# The Impact of Training Techniques in Small Neural Networks

Ilarya Franco Computer Science

## Background

As machine learning (ML) models become increasingly integrated into scientific research, industry, and everyday technologies, concerns about their energy consumption and environmental impact have grown. It is estimated that by 2030, over 30% of the world's energy consumption will be from artificial intelligence (AI) models (Jones, 2018). Simply training large models, such as ChatGPT's GPT-3 model, required the energy used in around 121 homes in the US in a year and produced around 550 tons of carbon dioxide (Patterson, 2021). As concerns about sustainability increase, researchers have begun exploring "Green AI," which emphasizes creating algorithms "capable of maximizing the energy efficiency and reducing the environmental impact of AI models, while supporting the use of this technology to respond to different environmental challenges" (Lannelongue, 2021).

While much of the existing research focuses on large models or hardware comparisons, further studies have suggested that training techniques, such as batch size, learning rate, and training duration, can significantly influence energy consumption (Geißler, 2024). Poorly chosen hyperparameters can cause identical models with the same accuracy to use five times more energy than optimized configurations (Geißler, 2024). However, little research has been done on how these training techniques affect the energy usage of small neural networks, which are widely used in education and lightweight applications, such as Internet of Things (IoT) devices. This project addresses that gap.

## Research Statement

This project investigates how two common training techniques, batch size and early stopping, affect the energy consumption, training time, and accuracy of small neural networks.

## Project Goals and Objectives

1. Measure the energy consumption of small neural networks under different batch sizes and early stopping configurations.

2. Compare how these two training techniques influence energy usage, accuracy, and training time.
3. Analyze the tradeoffs between energy efficiency and model performance.
4. Produce a dataset, visualizations, and a written research report summarizing the findings.

**Methodology**

This study will use a small convolutional neural network (CNN) trained on a Mixed National Institute of Standards and Technology (MNIST) dataset. MNIST datasets are commonly used to benchmark ML algorithms while being computationally manageable (Kim, 2018). This project will evaluate two training techniques individually: batch size and early stopping.

The computing environment will be installed and run on a personal laptop equipped with an Intel-based Central Processing Unit (CPU) and a standard Graphics Processing Unit (GPU). All experiments will be executed locally on this machine so that the energy-measurement tools can directly monitor hardware usage. The environment will be prepared by installing Python, PyTorch, and any other required libraries in a virtual environment. After the software installation, the system will be configured with CodeCarbon, Green Algorithms, and pyRAPL to measure energy consumption. A short pilot script will be used to verify that each tool records energy usage correctly and consistently, ensuring accurate measurements during the experiment.

The small CNN will be implemented with two to three convolutional layers and a fully connected classifier. The convolutional layers extract meaningful features, such as edges, curves, and stroke shapes, from the MNIST digit images to learn patterns. The fully connected classifier interprets these learned patterns and maps the features to one of the ten digit classes. This architecture will be kept constant throughout the study to ensure that any differences in energy consumption or performance are solely due to the training techniques being tested.

For each condition, the model will be trained five times to reduce “noise” from hardware fluctuations. To evaluate batch sizes, the model will be trained using 16 samples per batch and again using 128 samples per batch. To evaluate early stopping, the model will be trained with early stopping disabled versus with early stopping enabled using a patience value of three epochs. During each training run, the energy usage, training time, and final accuracy on the MNIST test set will be recorded. Energy consumption will be measured using CodeCarbon, Green Algorithms, and pyRAPL. CodeCarbon is a software tool that tracks the carbon emissions generated by a codebase (Bolón-Canedo, 2024). Green Algorithms is another tool that estimates the carbon footprint of a computation, taking into account the hardware configurations without interfering with the code (Lannelongue, 2021). pyRAPL is a toolkit that measures the CPU energy usage of a host machine while executing Python code, which is commonly used for machine learning algorithms (pyRAPL, 2026). These three tools are used to measure both hardware and software energy usage and allow for cross-validation of results to reduce measurement bias.

After completing all runs, the collected data will be aggregated and visualized using tables and graphs. These visualizations will be used to compare how batch size and early stopping influence the energy efficiency, training time, and performance of the model. The results will also be used to identify the tradeoffs when balancing energy and performance.

## Anticipated Outcomes

This project is expected to produce a complete dataset containing the energy consumption, training time, and accuracy results for each experimental condition. Additionally, visualizations that compare how batch size and early stopping influence the efficiency and performance of the small neural network will be produced. By analyzing these results, the study will likely show that larger batch sizes reduce overall energy usage by improving hardware utilization and shortening training time. Meanwhile, early stopping decreases energy consumption by preventing unnecessary computation after the model has converged.

