engine.

5.1 Data Collection and Preprocessing

The foundation of evaluating machine learning systems lies in obtaining and readying reliable data. In the context of Green AI, datasets come from reliable public sources like the UCI Machine Learning Repository or Kaggle, or sometimes from contributions by users. Once gathered, the data receives careful preprocessing to prepare it for training models. This step involves handling gaps or inconsistencies in values, either by filling them in or removing them, along with dropping any irrelevant or messy parts. Numerical features get normalized too, so they scale evenly across the board. Categorical data then finds suitable encoding, and the whole set splits into training and testing portions, often in an 80:20 split or something close. Such a split helps preserve the data's quality and allows for fair checks on the models. Overall, this preprocessing aims to deliver clean, uniform data to all models, cutting out any biases from uneven inputs.

5.2 Model Training

After preparing the data, the Green AI approach moves into training models, where different machine learning methods get tested out. The setup covers a range of models to show varying levels of computational demands and results. Among them are Logistic Regression, Decision Tree, Random Forest, Support Vector Machine, or SVM, K-Nearest Neighbors, known as KNN, and even Convolutional Neural Network for deeper learning tasks. Each one trains on the identical dataset, under the same hardware and software conditions, to keep comparisons even. Experiments usually run on standard processors, with set GPU access and memory limits. Maintaining this steady environment means that differences in how models perform or use energy stem purely from the algorithms themselves, not from outside factors.

5.3 Performance Evaluation

Evaluating performance forms the core analysis stage, where trained models face detailed scrutiny on accuracy and resource use. Green AI draws on several key metrics for this. Accuracy measures how often predictions hit the mark overall. Precision looks at the share of positive predictions that turn out correct. Recall checks how well the model spots true positives. The F1-score combines precision and recall into a balanced average. Training time and inference time track the efficiency of computations. These together offer a full picture of how well models predict and what they demand in terms of time and power. In this way, the process highlights a model's worth beyond just accuracy, factoring in the effort needed to reach those results.

5.4 Energy Consumption Measurement

To gauge the ecological footprint of computing, Green AI employs tools like Code Carbon and pyRAPL for tracking energy. These open-source options capture live usage from CPU and GPU during training and testing. Code Carbon tallies energy in watt-hours and converts it to carbon emissions, based on local grid details that factor in renewable sources. PyRAPL, on the other hand, provides specifics on CPU power draw, which suits lighter models without much GPU involvement. Energy figures get logged methodically for clear reports and deeper review. Bringing in these tools enables precise tracking of each algorithm's sustainability impact through its energy patterns.

5.5 Green Efficiency Score (GES) Calculation

Green AI combines performance and energy data into one score called the Green Efficiency Score, or GES. This metric balances accuracy against energy use via a formula. It weights the accuracy of a model and its energy draw, with adjustable factors that often favor accuracy at 0.7 and energy at 0.3. Higher scores point to models that manage strong predictions with lower environmental costs. The GES underpins the system's suggestions, helping users spot options that maximize benefits while minimizing harm. Visualization and Recommendation. The final part focuses on displaying results and offering picks. Metrics like accuracy, training length, and energy appear in graphs, charts, and side-by-side views. Scatter plots might show accuracy plotted against energy, bar charts compare emissions across models, and rankings highlight GES leaders. From there, the system suggests the best model for balancing accuracy and energy needs. Users gain practical guidance to select fitting options for their goals on performance and green practices. The visuals also educate, showing links between compute demands and eco effects. In the end, Green AI weaves technical checks with environmental awareness. Embedding energy and carbon tracking throughout the machine learning process fosters ethical development without sacrificing results. This framework shows promise for research settings, teaching environments, and industry uses where efficiency and sustainability matter equally.

V. ADVANTAGES & DISADVANTAGES

6.1. Advantages.

The Green AI framework, designed for energy efficient machine learning model selection, brings several key benefits in academic, industrial, and environmental areas. It improves computational efficiency while introducing a new way to think about sustainable artificial intelligence. Several main advantages stand out.

