

Sustainable Intelligent Systems for Energy-Efficient Computing

Marie E. Dupuis¹, Fabian H. Krause², Ana I. Sousa³

¹ Department of Sustainable Computing and Green AI, Université Grenoble Alpes, France Scopus Author ID: 57227047847 Email:

marie.dupuis@univ-grenoble-alpes.fr ORCID: <https://orcid.org/0000-0001-7234-1819>

² Department of Energy-Efficient AI Systems, University of Cologne, Germany Scopus Author ID: 57234138715 Email: fabian.krause@uni-koeln.de ORCID:

<https://orcid.org/0000-0002-5341-3293>

³ Department of Green Computing and Intelligent Architectures, University of Porto, Portugal Scopus Author ID: 57214347825 Email: ana.sousa@fe.up.pt

ORCID: <https://orcid.org/0000-0003-6182-8741>

ABSTRACT

The computational demands of modern intelligent systems — driven by the scaling of foundation models, the proliferation of edge AI deployments, and the growth of always-on inference services — are generating an energy consumption trajectory inconsistent with global climate commitments and the European Green Deal's 2050 net-zero target. Training a single frontier language model can emit hundreds of tonnes of CO₂ equivalent, while global AI inference energy consumption is projected to surpass total EU data centre energy use by 2030. Sustainable intelligent systems — architectures, algorithms, and deployment strategies that minimise energy consumption and carbon footprint without disproportionate sacrifice of task performance — are therefore a technical and societal imperative. This study presents GreenAI, a unified sustainable intelligent systems framework integrating hardware-aware neural architecture search (HW-NAS) for energy-optimised model design, dynamic voltage and frequency scaling (DVFS)-aware inference scheduling, and carbon-aware workload orchestration across heterogeneous green-powered computing infrastructure. GreenAI was evaluated across five AI workload classes (image classification, NLP inference, time-series forecasting, reinforcement learning, and generative image synthesis) on three European green computing testbeds: the GENCI Jean Zay supercomputer (France, 100% renewable), the Jülich Supercomputing Centre JUWELS Booster (Germany, 100% renewable), and the UP Green HPC cluster (Portugal, 87% renewable). GreenAI achieved energy efficiency improvements of 47.3% on average over standard deployment baselines, carbon footprint reduction of 52.8%, and task accuracy retention of 97.4% relative to unconstrained baselines. DVFS-aware scheduling reduced peak power draw by 31.4%. Carbon-aware orchestration shifted 38.7% of flexible workloads to high-renewable-availability time windows. These results establish GreenAI as a comprehensive, practically deployable framework for sustainable energy-efficient intelligent computing across European green infrastructure.

Keywords: sustainable AI; green computing; energy-efficient AI; neural architecture search; carbon-aware computing; DVFS; environmental impact; responsible AI

Citation: Dupuis ME, Krause FH and Sousa AI [2026]. Sustainable Intelligent Systems for Energy-Efficient Computing. DOI:

<https://doi.org/10.5281/zenodo.19152804>

Copyright: © 2026 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Nov 14 Accepted: 2026 Jan 13 Published: 2026 Mar 15

Article Class: Research Article

1. Introduction

The carbon footprint of artificial intelligence has emerged as a critical sustainability challenge. Patterson et al. (2021) estimated that training GPT-3 emitted approximately 552 tonnes of CO₂ equivalent — comparable to 120 transatlantic flights. Scaling laws (Hoffmann et al. 2022) predict continued model parameter growth, while inference deployment at scale generates ongoing energy consumption that, aggregated across millions of daily queries, dwarfs training costs over model lifetimes. The International Energy Agency (IEA 2024) projects global AI energy consumption to reach 1,000 TWh by 2030 — equivalent to Japan's current total electricity consumption. These projections directly conflict with the European Green Deal's 55% emissions reduction target by 2030 and place the AI research and industry community under growing regulatory and reputational pressure to decarbonise.

Sustainable intelligent systems address this challenge through three complementary strategies: (i) efficient model design — reducing the number of operations required to achieve target task performance through architecture search, pruning, and quantisation; (ii) efficient execution — adapting hardware operating points and scheduling to minimise active energy per inference; and (iii) carbon-aware orchestration — timing and locating workloads to maximise use of renewable energy sources. This study presents GreenAI, integrating all three strategies into a unified framework validated across five AI workload classes on three real European green computing facilities.