These findings will be used to create a written research paper that explains the methodology, presents the analyzed data, and discusses the tradeoffs between energy efficiency and model accuracy. A poster will also be created for the university showcase that highlights the key results and practical recommendations for energy-conscious training practices.

These outcomes are relevant for undergraduate and graduate students involved in machine learning, researchers working with limited computational resources, and educators seeking to incorporate sustainability into ML coursework. Developers building lightweight models for laptops or Internet of Things devices may also benefit from understanding how training decisions affect energy usage. This project aims to provide insight into practical strategies for reducing computational cost while maintaining model performance.

## Project Significance

The significance of this project is in its contribution to the growing field of sustainable machine learning. As data centers continue to consume increasing amounts of electricity, reducing the energy consumption of computational tasks is an urgent priority. While the environmental impact of large neural networks has been studied, it is important to realize that even smaller models contribute to cumulative carbon emissions, especially when trained repeatedly in educational or research settings.

This project also aligns with the principles of Green AI, which calls for research that prioritizes efficiency and environmental responsibility (Schwartz, 2020). By focusing on training techniques, rather than model architecture, this project can provide evidence-based guidelines for energy-efficient training that has the potential to influence classroom practices, research workflows, and the development of lightweight machine learning applications. Its findings can promote responsible computing practices that reduce the environmental impact of machine learning models.

## Project Timeline

| Dates | Tasks |
|---|---|
| January 1, 2027 - January 14, 2027 | Conduct literature review; finalize training techniques; set up environment |
| January 15, 2027 - January 31, 2027 | Install and test CodeCarbon and pyRAPL; run pilot experiments |
| February 1, 2027 - February 28, 2027 | Run full training experiments for each technique; collect energy data |

| | |
|---|---|
| March 1, 2027 - March 15, 2027 | Analyze data; generate visualizations; compare techniques |
| March 16, 2027 - April 10, 2027 | Draft research paper; create poster for Showcase |
| April 11, 2027 - April 30, 2027 | Revise final report; prepare presentation |

**Budget**

| **Item** | **Units** | **Cost** | **Justification** |
|---|---|---|---|
| Cloud Computing Credits (Google Colab Pro) | 5 x $9.99 (per month) | $49.95 | Required for running repeated training experiments without overloading personal hardware. |
| External SSD (1 TB) | 1 | $60 | For storing datasets, model checkpoints, and energy logs. |
| Poster Printing (SGA Printing Services) | 1 | $20 | For Showcase presentation. |
| Cables/Adapters (USB-C to HDMI, USB-C power cable, spare USB-A cable) | 3 x $20 | $60 | Needed for connecting a laptop to external monitors and measurement hardware during experiments. |
| **Total** | | **$189.95** | |

## References

Bolón-Canedo, Verónica, et al. “A Review of Green Artificial Intelligence: Towards a More Sustainable Future.” *Neurocomputing*, vol. 599, Sept. 2024, p. 128096. *ScienceDirect*, https://doi.org/10.1016/j.neucom.2024.128096.

Geißler, Daniel, et al. “The Power of Training: How Different Neural Network Setups Influence the Energy Demand.” *Architecture of Computing Systems*, edited by Dietmar

Fey et al., Springer Nature Switzerland, 2024, pp. 33–47. *Springer Link*, https://doi.org/10.1007/978-3-031-66146-4_3.

Jones, Nicola. "How to Stop Data Centres from Gobbling up the World's Electricity." *Nature*, vol. 561, no. 7722, Sept. 2018, pp. 163–66. *www.nature.com*, https://doi.org/10.1038/d41586-018-06610-y.

Kim, Bitna, and Young Ho Park. "Beginner's Guide to Neural Networks for the MNIST Dataset Using MATLAB." *Korean Journal of Mathematics*, vol. 26, no. 2, June 2018, pp. 337–48. *kkms.org*, https://doi.org/10.11568/kjm.2018.26.2.337.

Lannelongue, Loïc, et al. "Green Algorithms: Quantifying the Carbon Footprint of Computation." *Advanced Science*, vol. 8, no. 12, June 2021, p. 2100707. *DOI.org (Crossref)*, https://doi.org/10.1002/advs.202100707.

Patterson, David, et al. "Carbon Emissions and Large Neural Network Training." arXiv:2104.10350, arXiv, 23 Apr. 2021. *arXiv.org*, https://doi.org/10.48550/arXiv.2104.10350.

Schwartz, Roy, et al. "Green AI." *Communications of the ACM*, vol. 63, no. 12, Nov. 2020, pp. 54–63. *DOI.org (Crossref)*, https://doi.org/10.1145/3381831.

