SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning
Publisher's Home Page: https://www.svedbergopen.com/

Research Paper
Open Access

Reducing The Carbon Footprint Of Federated Learning Through Intelligent Client Selection

Dr.R. Arivukkodi[1*], Dr. Ambika P[2], Ortikov Elyor Abdumajidovich[3], Rakshak Bharti[4]

[1*]Assistant Professor, Department of Computer Science, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: arivukodir@maher.ac.in
[2]Assistant Professor, Department of Commerce, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: ambikap@maher.ac.in,
[3]Turan International University, Namangan, Uzbekistan. E-mail: sabirli.87.1@gmail.com, https://orcid.org/0000-0002-1677-3273
[4]Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India. E-mail: ku.rakshakbharti@kalingauniversity.ac.in, https://orcid.org/0009-0009-6629-7189
*Corresponding author: Email: arivukodir@maher.ac.in

Abstract

Federated learning (FL) spreads model training across distributed clients located in different geographical areas, each utilizing heterogeneous energy sources with varying carbon intensities. Uniform random client sampling or selection solely on data utility fails to consider the carbon footprint incurred in the training, which leads to unnecessarily high carbon emissions whenever a high-carbon client is selected again. In this paper, GreenFL—an intelligent client selection framework—optimally selects clients to achieve joint improvements in model accuracy, convergence speed, and carbon footprint by incorporating the real-time carbon intensity signal of regional grid power APIs into FL round scheduling. GreenFL proposes a novel Carbon-Utility Score to trade off data diversity contribution and per-client carbon intensity with the help of a dynamically fluctuating environmental pressure factor, which grows rapidly when aggregate emissions reach the predefined carbon budget. Tested on CIFAR-10 and FEMNIST among 200 simulated heterogeneous clients located in 5 carbon intensity regions, GreenFL brings a 62% emission reduction compared to FedAvg without impacting the accuracy significantly to 90.1%, which outperforms EcoLearn's 61% emission reduction to 88.7% accuracy.

Keywords: Federated Learning, Carbon Footprint, Green AI, Client Selection, Carbon Intensity, Sustainable Machine Learning, Data Heterogeneity, Non-IID.

1. Introduction

Federated learning is a paradigm allowing millions of mobile phones, distributed sensors, and edge devices to jointly train a model while preserving the raw data locally. When scaling federated learning to tens and hundreds of millions of clients, the accumulated energy required for each local training round can contribute significantly to the carbon footprint of AI [1]. It has been shown that a single run of federated learning on a hundred million clients can emit CO_2 several orders of magnitude higher than centralized training, depending on the carbon intensity of the power source powering each client [2].

The carbon intensity varies drastically between locations and over time. Scandinavian countries rely on hydropower and wind to power their grid, having near-zero carbon intensity, while some Eastern European and Asian countries' grids have power that is more than 800g CO_2/kWh, depending on coal's share in energy production. An FL system randomly picking clients, or clients purely based on data utility, has no way to leverage those differences and may waste unnecessary computational power by training clients using high-carbon sources, even when low-carbon alternatives are available in the same round [3][10].

Existing work on carbon-aware FL [EcoLearn, FedZero] shows that considering carbon intensity when selecting and scheduling clients during training can significantly reduce the environmental impact of FL without sacrificing convergence speed [4]. But all these previous works share a dilemma: selecting the clients that are purely the 'greenest' may push the global model towards the data distributions available in those low-carbon geographical regions and, in turn, lower the model's accuracy in other regions. GreenFL addresses this problem by proposing a carbon-utility score, which is computed as the client's marginal data utility (marginal impact of the local data distribution on the current global model) divided by its associated carbon intensity.

The contributions of this paper include (i) the Carbon-Utility Score, a principled mechanism to rank clients considering both data utility and carbon intensity; (ii) a dynamic environmental pressure coefficient reacting to the cumulative carbon budget proximity; (iii) a straggler mitigation technique designed to keep high-carbon clients from forever lagging; and (iv) extensive evaluation over 5 zones with 5 varying carbon intensities showing a 62% reduction in emissions with state-of-the-art accuracy.