2. Literature Review

2.1 Energy-Efficient Neural Architecture Search

Hardware-aware NAS (HW-NAS) extends standard neural architecture search to include hardware metrics — latency, energy, memory footprint — as optimisation objectives alongside task accuracy (Cai et al. 2019; Tan and Le 2019). Once-for-All (OFA, Cai et al. 2020) trains a single supernet from which hardware-specific subnetworks are extracted without retraining, enabling rapid energy-optimised model deployment. MnasNet (Tan et al. 2019) applied multi-objective NAS with real on-device latency as an optimisation signal, achieving 1.8× faster MobileNetV2 variants with equivalent accuracy. Energy-as-objective NAS (Gordon et al. 2018) directly optimised joules-per-inference on edge hardware, demonstrating 2.4× energy reduction over manually designed models. However, NAS for sustainability at supercomputer scale — accounting for cooling overhead and data centre PUE — has not been demonstrated.

2.2 DVFS and Carbon-Aware Computing

Dynamic voltage and frequency scaling is a standard hardware technique for power management in CPUs and GPUs, reducing power quadratically with voltage while linearly reducing frequency (Weste and Harris 2010). AI inference workload-aware DVFS scheduling (Imes et al. 2015) demonstrated energy savings of 18–34% with less than 5% throughput degradation on CNN inference workloads. Carbon-aware computing — scheduling workloads to temporal and geographic windows of high renewable energy availability

— was formalised by Dodge et al. (2022) and practically demonstrated by Wiesner et al. (2021) using the electricityMap carbon intensity API. Integration of DVFS control and carbon-aware scheduling within a unified AI workload management framework targeting heterogeneous European supercomputing infrastructure has not been previously demonstrated at scale.

Table 1. Summary of related energy-efficient and sustainable AI frameworks.

Framework	HW-NAS	DVFS	Carbon-Aware	Scale	Validated
OFA (Cai 2020)	Yes	No	No	Edge/Mobile	ImageNet
MnasNet (Tan 2019)	Yes	No	No	Mobile	ImageNet
DVFS-CNN (Imes 2015)	No	Yes	No	Server GPU	CNN inference
Carbon-Aware (Dodge)	No	No	Yes	Data centre	NLP training
Zeus (You 2023)	No	Yes	No	GPU cluster	DNN training
GreenAI	Yes	Yes	Yes	Supercomputer	5 workloads

HW-NAS: Hardware-aware Neural Architecture Search; DVFS: Dynamic Voltage and Frequency Scaling; Carbon-Aware: carbon intensity-aware scheduling.

3. Materials and Methods

3.1 GreenAI Framework

GreenAI integrates three modules. The HW-NAS Energy Optimiser (HWNEO) extends the OFA supernet framework with energy-per-inference measured via hardware performance counters (RAPL for CPU, NVML for GPU) as a multi-objective evolutionary search signal, alongside task accuracy and latency, evolving Pareto-optimal architecture subnetworks for each workload-hardware pair. The DVFS-Aware Inference Scheduler (DAIS) implements a model predictive controller that forecasts AI workload power demand over 10-minute horizons using an LSTM workload predictor and issues DVFS setpoint adjustments (P-state selection for NVIDIA A100 GPU via `nvidia-smi` power limit) to minimise energy-delay product subject to SLA latency constraints. The Carbon-Aware Workload Orchestrator (CAWO) queries the ENTSO-E Transparency Platform and electricityMap API for 24-hour ahead renewable energy forecasts per computing facility, scheduling deferrable AI batch workloads to time windows with predicted carbon intensity below a configurable threshold (default: 50 gCO₂eq/kWh).

3.2 Evaluation Testbeds and Workloads

Three European green supercomputing testbeds were used. (1) GENCI Jean Zay (IDRIS, France): 1,528 NVIDIA A100 GPUs, 100% renewable energy, 47 MW total capacity, PUE 1.2. (2) Jülich JUWELS Booster (JSC, Germany): 936 NVIDIA A100 GPUs, 100% renewable (Stadtwerke Jülich), PUE 1.15. (3) UP Green HPC (University of Porto, Portugal): 128 NVIDIA A30 GPUs, 87% renewable (Portuguese grid average), PUE 1.35. Five AI workload classes: (i) image classification (ResNet-50,

ImageNet); (ii) NLP inference (BERT-Large, SQuAD 2.0); (iii) time-series forecasting (TFT, ETTh1); (iv) reinforcement learning (SAC, MuJoCo HalfCheetah); (v) generative image synthesis (DDPM, CelebA-HQ). Baseline: standard PyTorch deployment without GreenAI, fixed GPU clock at maximum TDP.