Green AI pushes the AI field toward more environmentally friendly approaches. It does this by making sustainability a measurable part of the process. Traditional methods usually just look at accuracy alone. Green AI uses its Green Efficiency Score, or GES, to include energy use and carbon effects in evaluations. Researchers and companies then consider the environmental side of their work. This fits with bigger goals like the Paris Climate Agreement and the United Nations Sustainable Development Goals.

The framework offers a balanced way to evaluate models. It combines accuracy with energy efficiency in a two part metric. The GES helps avoid over focusing on performance that wastes too much power. This balance supports practical and responsible use of models. It matters especially in places with limited energy, such as mobile devices or Internet of Things setups.

Energy efficient models cut down on resource needs. They use less GPU time, less cooling, and less electricity. Institutions and industries see lower costs as a result. In the long run, these models lead to big savings. Large scale AI systems, which run computations all the time, benefit the most.

Green AI also encourages better use of hardware. It recommends models that fit specific datasets and setups well. Developers look at both performance and energy data. They pick the most economical hardware options. This avoids extra power use and wear on equipment.

6.2. Disadvantages.

Green AI has many strong points. Still, it faces some limitations and challenges. Future work could tackle these issues. The main drawbacks appear next.

Energy measurements depend heavily on hardware sensors and tools like Code Carbon or pyRAPL. If the hardware does not support fine grained power tracking, results might not match real energy use. Readings can vary between devices too. This creates problems in comparing models fairly.

Adding real time energy monitoring brings a small extra load on computation. During training, it takes resources to record and save energy information. Runtime might slow a bit as a result. Smaller systems or low power machines feel this more.

Model performance can change with different datasets. Sizes and complexities vary a lot. A model efficient on one set might need much more power on another. Factors like input sizes, data preparation, or processing needs play a role. Green AI suggestions thus depend on the context. They need checking again for new uses.

No standard benchmarks exist yet for energy efficiency in machine learning models. Without them, evaluations differ across research groups. It makes uniform testing hard. A shared benchmark dataset could fix this down the line.

Green AI involves a trade off between accuracy and efficiency. In key areas like medical diagnosis or self driving cars, even small drops in accuracy might not work. Energy savings alone may not justify it. The framework might not suit every situation where precision comes first.

VI. RESULTS

The Green AI framework received a detailed assessment through several standard benchmark datasets. These included Iris, MNIST, and CIFAR-10. The goal was to check how well it managed model accuracy alongside energy use. Researchers picked these datasets because of their differences in size, complexity, and setup. That allowed testing over many kinds of machine learning models. All experiments ran on the same hardware and software setups. This ensured that differences in performance or energy came only from the algorithms. Not from outside factors. Results showed a clear pattern. Complex deep learning models often hit high accuracy marks. Yet simpler, traditional methods gave similar results. They did so with much less energy and computing power. Evidence like this backs up Green AI's main idea. Sustainable AI does not always need huge resources. Numbers from the tests appear in the table below. Analysis of Results. From the table, key points stand out. Logistic Regression struck a strong balance in accuracy and energy savings. It reached a Green Efficiency Score of 0.89. That topped all other models tested. The CNN did get a bit higher accuracy at 98 percent. Still, its training took about 55 seconds. It used 8.5 kWh of energy. Around 34 times more than Logistic Regression.




Fig. 4. Relationship between model accuracy and corresponding carbon emissions, showing the exponential environmental cost of achieving marginal accuracy gains




Fig. 5. Comparative plot showing CO₂ emissions across different machine learning models, emphasizing the steep energy–accuracy trade-off.

Studies indicate that attempts to gain even small improvements in accuracy often lead to much higher costs in computing resources and greater harm to the environment. The Decision Tree model performed in a similar way. It reached 95 percent accuracy while keeping energy use quite low. The Random Forest model also delivered solid accuracy levels. Still, it demanded far more energy because of the complex ensemble methods involved in its structure. Evidence points to clear diminishing returns here. Adding more layers, parameters, or resources just for slight boosts in accuracy does not pay off as expected. These trade-offs make it essential to factor in sustainability measures right next to predictive accuracy when assessing AI models overall. Graphs generated by the Green AI framework help confirm these key insights.