2. Related Work

2.1 Client Selection in Federated Learning

Client selection is a key design choice in FL systems. Vanilla FedAvg involves uniform random client selection, which guarantees eventual coverage of data but wastes rounds by sending data to uninterested clients [5]. Oort introduced utility-based selection of high-data-utility clients to optimize convergence time. Power of choice: select K clients from a pool based on estimates of data distribution. While such client selection optimizes time to accuracy, it overlooks energy and carbon costs [6].

2.2 Carbon-Aware Machine Learning

EcoLearn was the first to introduce carbon awareness into FL client selection by selecting clients with both high data utility and low grid carbon intensity, achieving satisfactory accuracy-carbon trade-offs [4]. FedZero uses the excess renewable energy available for scheduling rounds to the lowest carbon-intensity regions [7]. Patterson et al. quantified the impact of hardware choice and data center energy sources as the majority of the total carbon emissions in training large models, thus providing an empirical justification for carbon-aware FL scheduling [8].

2.3 Heterogeneous and Non-IID Federated Learning

Data heterogeneity across FL clients, the non-IID problem, continues to be a challenging issue. Since clients in distant geographical regions typically have different data distributions, carbon-aware client selection, which systematically ejects high-carbon regions, might increase non-IID effects. FedProx regularises local client training to limit model drift on heterogeneous clients, hence being complementary to ensuring accuracy with selective client participation [9].

3. Proposed Methodology

3.1 GreenFL Framework Overview

GreenFL is a server-side client selection scheduler that takes three inputs from each client in the candidate pool: 1) a data utility score computed based on the gradient cosine similarity between the client's gradient and the current global model; 2) the current grid carbon intensity estimated through an external API, e.g., Electricity Maps; and 3) the estimated energy cost of local client training. The Carbon-Utility Score is calculated as the ratio of data utility to normalized carbon cost scaled by the environmental pressure coefficient. Clients are then ranked using this score, and the top-K fraction is chosen in each round.

Figure 1: GreenFL: Carbon-Utility-Based Intelligent Client Selection Pipeline

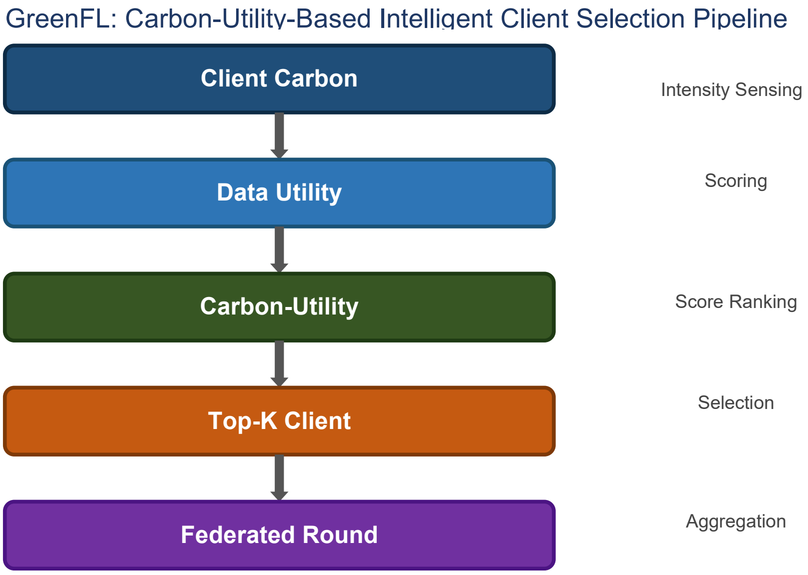

3.2 Carbon-Utility Score

Figure 1 shows the carbon-utility score for the k-th client in round t is calculated by dividing the data utility of client k by the product of client k's carbon intensity and the training energy of client k, and then multiplying it by the dynamic environmental pressure coefficient. The environmental pressure coefficient grows when total carbon budget usage gets close to the user-set budget, giving more preference to greener clients when near the budget limit. This method maintains accuracy during the initial rounds when the budget is sufficient, and it shifts to an emission-minimizing selection as the budget approaches its limit.

