UNIVERSITY OF COPENHAGEN

EC-NAS: Energy Consumption Aware Tabular Benchmarks for Neural Architecture Search

Bridging Energy Efficiency in NAS Research

Pedram Bakhtiarifard[a], Christian Igel[a] and Raghavendra Selvan[ab]

a. Department of Computer Science, University of Copenhagen
b. Department of Neuroscience, University of Copenhagen

Contact Information:
Universitetsparken 1, 1.1.N121
2100 København

Phone: +45 60219910
Email: pba@di.ku.dk

Keywords: Neural architecture search, sustainable machine learning, multi-objective optimization

ENERGY consumption during deep learning model selection, training, and deployment has surged recently [2, 3]. In response, we introduce EC-NAS [1], a specialized tabular benchmark for Neural Architecture Search (NAS) that foregrounds energy efficiency. Tabular benchmarks, by virtue of pre-computed performance metrics, enable cost-effective NAS evaluations. EC-NAS goes a step further, integrating energy consumption data across diverse architectures with Carbontracker.

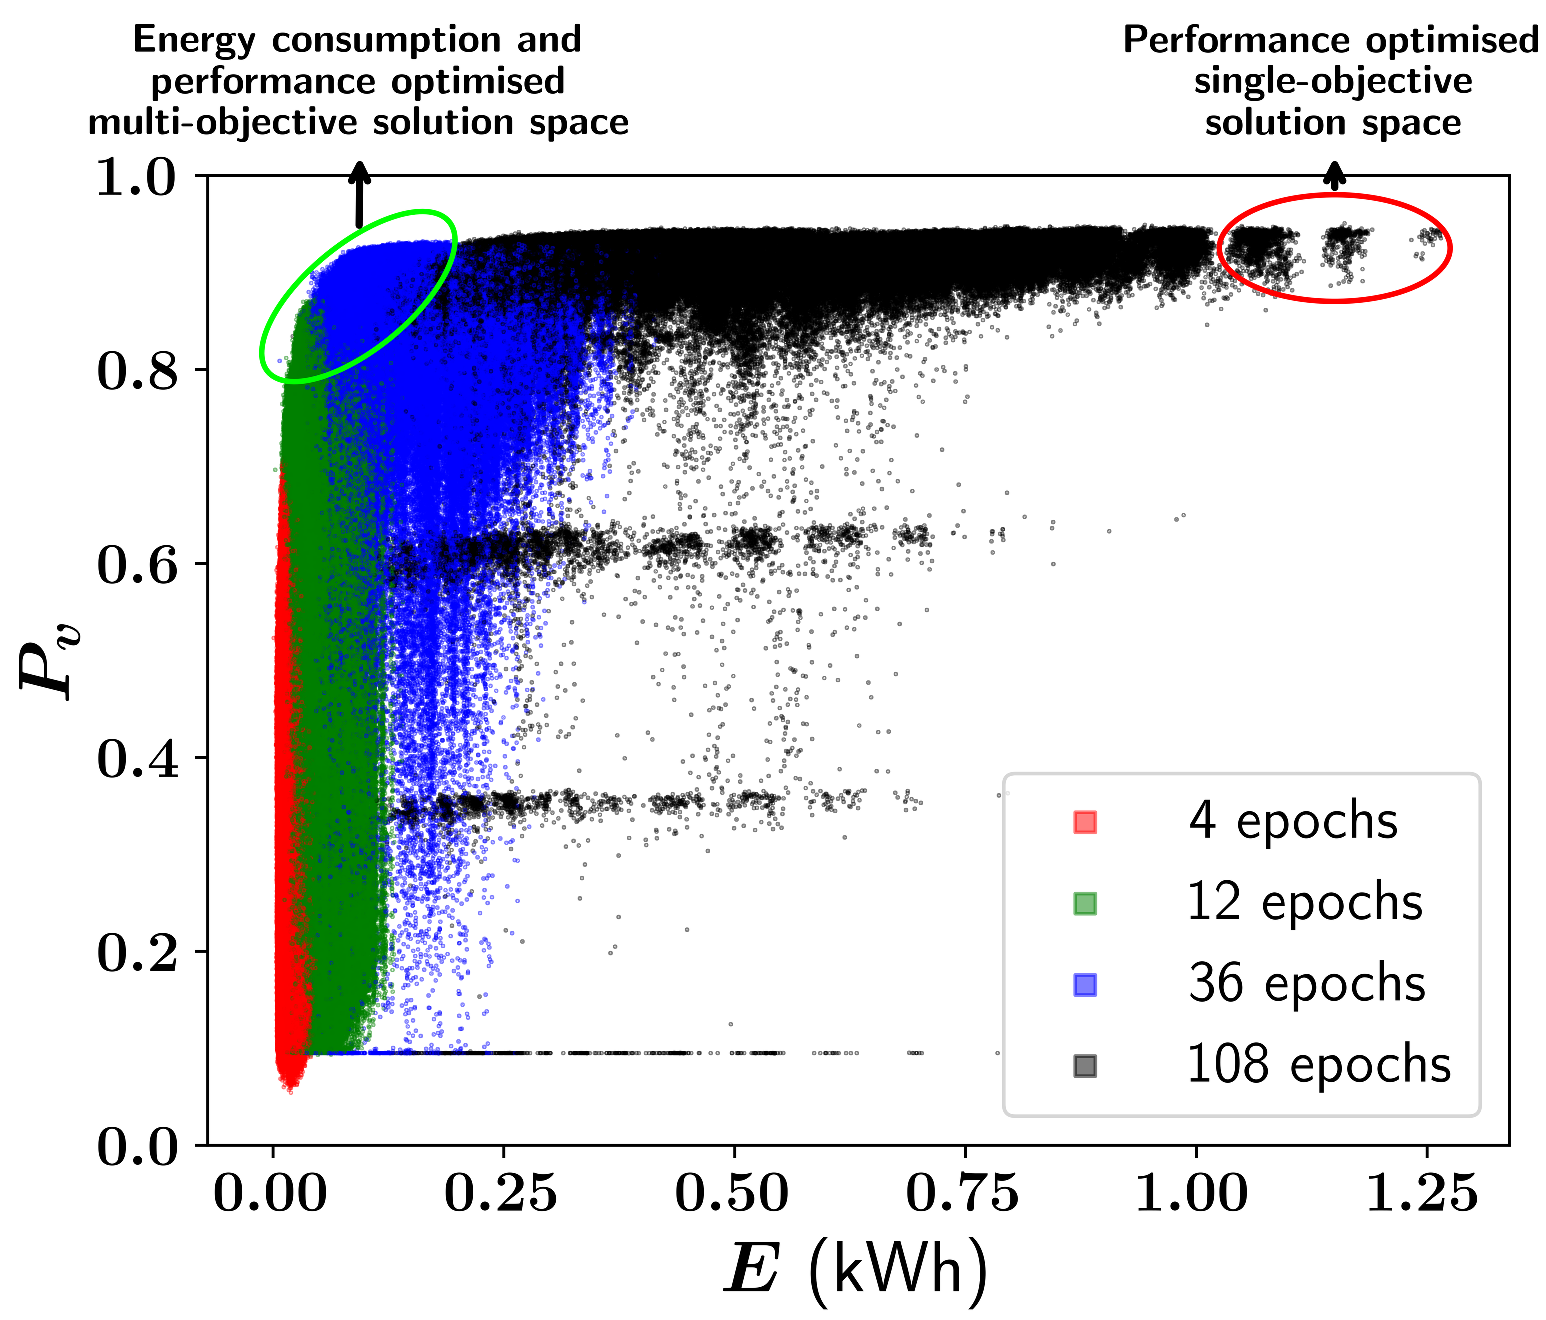

Figure 1: Scatter plot of CNNs: energy (E) vs. performance (P_v). Red ellipse: high performance, high energy. Green ellipse: optimized energy with minor performance drop.

Surrogate Model for Energy Estimation

Our EC-NAS benchmark employs a surrogate model to efficiently predict energy consumption for models within the NAS-Bench-101 dataset [4].