*Welcome to pyRAPL's Documentation! — pyRAPL 0.2.0 Documentation*. https://pyrapl.readthedocs.io/en/latest/index.html. Accessed 7 Apr. 2026.

***********************************************************************

## The Impact of Training Techniques in Small Neural Networks (OLD)

Ilarya Franco
Computer Science

### Background

As machine learning (ML) models become increasingly integrated into scientific research, industry, and everyday technologies, concerns about their energy consumption and environmental impact have grown. It is estimated that by 2030, over 30% of the world's energy consumption will be from artificial intelligence (AI) models (Jones, 2018). Simply training large models, such as ChatGPT's GPT-3 model, required the energy used in around 121 homes in the US in a year and produced around 550 tons of carbon dioxide (Patterson, 2021). As concerns about sustainability increase, researchers have begun exploring "Green AI," which emphasizes creating algorithms "capable of maximizing the energy efficiency and reducing the environmental impact of AI models, while supporting the use of this technology to respond to different environmental challenges" (Lannelongue, 2021).

While much of the existing research focuses on large models or hardware comparisons, further further studies have suggested that training techniques, such as batch size, learning rate, and training duration, can significantly influence energy consumption (Geißler, 2024). Poorly chosen hyperparameters can cause identical models with the same accuracy to use five times more energy than optimized configurations (Geißler, 2024). However, little research has been done on how these training techniques affect the energy usage of small neural networks, which are widely used in education and lightweight applications, such as Internet of Things devices. This project addresses that gap.

### Research Statement

This project investigates how common training techniques, such as batch size, learning rate, early stopping, and number of epochs, affect the energy consumption, training time, and accuracy of small neural networks.

### Project Goals and Objectives

1. Measure the energy consumption of small neural networks under different training techniques.
2. Compare how batch size, learning rate, early stopping, and number of epochs influence energy usage, accuracy, and training time.
3. Analyze the tradeoffs between energy efficiency and model performance.

**Methodology**

This study will use a small convolutional neural network trained on an MNIST dataset, which is used to “benchmark the performance of machine learning algorithms” while being computationally manageable (Kim, 2018). Four commonly-used training techniques will be evaluated individually: batch size, learning rate, early stopping, and number of epochs.

Hyperparameters such as batch size (16 vs 128 samples) and learning rate (constant, step decay, cosine decay) directly influence the number of optimization steps required, which affects total energy consumption (Ariyanti, 2024). Early stopping has been identified as a key strategy for preventing unnecessary computation, as many machine learning experiments waste energy by continuing training after convergence (Henderson, 2022). The number of epochs (short vs extended training) is examined because training duration determines how long computation resources are used, and longer schedules lead ot an increase in energy consumption (Ariyanti, 2024). The model will be trained several times on each variable individually to isolate its effect. The energy consumption, training time, and accuracy will be recorded.

Energy consumption will be measured using CodeCarbon, Green Algorithms, and pyRAPL. CodeCarbon is a software tool that tracks the carbon emissions generated by a codebase (Bolón-Canedo, 2024). Green Algorithms is another tool that estimates the carbon footprint of a computation, taking into account the hardware configurations without interfering with the code (Lannelongue, 2021). pyRAPL is a toolkit that measures the CPU energy usage of a host machine while executing Python code, which is commonly used for machine learning algorithms (pyRAPL, 2026). These three tools are used to measure both hardware and software energy usage and allow for cross-validation of results to reduce measurement bias.

## Anticipated Outcomes

Based on prior research into sustainable AI, larger batch sizes are expected to reduce energy consumption by increasing the GPU utilization, which reduces overall training time (Geißler, 2024). Early stopping is also anticipated to substantially lower energy consumption without reducing accuracy, as most models continue training long after performance has plateaued, which results in unnecessary energy usage (Ariyanti, 2024). Learning rates such as step decay or cosine decay are expected to improve accuracy while having reasonable energy costs as they increase the speed to convergence (Atiénzar, 2026). The greater number of epochs, the greater the energy consumption is expected to be as training duration directly relates to an increase in energy usage (Qiu, 2021). This project is expected to reveal the tradeoffs between accuracy and energy efficiency through various training techniques.

These outcomes are relevant for undergraduate and graduate students involved in machine learning, researchers working with limited computational resources, and educators seeking to incorporate sustainability into ML coursework. Developers building lightweight models for laptops or Internet of Things devices may also benefit from understanding how training decisions affect energy usage. This project aims to provide insight into practical strategies for reducing computational cost while maintaining model performance.

## Project Significance

The significance of this project is in its contribution to the growing field of sustainable machine learning. As data centers continue to consume increasing amounts of electricity, reducing the energy consumption of computational tasks is an urgent priority. While the environmetal impact of large neural networks has been studied, it is important to realize that even smaller models contribute to cumulative carbon emissions, especially when trained repeatedly in educational or research settings.