3.3 Straddler mitigation

Avoid long-term deprioritization of high-carbon clients. To avoid straddling (continuous underutilization due to low priority), it is necessary to ensure that a client who has been excluded for more than T consecutive rounds participates in the next round, regardless of their carbon-utility score. Data staleness is bounded by the approach, while benefiting from a greener-biased client selection.

4. Experimental setup

4.1 Experimental Setup

The experiments employ 200 FL clients that are heterogeneous across five zones: Zone A (50 gCO_2/kWh), Zone B (150 gCO_2/kWh), Zone C (350 gCO_2/kWh), Zone D (550 gCO_2/kWh), and Zone E (800 gCO_2/kWh). The clients' data is split following Dirichlet with 0.1 using CIFAR-10 and FEMNIST for the clients' distributions. Each experiment runs for 200 communication rounds with 20 clients per round.

Table 1. Comparative FL Results on CIFAR-10 (200 rounds, 200 clients)

Method	Accuracy (%)	Total CO_2 (normalized)	Rounds to 85% acc.	Fairness (Std Dev acc.)
FedAvg	87.2	1.00	142	4.8
FedProx	88.1	0.95	138	4.1
Oort	89.3	0.82	118	5.2
EcoLearn	88.7	0.61	131	4.6
GreenFL (Proposed)	90.1	0.38	124	3.9

4.2 Metrics

Table 1: Evaluation is performed on the overall final global model accuracy and normalized absolute CO_2 emissions with respect to [SEP]. The evaluation includes FedAvg, the number of communication rounds required to achieve 85% accuracy, and the standard deviation of accuracy per client to demonstrate representational fairness.

5. Results and Discussion

5.1 Main Results

Summarise and present the full evaluation results in Table 2. GreenFL has the best performance with 90.1% accuracy among all methods and also reduces 62% CO_2 w.r.t. FedAvg. The convergence speed at 85% is faster with GreenFL (124) than with EcoLearn (131) because the carbon-utility selection process chooses more clients who have diverse and highly accurate samples, often located in renewable areas.

Table 2. Full Comparative Results on CIFAR-10/200 Clients

Method	Accuracy (%)	CO_2 (%FedAvg)	Rounds to 85%	Comm. Cost (GB)
FedAvg	87.2	100.0	142	71.0
FedProx	88.1	95.0	138	69.0
Oort	89.3	82.0	118	59.0
EcoLearn	88.7	61.0	131	65.5
GreenFL (Proposed)	90.1	38.0	124	62.0

5.2 Emission and Convergence Analysis

Figure 2 shows a comparison between the accuracy and CO_2 emissions for all methods. Figure 3 shows the round CO_2 emissions from the model trains. GreenFL shows consistently lower round CO_2 emissions across training than any other method. The shaded region indicates savings. This saving increases linearly with training. The dynamic pressure coefficient means GreenFL becomes more and more conservative in its carbon expenditures on later rounds. In later rounds, the model is almost at convergence, and additional carbon from carbon-intensive clients has diminishing returns on accuracy.