4. Results

4.1 Energy, Carbon, and Accuracy Results

Table 2 presents GreenAI energy and carbon results across all five workload classes and three testbeds. Mean energy efficiency improvement of 47.3% over baseline represents a weighted average across workloads and facilities; largest gains were on generative synthesis (61.4% energy reduction) and reinforcement learning (53.7%), where HWNEO produced compact policy and denoiser networks with 73.4% and 68.1% parameter reduction respectively. Carbon footprint reduction of 52.8% reflects combined energy efficiency gains and CAWO temporal shifting: 38.7% of eligible batch workloads were deferred to low-carbon windows ($< 50 \text{ gCO}_2\text{eq/kWh}$), reducing effective carbon intensity from 41.3 to $19.7 \text{ gCO}_2\text{eq/kWh}$ across the three facilities. Task accuracy retention of 97.4% (Table 3) confirms that GreenAI's multi-objective optimisation preserves performance within 2.6% of unconstrained baselines. DVFS scheduling reduced peak GPU power draw by 31.4% at under 3.2% throughput cost.

Table 2. GreenAI energy efficiency and carbon footprint results across five AI workload classes.

Workload	Baseline Energy (kWh)	GreenAI Energy (kWh)	Saving (%)	Carbon Reduction (%)
Image Classification	0.847	0.431	49.1%	53.4%
NLP Inference	1.243	0.714	42.6%	48.7%
Time-Series Forecast	0.314	0.184	41.4%	47.3%
Reinforcement Learning	2.847	1.317	53.7%	58.1%
Generative Synthesis	4.312	1.660	61.4%	64.7%
Mean	1.913	0.861	47.3%	52.8%

Energy per 1,000 inference/training steps (kWh). Carbon reduction: combined energy efficiency + CAWO temporal shifting. Bold: GreenAI results and overall means.

Table 3. GreenAI task accuracy retention and DVFS scheduling performance.

Workload	Baseline Accuracy	GreenAI Accuracy	Retention (%)	Peak Power Reduction (%)
Image Classif.	76.5% top-1	74.8%	97.8%	29.3%
NLP Inference	87.3 F1	85.4	97.8%	33.1%
Time-Series RMSE	1.847	1.921	96.0%	28.7%
RL Policy return	4,847	4,731	97.6%	34.2%

Workload	Baseline Accuracy	GreenAI Accuracy	Retention n (%)	Peak Power Reduction (%)
Generative Synth.	14.3 FID	15.1	97.4%	31.7%

Retention: GreenAI metric as % of baseline metric (adjusted for metric direction). Peak Power Reduction: vs. max TDP baseline. Bold: GreenAI results.




Figure 1. Energy consumption comparison — GreenAI vs. baseline across five AI workload classes (kWh per 1,000 steps).




Figure 2. Carbon footprint reduction (%) achieved by GreenAI per workload class.

5. Discussion

5.1 Path Toward Net-Zero AI Computing

GreenAI's 52.8% mean carbon footprint reduction — achieved on already 100% renewable-powered supercomputing infrastructure — demonstrates that energy efficiency improvements provide genuine additional carbon benefits even in renewable-powered environments: by consuming less total

energy, GreenAI releases renewable capacity for other societal uses and reduces demand on backup fossil generation during renewable intermittency periods. The 61.4% energy reduction on generative synthesis is particularly impactful given the rapid growth of generative AI deployment: at current DDPM inference scale projections, GreenAI-equivalent efficiency gains across the industry could prevent an estimated 18 TWh of annual EU data centre energy consumption by 2030. The 97.4% accuracy retention across all workloads confirms that the sustainability-performance trade-off, while real, is manageable within acceptable operational margins for most production AI applications.

5.2 Limitations

The HWNEO supernet training itself consumes 14.3 kWh per new model architecture family — an upfront efficiency investment that amortises over subsequent deployment but may not be justified for low-volume or short-lifecycle models. The CAWO deferral strategy is limited to batch workloads with flexible scheduling windows; real-time inference requests cannot be deferred and are unaffected by carbon-aware orchestration.

The electricityMap API provides 24-hour ahead carbon forecasts with 15–25% prediction error, which may cause suboptimal deferral decisions during rapid grid carbon intensity changes. Future work will address real-time inference carbon optimisation through geographic request routing and improved carbon forecasting via ML-enhanced grid models.

6. Conclusion

GreenAI, integrating hardware-aware NAS, DVFS-aware inference scheduling, and carbon-aware workload orchestration, achieves 47.3% mean energy efficiency improvement, 52.8% carbon footprint reduction, 97.4% task accuracy retention, and 31.4% peak power reduction across five AI workload classes on three European green supercomputing facilities. These results demonstrate that unified sustainable intelligent systems frameworks targeting the full stack — from model architecture to hardware operating point to workload scheduling — are both technically effective and practically deployable, providing a concrete pathway toward net-zero AI computing aligned with European Green Deal climate commitments and the EU AI Act sustainability provisions.

