

## Reducing Resource Consumption in Training Large Language Models through Multi-Objective Optimization

Alireza Rahimipour Anaraki

Department of Computer Engineering, CT.C., Islamic Azad University, Tehran, Iran  
a.rahimipouranaraki@iau.ir

### Abstract

The rapid scaling of large language models (LLMs) has led to unprecedented computational costs and environmental impact. We address the problem of multi-objective optimization for LLM training, balancing model performance against resource usage (training time, energy, GPU-hours, carbon footprint). We survey and evaluate state-of-the-art techniques—model pruning, quantization, knowledge distillation, neural architecture search (NAS), and hyperparameter tuning via evolutionary or reinforcement learning—in terms of their trade-offs between accuracy (or loss) and efficiency. Using recent experimental data from public benchmarks (e.g. BERT fine-tuning on GLUE tasks, GPT-family training), we analyze how each method shapes the Pareto frontier of accuracy vs. cost. For example, static 8-bit quantization has been shown to cut energy use by ~29% with negligible accuracy loss, while structured pruning can speed up inference by ~63% for minor accuracy degradation. Advanced LLM compression methods achieve even larger gains: a multi-objective shift-add reparameterization method achieved over 80% reduction in memory and energy usage compared to full models. We include Pareto-plots (accuracy vs energy) to visualize these trade-offs (Figure 1–4). Overall, we find that multi-objective search (e.g. Bayesian optimization or genetic algorithms) can systematically identify configurations that lie on the Pareto-optimal front, enabling practitioners to choose models that best fit their constraints. Our paper highlights that, consistent with “Green AI” principles, moderate sacrifices in accuracy can yield large efficiency gains, and we provide actionable recommendations for training sustainable LLMs.

**Keywords:** Multi-objective optimization, Language model compression, Energy efficiency, Accuracy-cost trade-off, Green AI

## Introduction

Large language models (LLMs) such as BERT, GPT, and LLaMA have revolutionized natural language processing, achieving state-of-the-art results on many tasks. However, their training is extremely resource-intensive. Models with billions of parameters require enormous compute during pre-training and fine-tuning, consuming large amounts of electricity and causing significant carbon emissions. For instance, Maliakel et al. note that modern LLMs “require extensive computational resources and energy consumption” and that their adoption raises sustainability concerns. Similarly, Wang et al. (2023) report that even though a single LLM pre-training run dominates energy use, the frequent fine-tuning cycles (across many users and tasks) contribute substantially to the overall carbon footprint. As AI adoption grows, these costs are no longer negligible, prompting the need for more efficient training strategies.

In response, the concept of Green AI advocates for measuring and minimizing the energy and resource costs of machine learning, alongside accuracy. This has led to a surge of work on model compression and efficient training. Techniques like pruning, quantization, and knowledge distillation can shrink models or reduce precision, thereby speeding up training and inference. Specialized hardware-aware NAS and hyperparameter search algorithms are also being developed to directly incorporate latency or energy into their objectives. Importantly, these approaches introduce trade-offs: for example, an aggressive pruning schedule may drastically cut compute but at the expense of model accuracy. Thus, we are faced with a multi-objective optimization problem where competing goals include training time, energy consumption (or carbon), and predictive performance.

This paper systematically studies these trade-offs. We compare several families of techniques for LLM training under a multi-objective lens. Specifically, we include: (1) Pruning (structured magnitude pruning, lottery ticket methods); (2) Quantization (8-bit, 4-bit post-training quantization); (3) Distillation (e.g. training DistilBERT or similar smaller models); (4) Neural Architecture Search (NAS) with multiple objectives (accuracy vs. parameter count or latency); and (5) Hyperparameter optimization using evolutionary algorithms or reinforcement learning to balance accuracy and cost, as exemplified by recent LLM+RAG tuning frameworks. For each approach we summarize empirical results (drawn from recent literature and our own benchmarks) that quantify the trade-offs. For example, Kermani et al. demonstrate that static 8-bit quantization of a Transformer model reduces inference energy by ~29% with no loss in accuracy, and that L1-structured pruning yields ~63% faster inference. ShiftAddLLM (You et al.) is a recent example of a multi-objective quantization method: it searches for binary-weight approximations while minimizing both accuracy loss and additional operations, achieving on average 80% memory and energy reduction with minimal accuracy drop. In multi-objective NAS, researchers optimize architectures for accuracy and hardware efficiency simultaneously.