Fig. 6. Comparison of predictive accuracies for Logistic Regression, CNN, and Transformer models under uniform experimental conditions.

The graphs illustrating accuracy in relation to energy consumption exhibited a sharp rise in energy usage after reaching a certain accuracy threshold—this curve represents what can be termed the accuracy–energy saturation point. This saturation point signifies the phase where further computational efforts yield minimal improvements in accuracy while considerably escalating energy requirements. Identifying this threshold is crucial for both researchers and practitioners, as it enables them to make well-informed choices about when model optimization becomes neither computationally nor environmentally sustainable. Furthermore, scatter plots that compare GES scores across various models clearly demonstrated the inverse relationship between energy consumption and overall efficiency. Additionally, bar charts depicting carbon emissions based on Code Carbon data indicated that models like CNNs and large ensemble methods disproportionately contribute to the overall carbon footprint of the system.




Fig. 7. Comparison of estimated carbon emissions for different machine learning models, highlighting Logistic Regression as the most eco-efficient.

The visual results provide valuable understanding of the various trade-offs in performance. They highlight how Green AI serves as a tool for analysis and learning. Evidence indicates that putting Green AI into practice carries significant effects for academic and industry areas alike. Within university research and teaching environments, this system helps cut down the carbon emissions from lab work. It does so by letting students and faculty pick algorithms that deliver sufficient precision while using less power. Such choices result in lower overall electricity consumption in common computing spaces. They also build greater knowledge of eco-friendly computing habits among learners. In business settings, Green AI offers a key advantage for rolling out AI in sustainable ways. Companies can weave in choices based on energy use right into their regular production steps. This approach lowers costs tied to powering servers, managing cooling, and maintaining equipment. Even more important, weaving sustainability measures into how models are built helps match company AI plans with worldwide goals for the environment. These include the United Nations Sustainable Development Goal 13 on climate action.

Experimental results confirm that AI setups can balance energy savings with strong accuracy. This happens through proper tools like Green AI. Studies show that basic machine learning options often work well in place of heavy computational designs. They keep up efficiency and care for the environment. At the same time, they do not sacrifice needed precision. Overall, the Green AI setup points to real advances in creating fair, evidence-based methods for green artificial intelligence.

VII. CONCLUSION

The Green AI initiative represents a key step forward in how we assess and put into practice machine learning models. It introduces the Green Efficiency Score, or GES, which combines accuracy with energy use in a single measure. Evidence shows this kind of unified approach reflects a growing recognition that tech progress needs to align with care for the environment. Studies from this project indicate that building energy-saving AI is not just possible, it seems practical too, without sacrificing how well the models perform.

This system tracks the power demands of various algorithms and feeds that information back into choices about which models to use. In that way, it encourages picking options that deliver solid results while cutting down on harm to the planet. Developers, along with students and researchers, get practical tools here to adopt habits in computing that prioritize sustainability. A major strength lies in how Green AI helps push the world toward AI that keeps carbon emissions in mind. Beyond classrooms and research labs, the setup adapts readily to business settings, where it could cut waste in energy, trim expenses, and ease the strain on hardware setups.

Relying on hard numbers for energy tracking turns what was once an abstract concern into something concrete that shapes decisions on models and tweaks to them. Organizations taking up these ideas stand to boost their overall operations in computing, all while meeting goals that match global pushes for sustainability. The framework sets up

promising paths for what comes next in eco-friendly machine learning. Future work might include things like automatic tuning of models through AutoML, or retraining them based on when green energy is available, even linking up with cloud networks designed for lower emissions. Such steps could make the Green AI idea more real, turning sustainability into a core part of building AI, not something tacked on later. At its heart, this effort aims for a world where model accuracy lives alongside low environmental cost. That change in thinking reminds us that real advances in tech get judged by more than raw speed or power, they have to move forward in ways that are thoughtful and responsible. With tools like Green AI, those in the field can lead the way to solutions driven by data yet kind to the earth, making sure AI growth helps build a healthier planet.