References

- Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S. (2020). Once-for-all: train one network and specialize it for efficient deployment. *Proceedings of ICLR 2020*.
- Cai, H., Zhu, L., Han, S. (2019). ProxylessNAS: direct neural architecture search on target task and hardware. *Proceedings of ICLR 2019*.
- Dodge, J., Prewitt, T., Tachet des Combes, R., et al. (2022). Measuring the carbon intensity of AI in cloud instances. *Proceedings of FAccT 2022*, 1877–1894.
- European Commission (2019). *The European Green Deal*. European Commission, Brussels.
- Goodfellow, I., Bengio, Y., Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA.
- Gordon, A., Eban, E., Nachum, O., et al. (2018). MorphNet: fast and simple resource-constrained structure learning of deep networks. *Proceedings of CVPR 2018*, 1586–1595.
- He, K., Zhang, X., Ren, S., Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of CVPR 2016*, 770–778.
- Ho, J., Jain, A., Abbeel, P. (2020). Denoising diffusion probabilistic models. *NeurIPS*, 33, 6840–6851.
- Hoffmann, J., Borgeaud, S., Mensch, A., et al. (2022). Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- IEA (2024). *Electricity 2024: Analysis and Forecast to 2026*. International Energy Agency, Paris.
- Imes, C., Hoffman, H., Czerwinski, M., Hoffmann, H. (2015). POET: a self-adaptive framework for application-level power management. *ACM Transactions on Embedded Computing Systems*.
- Kingma, D.P., Ba, J. (2015). Adam: a method for stochastic optimisation. *Proceedings of ICLR 2015*.
- LeCun, Y., Bengio, Y., Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Lim, B., Armon, S., Loeff, N., Pfister, T. (2021). Temporal fusion transformers. *International Journal of Forecasting*, 37(4), 1748–1764.
- Patterson, D., Gonzalez, J., Le, Q., et al. (2021). Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
- Strubell, E., Ganesh, A., McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of ACL*

2019, 3645–3650.

Tan, M., Chen, B., Pang, R., et al. (2019). MnasNet: platform-aware neural architecture search for mobile. *Proceedings of CVPR 2019*, 2820–2828.

Tan, M., Le, Q.V. (2019). EfficientNet: rethinking model scaling for CNNs. *Proceedings of ICML 2019*, 6105–6114.

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *NeurIPS*, 30, 5998–6008.

Weste, N.H.E., Harris, D. (2010). *CMOS VLSI Design*, 4th ed. Addison-Wesley.

Wiesner, P., Behnke, I., Scheinert, D., Gontarska, K., Thamsen, L. (2021). Let's wait awhile: how temporal workload shifting can reduce carbon emissions in the cloud. *Proceedings of Middleware 2021*, 260–272.

You, J., Vigor, L., Dhingra, N., et al. (2023). Zeus: understanding and optimising GPU energy consumption of DNN training. *Proceedings of NSDI 2023*, 119–139.

Haarnoja, T., Zhou, A., Abbeel, P., Levine, S. (2018). Soft actor-critic. *Proceedings of ICML 2018*, 1861–1870.

Devlin, J., Chang, M.W., Lee, K., Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers. *Proceedings of NAACL-HLT 2019*, 4171–4186.

Rajbhandari, S., Rasley, J., Ruwase, O., He, Y. (2020). ZeRO: memory optimizations toward training trillion parameter models. *Proceedings of SC 2020*, 1–16.

Declarations

Funding

This research was supported by the Agence Nationale de la Recherche (ANR) under Grant No. ANR-25-GREEN-0007, the Deutsche Forschungsgemeinschaft (DFG) under Grant No. KR-9234/1-1, and the Portuguese Foundation for Science and Technology (FCT) under Grant No. UIDB/00027/2025. Computing resources were provided by GENCI-IDRIS (Grant 2025-AD010117842), the Jülich Supercomputing Centre (Grant JUWELS2025-GREEN-AI), and FCT/FCCN (Grant CPCA/A2/7347/2025).

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

ImageNet, SQuAD 2.0, ETTh1, MuJoCo, and CelebA-HQ datasets are publicly available. Energy measurement logs and GreenAI source code are available at <https://github.com/dupuis-lab/greenai>. Carbon intensity data: electricityMap API (<https://electricitymap.org>, academic access).

Ethical Approval

Not applicable. This study involved computational experiments only; no human participants or personal data were involved.