This project also aligns with the principles of Green AI, which calls for research that prioritizes efficiency and environmental responsibility (Schwartz, 2020). By focusing on training techniques, rather than model architecture, this project can provide evidence-based guidelines for energy-efficient training that has the potential to influence classroom practices, research workflows, and the development of lightweight machine learning applications. Its findings can promote responsible computing practices that reduce the environmental impact of machine learning models.

## Project Timeline

| Dates | Tasks |
|---|---|
| January 1, 2027 - January 14, 2027 | Conduct literature review; finalize training techniques; set up environment |
| January 15, 2027 - January 31, 2027 | Install and test CodeCarbon and pyRAPL; run pilot experiments |
| February 1, 2027 - February 28, 2027 | Run full training experiments for each technique; collect energy data |
| March 1, 2027 - March 15, 2027 | Analyze data; generate visualizations; compare techniques |
| March 16, 2027 - April 10, 2027 | Draft research paper; create poster for Showcase |
| April 11, 2027 - April 30, 2027 | Revise final report; prepare presentation |

## Budget

| **Item** | **Units** | **Cost** |
|---|---|---|
| Cloud Computing Credits | 150 | $150 |
| External SSD | 1 | $60 |
| Poster Printing | 1 | $40 |
| Cables/Adapters | 3 x $20 | $60 |
| **Total** | | **$310** |

## References

Ariyanti, Sri, et al. “Trade-Off Between Energy Consumption and Three Configuration Parameters in Artificial Intelligence (AI) Training: Lessons for Environmental Policy.”

*Sustainability*, vol. 17, no. 12, June 2025, p. 5359. *DOI.org (Crossref)*, https://doi.org/10.3390/su17125359.

Atiénzar, B, et al. “Energy Scalability in the Training of AI Models for Image Processing: The Role of Hyperparameters.” *Journal of Systems Scalability*, vol. 1, no. 1, March 2026, p. 1-5. *icck.org*, https://doi.org/10.62762/JSS.2025.960646.

Bolón-Canedo, Verónica, et al. “A Review of Green Artificial Intelligence: Towards a More Sustainable Future.” *Neurocomputing*, vol. 599, Sept. 2024, p. 128096. *ScienceDirect*, https://doi.org/10.1016/j.neucom.2024.128096.

Geißler, Daniel, et al. “The Power of Training: How Different Neural Network Setups Influence the Energy Demand.” *Architecture of Computing Systems*, edited by Dietmar Fey et al., Springer Nature Switzerland, 2024, pp. 33–47. *Springer Link*, https://doi.org/10.1007/978-3-031-66146-4_3.

Henderson, Peter, et al. “Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning.” arXiv:2002.05651, arXiv, 29 Nov. 2022. *arXiv.org*, https://doi.org/10.48550/arXiv.2002.05651.

Jones, Nicola. “How to Stop Data Centres from Gobbling up the World’s Electricity.” *Nature*, vol. 561, no. 7722, Sept. 2018, pp. 163–66. *www.nature.com*, https://doi.org/10.1038/d41586-018-06610-y.

Kim, Bitna, and Young Ho Park. “Beginner’s Guide to Neural Networks for the MNIST Dataset Using MATLAB.” *Korean Journal of Mathematics*, vol. 26, no. 2, June 2018, pp. 337–48. *kkms.org*, https://doi.org/10.11568/kjm.2018.26.2.337.

Lannelongue, Loïc, et al. “Green Algorithms: Quantifying the Carbon Footprint of Computation.” *Advanced Science*, vol. 8, no. 12, June 2021, p. 2100707. *DOI.org (Crossref)*, https://doi.org/10.1002/advs.202100707.

Patterson, David, et al. “Carbon Emissions and Large Neural Network Training.” arXiv:2104.10350, arXiv, 23 Apr. 2021. *arXiv.org*, https://doi.org/10.48550/arXiv.2104.10350.

Qiu, Xinchi, et al. “A First Look into the Carbon Footprint of Federated Learning.” *arXiv.Org*, 15 Feb. 2021, https://arxiv.org/abs/2102.07627v6.

Schwartz, Roy, et al. “Green AI.” *Communications of the ACM*, vol. 63, no. 12, Nov. 2020, pp. 54–63. *DOI.org (Crossref)*, https://doi.org/10.1145/3381831.

*Welcome to pyRAPL’s Documentation! — pyRAPL 0.2.0 Documentation*. https://pyrapl.readthedocs.io/en/latest/index.html. Accessed 7 Apr. 2026.