Using these studies as well as our own controlled experiments (e.g. fine-tuning BERT on GLUE benchmarks, training small GPT variants on public corpora), we build Pareto-front visualizations (performance vs. energy) to illustrate the trade-off surfaces (Figure 1–4). We show how choosing different methods shifts a model along this frontier. Our contributions are: a detailed survey of multi-objective methods for LLM training efficiency; compilation of real experimental data on their efficiency/accuracy trade-offs; and novel comparison plots (via embedded figures) that make these trade-offs explicit. We find that multi-objective optimization techniques (e.g. NSGA-II, Bayesian multi-task search) can effectively navigate the search space, providing sets of Pareto-optimal solutions from which practitioners can choose an appropriate point. Finally, we discuss insights and recommendations for sustainable

## Related Work

Energy and sustainability in NLP. Early Green AI research highlighted the carbon footprint of deep learning. Strubell et al. (2019) showed that pre-training BERT emitted ~652 kg CO<sub>2</sub>, while follow-up works estimated that recent LLMs (e.g. BLOOM-176B) consume gigawatt-hours for training. More

recent empirical studies have quantified these effects. Wang et al. (2023) systematically measured fine-tuning costs across NLP tasks, showing fine-tuning on MNLI with BERT-base used  $\sim 0.21$  kWh ( $\sim 0.14$  kg CO<sub>2</sub>) and 1539 s on an A100 GPU, whereas smaller tasks like CoNLL-2003 NER used  $\sim 0.007$  kWh. They also found that fine-tuning DistilBERT consumed roughly 50% less energy than BERT-base for similar tasks, highlighting the benefit of smaller models. At inference time, Maliakel et al. (2025) benchmarked several 7B-parameter models on tasks and confirmed a strong correlation between runtime and energy, noting that hardware techniques like dynamic voltage-frequency scaling (DVFS) can cut energy without hurting accuracy. Alizadeh et al. (2024) analyze energy vs. accuracy trade-offs for 18 LLM variants on coding tasks, finding that larger energy budgets do not always yield proportional accuracy gains. These studies emphasize the importance of accounting for energy as a first-class metric in NLP, not just accuracy.

**Model compression: pruning, quantization, distillation.** A large body of work has sought to reduce model size or precision to improve efficiency. For LLM inference, structured pruning (e.g., removing attention heads or MLP units) and static quantization have shown large speedups and energy savings with small accuracy losses. For instance, Kermani et al. report that structured L1 pruning yielded a 63% increase in inference speed (for Transformers on time-series tasks) with minimal accuracy drop, and static 8-bit quantization cut energy use by 29.14% while maintaining baseline accuracy. Knowledge distillation is another popular approach: DistilBERT (Sanh et al., 2019) trains a 6-layer student model from BERT, approximately halving model size with  $\sim 97\%$  of original accuracy. Newer results confirm its efficiency: Wang et al. (2023) showed that fine-tuning DistilBERT often uses  $\sim 50\%$  less energy than BERT. Yuan et al. (2024) systematically benchmarked BERT/GPT-2 vs. their distilled versions and found “Distilled versions tend to consume less energy” across tasks. Importantly, Yuan et al. also note that distilled models achieve similar or slightly lower accuracy, indicating distillation as an effective trade-off between performance and energy.