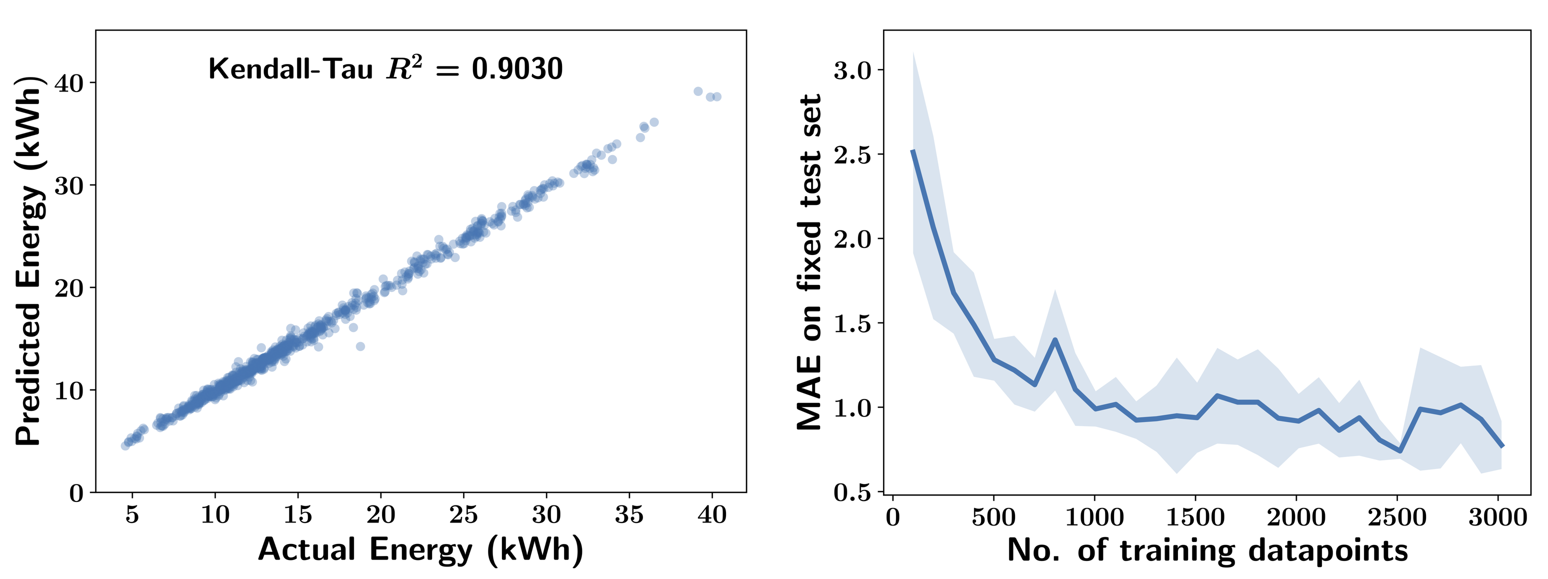

Figure 2: Scatter plot depicting the Kendall-Tau correlation coefficient between predicted and actual energy consumption (left) and the influence of training data size on test accuracy (right). Error bars are based on 10 random initializations.

Dataset Insights

Operation substitutions in deep learning architectures affect performance and energy, showing energy efficiency extends beyond training time.