REFERENCES

- [1] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, Florence, Italy, pp. 3645–3650, 2019.
- [2] P. Henderson, J. Hu, J. Romoff, J. Pineau, and E. Brunskill, "Towards the systematic reporting of the energy and carbon costs of machine learning," *J. Mach. Learn. Res.*, vol. 21, no. 248, pp. 1–43, 2020.
- [3] D. Patterson *et al.*, "Carbon emissions and large neural networks: A technical analysis," *IEEE Trans. Sustain. Comput.*, vol. 7, no. 2, pp. 346–357, Apr.–Jun. 2022.
- [4] S. Schmidt, S. Raschka, and S. Roth, "Reducing the carbon footprint of machine learning models," *Sustain. Comput.: Informatics Syst.*, vol. 35, pp. 100752, 2022.
- [5] J. Rae, B. Lascelles, and A. Wong, "Carbon-aware scheduling for AI workloads in cloud systems," *Proc. IEEE Int. Conf. Sustain. Comput. (SustainCom)*, Chongqing, China, pp. 112–119, 2022.
- [6] C. Wu and Y. Jin, "Energy-aware machine learning: Survey and future directions," *IEEE Access*, vol. 9, pp. 119915–119935, 2021.
- [7] H. Zhou, X. Dong, and J. Wang, "Lightweight deep learning models for resource-constrained environments," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 7, pp. 2984–2998, Jul. 2022.
- [8] Y. Li, Z. Zhang, X. Wang, and J. Lin, "A survey on green deep learning: Methods and challenges," *ACM Comput. Surveys*, vol. 54, no. 9, pp. 1–35, 2021.
- [9] R. Gupta, P. Kaur, and S. Arora, "Energy-efficient model selection framework for machine learning applications," *IEEE Access*, vol. 11, pp. 150560–150573, 2023.
- [10] A. Malik and D. Mehta, "Comparative analysis of ML algorithms based on energy efficiency and accuracy," *IEEE Trans. Sustain. Comput.*, vol. 9, no. 1, pp. 66–75, Jan.–Mar. 2024.
- [11] L. F. W. Anthony, B. Kanding, and R. Selvan, "Carbontracker: Tracking and predicting the carbon footprint of training deep learning models," *IEEE Access*, vol. 8, pp. 214417–214428, 2020.
- [12] Y. Liu, H. Luo, and T. Zhang, "Efficient machine learning via hardware–software co-design," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 40, no. 12, pp. 2390–2402, Dec. 2021.
- [13] K. Henderson and S. Kim, "The role of renewable energy in sustainable AI development," *Proc. IEEE Green Technol. Conf. (GreenTech)*, Denver, CO, USA, pp. 1–6, 2021.
- [14] J. Xu, W. Li, and J. Chen, "Optimizing AI energy consumption through model compression and pruning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 4, pp. 1426–1438, Apr. 2022.
- [15] T. Patterson, M. So, and Q. V. Le, "Efficiency improvements in AI computation: From hardware to algorithms," *IEEE Trans. Sustain. Comput.*, vol. 8, no. 3, pp. 540–551, Jul.–Sep. 2023.
- [16] S. Lee, J. Park, and K. Choi, "Measuring and mitigating carbon impact in federated learning systems," *IEEE Internet Things J.*, vol. 10, no. 4, pp. 3224–3235, Feb. 2023.
- [17] M. Khan, D. Patel, and A. R. Singh, "AI sustainability: Modeling energy-performance trade-offs," *IEEE Trans. Artif. Intell.*, vol. 5, no. 1, pp. 58–70, Jan. 2024.
- [18] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green AI," *Commun. ACM*, vol. 63, no. 12, pp. 54–63, 2020.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details