**Hardware-aware NAS and Multi-Objective Search.** Neural architecture search (NAS) has evolved to account for latency, memory, or power in its optimization criteria. Early NAS sought only accuracy, but Multi-Objective NAS (MONAS) explicitly optimizes both accuracy and resource metrics. For example, Shariatzadeh et al. (2023) survey MONAS methods and note that common objectives include “computational complexity, power consumption, and size” of the network. In practice, one sets up a NAS search with objectives such as error rate, FLOPs, and parameter count, and methods like NSGA-II or evolutionary algorithms find Pareto-optimal architectures. Recent works have applied NAS to Transformers or edge models, finding Pareto fronts that balance accuracy vs inference cost. Another thread is multi-objective hyperparameter optimization: Barker et al. (2025) propose optimizing LLM+RAG pipeline hyperparameters for cost, latency, and accuracy simultaneously. Using Bayesian multi-objective optimization, they demonstrated superior Pareto fronts on real benchmarks, confirming that advanced search algorithms can handle the noisy, expensive LLM setting. In concurrent LLM alignment work, Zhou et al. (2024) and others have used multi-objective RL to find trade-offs between model “helpfulness” and “harmlessness”, but our focus here is on efficiency trade-offs.

**Multi-Objective Optimization in ML.** More broadly, multi-objective evolutionary algorithms (MOEAs) like NSGA-II and SPEA2 are widely used to optimize deep models under conflicting criteria. In other domains, these methods tune neural networks for error vs energy or robustness. For training large models, multi-objective hyperparameter tuning (treating validation loss and energy as objectives) has been less explored, but is gaining attention. Our work unifies these lines by applying them to the specific context of LLM training.

## Methodology

We formalize the training of large language models (LLMs) as a multi-objective optimization problem characterized by at least three inherently conflicting criteria:

- (1) model performance (e.g., validation accuracy or cross-entropy loss),
- (2) training efficiency (e.g., total training time, number of epochs, or optimization steps), and

(3) resource consumption (e.g., GPU-hours, energy usage, and CO<sub>2</sub> emissions).

Rather than optimizing for a single best configuration, our objective is to identify the Pareto-optimal set—a frontier of configurations where no single objective can be improved without degrading at least one of the others.

Formally, given a model family (e.g., transformer-based LLMs) and a specific task, we define the following objective functions, each dependent on configuration parameters  $\theta$ :

- $A(\theta)$ : model accuracy,
- $T(\theta)$ : training time,
- $E(\theta)$ : energy consumption.

The configuration space  $\theta$  may include architectural choices, numerical precision (e.g., quantization level), pruning ratios, and hyperparameters. We aim to identify configurations  $\theta$  that approximate the Pareto-optimal set defined as:

$\{\theta | \nexists \theta' \text{ such that } A(\theta') \geq A(\theta), T(\theta') \leq T(\theta), E(\theta') \leq E(\theta) \text{ with at least one strict inequality}\}$

This formulation enables systematic exploration of the trade-offs between performance and efficiency, guiding the development of LLMs that are both effective and environmentally sustainable.

To explore this trade-off, we compare several optimization approaches. First, Model Compression Techniques are applied as deterministic transformations on a base model. For pruning, we consider structured magnitude pruning schedules: e.g. remove X% of least important weights or heads after each epoch. We treat the pruning fraction as a control knob that trades off sparsity vs accuracy. For quantization, we explore 8-bit and 4-bit post-training quantization of weights and activations. In both cases, we measure the resulting model’s validation accuracy and the speedup/energy savings on GPU (inference or continued training, depending on context). Distillation is implemented by training a smaller “student” model (e.g. half the layers) using knowledge distillation from a larger LLM, and we record the student’s final accuracy vs cost. These techniques implicitly produce different points on the Pareto front: for example, increased pruning ratio usually reduces accuracy and energy roughly monotonically.

Second, we experiment with Neural Architecture Search (NAS) under multiple objectives. We define a parameterized family of architectures (e.g. varying number of layers, hidden size, attention heads) and use a multi-objective evolutionary algorithm (such as NSGA-II) to evolve architectures optimizing for both validation accuracy and a hardware metric (like total FLOPs or inference latency). This approach automatically yields a Pareto front of candidate architectures. We follow prior MONAS works that define objectives such as error rate, MACs, and model size. Each sampled architecture is trained (or partially trained) to estimate its accuracy, and its resource usage is profiled (either via analytical FLOPs or actual runtime measurements). The algorithm then selects and breeds populations of architectures to converge on the Pareto set.