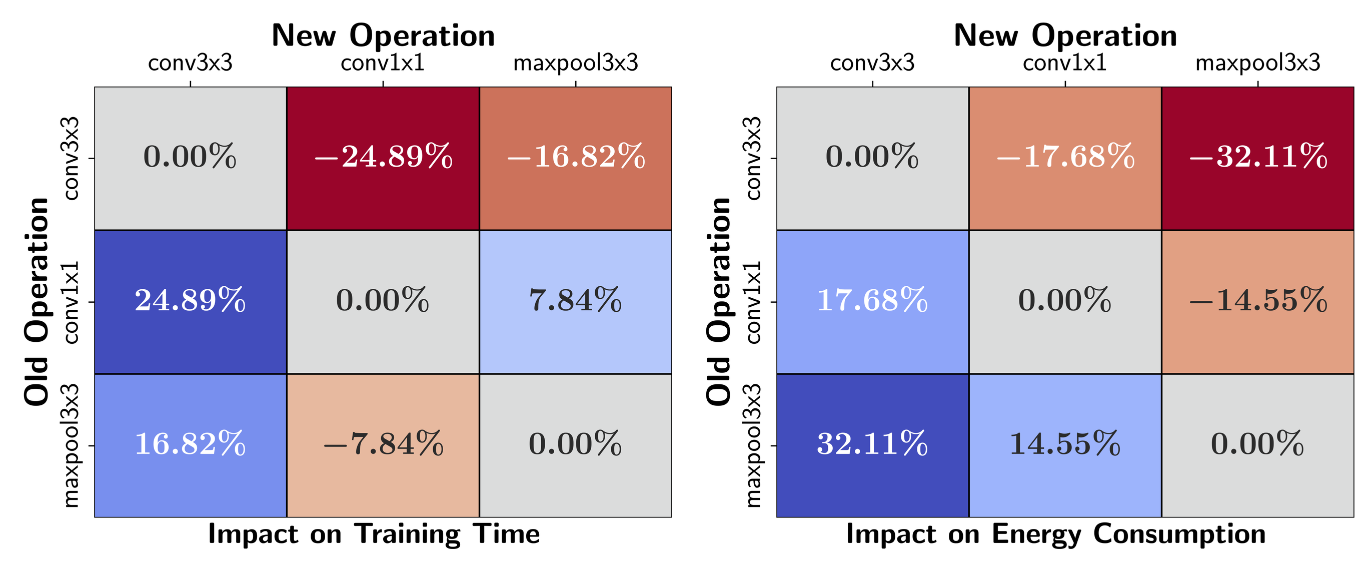

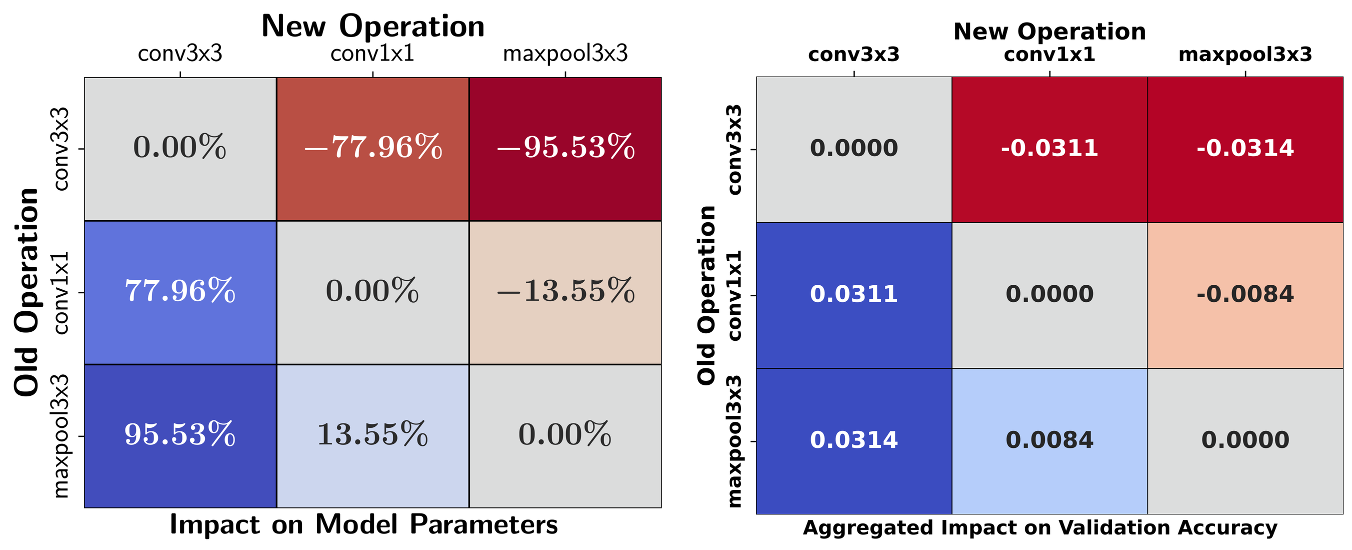

Figure 3: The impact of swapping one operator for another on energy consumption, training time, validation accuracy, and parameter count. The figure illustrates how changing a single operator can affect the different aspects of model performance, emphasizing the importance of selecting the appropriate operators to balance energy efficiency and performance.