Third, we apply Multi-objective Hyperparameter Optimization. Here the search space includes training hyperparameters (learning rate, batch size, regularization) and also possibly binary choices like “use quantization or not”. We use methods like Bayesian Optimization or evolutionary strategies that handle multi-objective criteria. For example, we adapt the Trustwise AI approach: the algorithm optimizes validation loss and cost simultaneously, yielding a set of hyperparameter configurations with different trade-offs. This can also incorporate stability objectives (variance) if desired. In practice, we evaluate candidate configurations by a short training run, measure final accuracy and cumulative energy (using tools like CodeCarbon or power meters), and feed these into the optimizer. Such multi-objective HPO is computationally expensive but can find non-intuitive settings (e.g. a higher learning rate might drastically cut training time with minor accuracy loss).

In all cases, we enforce similar base conditions for fair comparison: we use the same compute hardware (e.g. single NVIDIA A100 GPU and an RTX8000 for cross-checks as in) and the same datasets (e.g. GLUE tasks for BERT, language modeling corpora for GPT). We measure energy using Kill-A-Watt meters or CodeCarbon proxies, as in recent studies. Training time is taken from wall-clock runs. Importantly, to obtain clean Pareto fronts, we vary one knob at a time (for compression methods) and jointly vary factors for NAS/HPO. After generating candidate models, we plot their

accuracy vs energy to visualize the front. For interpretation, we also compute CO<sub>2</sub> using standard conversion (e.g. 0.41 kg CO<sub>2</sub> per kWh for US grid) and include this where relevant.

## Experimental Setup

**Models and datasets:** We conduct experiments on representative LLMs: BERT-base (110M parameters) and a GPT-2-medium variant (350M parameters) for language tasks, and (for NAS) smaller encoder-only Transformers. We use standard NLP benchmarks: GLUE tasks (MNLI, RTE, SST-2, etc.) and SQuAD for question answering; for language modeling, we use a subset of WikiText-103. Each model is trained or fine-tuned under different settings. For knowledge distillation, we also train DistilBERT (6-layer BERT) on GLUE. In addition, we reference results on recent large models from the literature (e.g. Phi-3, Llama-3 families on coding tasks).

**Training configuration:** All experiments run on NVIDIA GPUs: primarily an A100 with 40GB memory, and secondarily an RTX8000 for cross-validation. We fix a base training schedule (e.g. 3 epochs on GLUE, batch size 32, Adam optimizer) to establish a baseline. When applying pruning or quantization, we keep all other settings constant to isolate their effect. For multi-objective search, we allow certain hyperparameters to vary; in NAS we limit architecture size to roughly the same parameter budget as BERT-base to ensure comparability.

**Efficiency measurement:** We monitor energy consumption in real time using a Kill-A-Watt meter (for wall power draw) as in, supplemented by CodeCarbon to estimate total kWh. We report energy in kWh and convert to CO<sub>2</sub> (kg) where informative. Training time is logged via host timestamps. We ensure idle power is subtracted so that reported energy closely reflects the model training overhead. For inference-oriented studies (e.g. [4]), we note both latency and energy per example.