Hardware Consistency

Energy trends remain consistent when evaluated across multiple generations of hardware configurations.

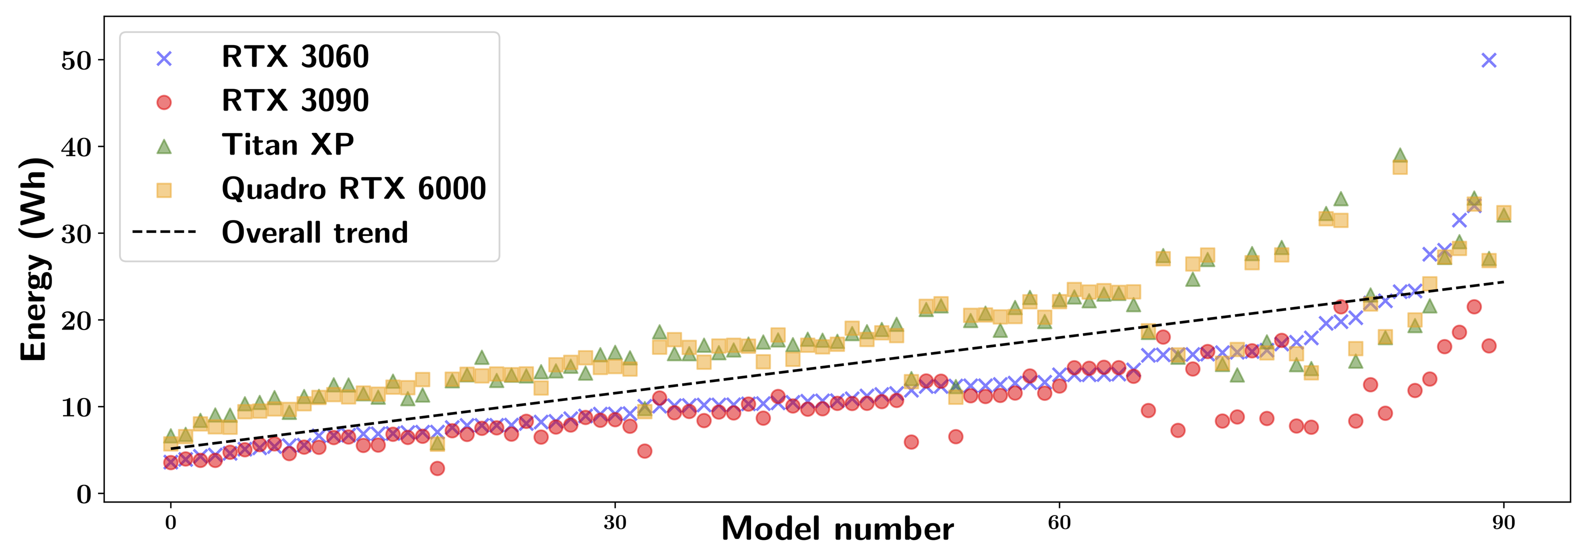

Figure 4: Energy consumption of models with DAGs where $|V| \leq 4$ on different GPUs. Models are organized by their average energy consumption for clarity.

Multi-Objective Optimization in NAS

Through EC-NAS and algorithms like SEMOA, we delineated the Pareto front, capturing the balance between energy consumption and accuracy.

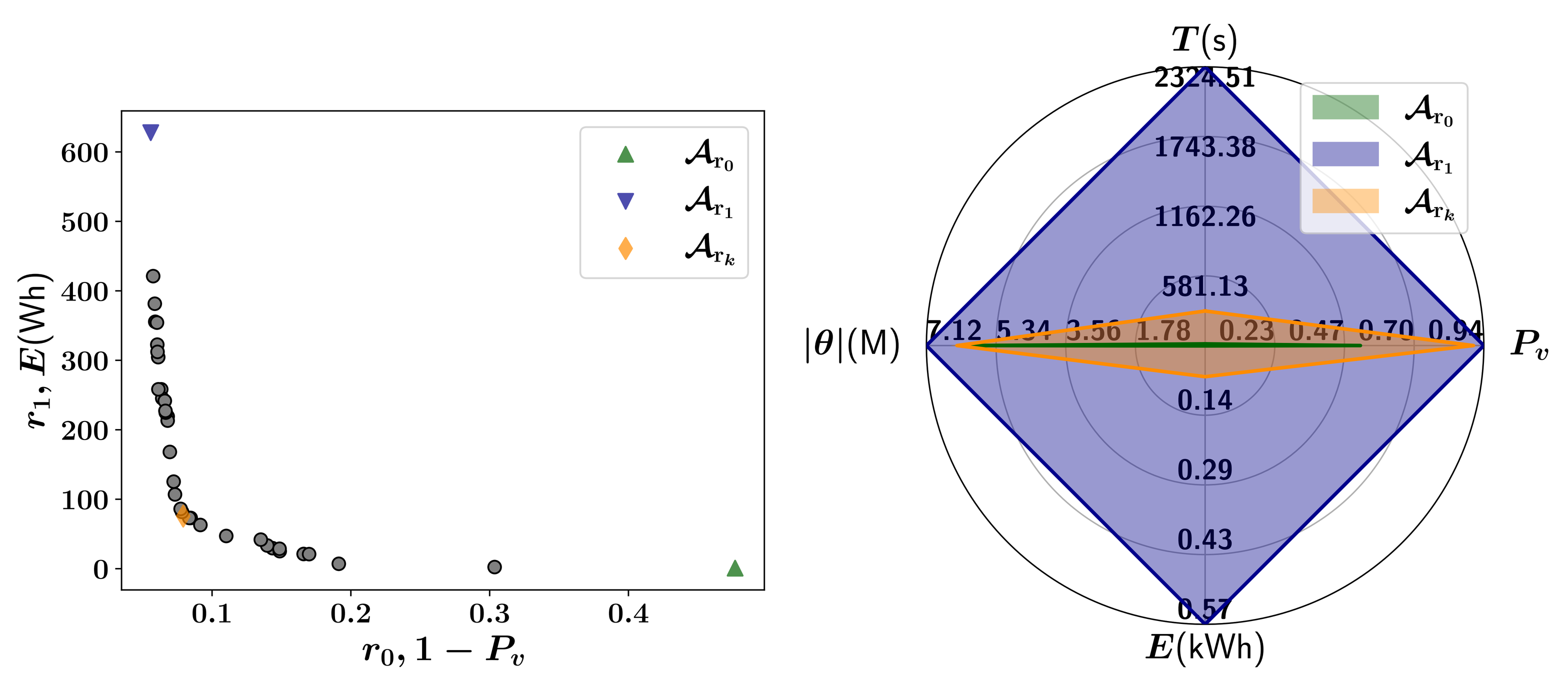

Figure 5: (Left) A representation of the Pareto front for one of SEMOA's runs. (Right) Summary of metrics for the extrema and knee point architectures for SEMOA from one of the runs.

Conclusion

- **Prioritizing Energy**: We stress the imperative of energy efficiency as a comprehensive indicator than training time, especially in the context of environmental sustainability.
- **Benchmark Utility**: Our EC-NAS benchmark offers researchers a valuable dataset, promoting energy-aware decisions in NAS.
- **Balancing Act**: Multi-objective optimization techniques allow for a clearer understanding of the trade-offs between energy consumption and model performance.

References

[1] P. Bakhtiarifard, C. Igel, and R. Selvan. Ec-nas: Energy consumption aware tabular benchmarks for neural architecture search. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.

[2] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen, and X. Wang. A comprehensive survey of neural architecture search: Challenges and solutions. *ACM Computing Surveys*, 54(4):1–34, 2021.

[3] J. Sevilla, L. Heim, A. Ho, T. Besiroglu, M. Hobbhahn, and P. Villalobos. Compute trends across three eras of machine learning. In *International Joint Conference on Neural Networks (IJCNN)*, 2022.

[4] C. Ying, A. Klein, E. Christiansen, E. Real, K. Murphy, and F. Hutter. NAS-Bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning (ICML)*, 2019.