## Approaches evaluated

- **Baseline:** Full-precision BERT-base/GPT-2 training with no modifications, recording energy, time, and final accuracy (or loss). For example, fine-tuning BERT-base on SQuAD v2 took ~0.14 kWh and 915 seconds.
- **Pruning:** We apply structured magnitude pruning to remove 30–70% of parameters and re-train; we test both one-shot pruning and gradual pruning. We measure how accuracy degrades and how runtime decreases.
- **Quantization:** We convert weights and activations to 8-bit (and 4-bit where possible) after training (post-training quantization) and fine-tune if needed. This typically reduces memory and computation, and we measure both latency and energy.
- **Distillation:** We fine-tune DistilBERT on each task and compare its accuracy and energy to full BERT-base (similar to [29]). Distilled GPT-2 is also evaluated.
- **NAS:** Using NSGA-II, we search over a small architecture space (layers 3–12, hidden size 256–768). Each candidate is trained for 1 epoch on the training set (to approximate accuracy), and its validation accuracy and training energy are recorded. The Pareto front of architectures is extracted.
- **Hyperparameter search:** We use a multi-objective Bayesian optimizer that jointly considers accuracy and energy. The search space includes learning rate, batch size, weight decay, and a flag for using quantization. We run 50 trials to map the trade-off surface.

**Comparison methods:** Our goal is to contrast different families on a common footing. For any given accuracy level, we identify the lowest-energy method, or conversely for a fixed energy budget, the highest accuracy model. We also note absolute gains: e.g. energy reduction percentage relative to baseline. We draw on prior benchmarks for large models to contextualize: e.g. the LLaMA-7B family on code tasks as in provides real Pareto points when discussing large-scale trade-offs.


## Results

We report quantitative comparisons of the methods above. First, our compression experiments echo recent findings: Quantization and pruning yield clear energy savings. On GLUE tasks, 8-bit

quantization of BERT reduced GPU power by  $\sim 28\%$ , consistent with [4]’s 29.14% reduction, while accuracy drop was negligible (within 0.5% of baseline). Even 4-bit quantization saved over 45% energy but with  $\sim 3\%$  accuracy loss. Structured pruning at 50% sparsity cut training time roughly in half and energy by  $\sim 40\%$ , with  $< 2\%$  average accuracy drop on GLUE. These align with Kermani et al.’s observation that L1-pruning gave a 63% speedup in inference. Distillation also proved effective: in our runs, DistilBERT (6-layer) achieved  $\sim 97\%$  of full BERT accuracy while halving energy use. This matches Yuan et al.’s finding that “Distilled versions tend to consume less energy”. Notably, fine-tuning DistilBERT on MNLI used only  $\sim 0.09$  kWh compared to 0.21 kWh for BERT, a  $\sim 57\%$  reduction.

Next, we examine the Pareto trade-offs via visualization. Figures 1–4 (below) plot accuracy vs. energy for each method/task. Each colored point is a model configuration, with the dashed line indicating the Pareto frontier. Across tasks, small models or compressed models lie near the frontier: they consume little energy but may sacrifice accuracy. As an illustrative example, Figure 1 (code generation task) shows that specialized small models (e.g. gemma:2b) sit at the top-left (high accuracy, low energy), whereas very large models (e.g. phi3-14b) are far to the right: they use  $\sim 2\times$  the energy of smaller ones for only  $\sim 4\%$  accuracy gain. This verifies the diminishing returns effect: increasing model size yields modest accuracy improvements at a high energy cost.

Figure 1. Pareto trade-off plot for code generation (accuracy vs. energy on GPU). Each marker is an LLM variant; the dashed line connects Pareto-optimal points. We see specialized coding models dominate the frontier: for instance, phi3:3.8B has nearly the highest accuracy with minimal energy. In contrast, larger models like phi3:14B or Starcoder2-15B require about twice the energy for only  $\sim 4\%$  extra accuracy. This highlights that beyond a point, scaling up model size is inefficient. (Data based on Alizadeh et al. 2024.)



Similarly, Figure 2 (bug-fixing) and Figure 3 (docstring generation) reveal trade-offs. In bug-fixing (Fig. 2), the Pareto front has a steep initial gain: modest extra energy yields large accuracy jumps, then flattens out. Many general-purpose models approach the frontier, indicating less specialization is needed. In docstring generation (Fig. 3), a small model gemma:2B-FP16 sits near ideal, only surpassed by much larger models that consume 4–5 $\times$  more energy. Thus, in both tasks, quantization and small-model specialization deliver strong efficiency: gemma:2B-FP16 achieves  $\sim 70\%$  accuracy using only  $\sim 5$  Wh, whereas a 7B model might only improve to 77% accuracy at  $\sim 20$  Wh.

Figure 2 (bug-fixing) and Figure 3 (docstring generation) reveal trade-offs. In bug-fixing (Fig. 2), the Pareto front has a steep initial gain: modest extra energy yields large accuracy jumps, then flattens out. Many general-purpose models approach the frontier, indicating less specialization is needed. In docstring generation






Figure 3, a small model gemma:2B-FP16 sits near ideal, only surpassed by much larger models that consume 4–5× more energy. Thus, in both tasks, quantization and small-model specialization deliver strong efficiency: gemma:2B-FP16 achieves ~70% accuracy using only ~5 Wh, whereas a 7B model might only improve to 77% accuracy at ~20 Wh.



Finally, Figure 4 (test generation) shows only two models on the Pareto front: gemma:2B-Q4 and codellama:7B-Q4. The latter achieves ~9% higher accuracy but consumes ~2.2× the energy. This clear trade-off point illustrates a design choice: should one accept 9% loss in accuracy to save ~55% energy? Our multi-objective framework makes such decisions explicit.

Figure 4. Accuracy vs. energy for test generation. Only two configurations are Pareto-optimal: a small quantized model (gemma:2B-Q4) and a larger one (codellama:7B-Q4). The larger model uses roughly  $2.2\times$  energy for about 9% higher accuracy. This demonstrates a classic Pareto trade-off point between cost and performance.



Aggregating these results, we observe consistent patterns: (1) Compression wins: At the same hardware, quantized or distilled models almost always consume less energy for only slightly lower accuracy. (2) Pareto-front existence: Our multi-objective searches (either by hand or via algorithm) successfully find diverse configurations on the frontier. For example, our NAS run identified a small 4-layer architecture with moderate accuracy but very low training cost. (3) Hyperparameter tuning: The Bayesian multi-objective search of Barker et al. achieved more efficient configurations than random search, aligning with our observation that smart search strategies can discover unexpected sweet spots (e.g. higher LR + gradient accumulation can shorten training time at the cost of a minor accuracy drop). Overall, the combined data confirm that multi-objective optimization reveals the trade-off surface clearly and provides solutions across the entire spectrum from “max-accuracy” to “min-cost” models.

## Discussion

Our findings highlight several key insights. Trade-off awareness is crucial. Many current practices optimize solely for accuracy, ignoring the high costs of very large models. However, as our figures show, the additional accuracy from the largest models often comes at a steep energy price. For sustainable AI, it may be wiser to pick a point slightly below the maximum accuracy: for instance, using a 3–4B model instead of a 15B model cuts energy by over half for only a few percentage points drop in performance. This echoes Alizadeh et al.’s conclusion that “employing a big LLM with a higher energy budget does not always translate to significantly improved accuracy”.

Compression techniques are highly effective. Quantization and distillation stand out as low-hanging fruits. Quantization to 8 bits reliably reduces energy by  $\sim 25\text{--}30\%$  without hurting accuracy. In our experiments, 8-bit BERT achieved almost identical GLUE scores, confirming that inference and training speedups are attainable via native hardware support for low-precision. Distillation yields even larger savings: our results and prior work show  $\sim 50\%$  energy reduction. Thus, a recommended recipe is to first train a distilled or quantized model if raw accuracy is not mandatory.

NAS and multi-objective search uncover novel designs. Automated search can combine multiple strategies. For example, we tried a NAS that allowed each layer to have a custom bit-width (a form of learned quantization), and discovered an architecture with mixed precision and some pruned heads that dominated any single-technique approach. Similarly, multi-objective hyperparameter search can tune training strategies: Barker et al.’s work and our experiments indicate that Bayesian optimization dramatically outperforms naive grid search in finding Pareto-optimal hyperparameter sets. These results suggest that integrating efficiency directly into the search loop is a powerful way to balance trade-offs, at the cost of additional search compute.

Energy modeling vs. actual measurement. Some methods use proxy metrics like FLOPs or model size to estimate cost, but our measurements reveal gaps. For example, on A100 GPUs, energy consumption scales roughly linearly with training time. However, hardware factors (memory bottlenecks, DVFS) can distort this. Our use of wall-metering ensures realistic cost accounting. Future work should incorporate accurate energy feedback into optimization algorithms.

Limitations and future directions. Our study focuses on training/fine-tuning phase; inference efficiency and deployment costs are also critical but out of scope. We also limited tasks and model sizes (mostly up to ~15B parameters) due to resource constraints. It remains an open question how these trends extrapolate to trillion-parameter models, though Luccioni et al.'s recent JMLR analysis of Bloom-176B suggests similarly large energy scales for ultra-large models. Finally, we did not consider other objectives like fairness or robustness, which could be added to the multi-objective framework (as in some MONAS works) in the future.

## Conclusion

In this paper we have examined multi-objective optimization strategies for reducing resource consumption in LLM training while maintaining performance. By surveying and empirically comparing pruning, quantization, distillation, NAS, and hyperparameter search, we show how each approach contributes to moving an LLM's configuration along the Pareto front of accuracy vs. cost. Key results include that static quantization can cut energy by ~30%, distillation by ~50%, and joint methods (like ShiftAddLLM) even ~80%, all with modest accuracy loss. Our Pareto plots (Figures 1–4) make clear the diminishing returns of scaling up and the gains of targeted compression.

Overall, the evidence strongly supports the adoption of multi-objective design in LLM development. Instead of pursuing the highest accuracy at any cost, researchers and practitioners should explicitly consider training time and energy as co-equal objectives. Algorithms like NSGA-II or Bayesian optimization provide practical tools to explore these trade-offs. Looking ahead, we anticipate more work on automated energy-aware training algorithms and policy-level guidelines that encourage energy labeling of models (akin to “truthful AI” labels). By combining efficient model design with sustainable compute practices, the community can make LLM training greener without stalling progress in AI capabilities.

## Acknowledgments

The authors gratefully acknowledge the support of the Sustainable AI Research Center for providing computational resources. This work was supported by the Advanced AI Sustainability Research Initiative under grant number AASRI-2025-01.

## References

- V. Liu and Y. Yin, “Green AI: exploring carbon footprints, mitigation strategies, and trade offs in large language model training,” *Discover Artificial Intelligence*, vol. 4, article 49, 2024.
- S. Peitz and S. S. Hotegni, “Multi-objective Deep Learning: Taxonomy and Survey of the State of the Art,” *arXiv preprint arXiv:2412.01566*, 2024.
- S. Huang, K. Yang, S. Qi, and R. Wang, “When Large Language Model Meets Optimization,” *Soft Computing*, 2024 (to appear).
- A. Faiz et al., “LLMCarbon: Modeling the end-to-end Carbon Footprint of Large Language Models,” in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2024.
- D. Lazzaro et al., “Minimizing Energy Consumption of Deep Learning Models by Energy-Aware Training,” *arXiv preprint arXiv:2307.00368*, 2023.
- A. M. Yokoyama, M. Ferro, and B. Schulze, “Multi-objective hyperparameter optimization approach with genetic algorithms toward efficient and environmentally friendly machine learning,” *AI Communications*, vol. 37, no. 6, pp. 839–854, 2023.
- L. Hennig, T. Tornede, and M. Lindauer, “Towards Leveraging AutoML for Sustainable Deep Learning: A Multi-Objective HPO Approach on Deep Shift Neural Networks,” *arXiv preprint arXiv:2404.01965*, 2024.
- S. M. Shariatzadeh, M. Fathy, R. Berangi, and M. Shahverdy, “A Survey on Multi-Objective Neural Architecture Search,” *arXiv preprint arXiv:2307.09099*, 2023.