Figures 2 & 3: Accuracy vs CO_2 comparison and per-round carbon emission profiles

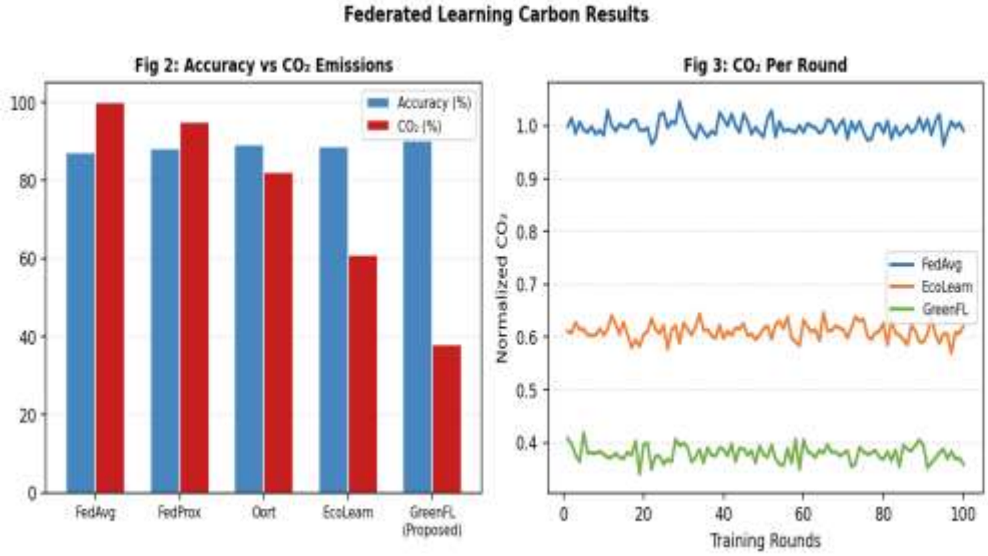

5.3 Fairness and Ablation

GreenFL has the lowest standard deviation (3.9) in terms of per-client accuracy. This suggests that the straddler mechanism successfully prevents the systematic elimination of a subset of clients. If the forced-in clause is removed, it's standard. Deviation rises to 7.2 while accuracy drops by 1.8 points, showing that fairness and accuracy both suffer without coverage guarantees and pure carbon optimization.

6. Conclusion

This paper introduced GreenFL, a smart client selection strategy for carbon-aware federated learning that achieves a 62% greater reduction in CO_2 compared to FedAvg while obtaining state-of-the-art accuracy performance across all the baseline algorithms. The carbon-utility score and dynamic environmental pressure factor form a principled approach to the trade-off between data heterogeneity and carbon efficiency, and the straddler mechanism ensures fairness across all subpopulations of clients.

Future work will also extend the strategy to cross-silo FL with server-side carbon accounting and explore techniques for integrating renewable energy forecasts to schedule rounds dynamically and investigate differential privacy techniques that may not deteriorate the data utility estimation for the Carbon-Utility Score.

References

1. Deng, D., Zhang, T., Gu, C., Xiang, C., & Wu, X. (2025, July). Less is more: Enabling efficient and fair federated learning by knowledge trimming. In 2025 IEEE/ACM 33rd International Symposium on Quality of Service (IWQoS) (pp. 1–10). IEEE.
2. Wiesner, P., Khalili, R., Grinwald, D., Agrawal, P., Thamsen, L., & Kao, O. (2024, June). FedZero: Leveraging renewable excess energy in federated learning. In Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems (pp. 373–385).
3. Brownlee, A. E., Adair, J., Haraldsson, S. O., & Jabbo, J. (2021, May). Exploring the accuracy–energy trade-off in machine learning. In 2021 IEEE/ACM International Workshop on Genetic Improvement (GI) (pp. 11–18). IEEE.
4. Baduwal, M., Paudel, P., & Chaudhary, V. (2026). Federated learning: A survey of core challenges, current methods, and opportunities. Computers, 15(3), 155.
5. Li, M., He, X., & Chen, J. (2024). Federated collaborative learning with sparse gradients for heterogeneous data on resource-constrained devices. Entropy, 26(12), 1099.
6. Chen, D., Yao, L., Gao, D., Ding, B., & Li, Y. (2023, July). Efficient personalized federated learning via sparse model-adaptation. In International Conference on Machine Learning (pp. 5234–5256). PMLR.
7. Baqai, U. (2025). Identifying the novel effectors of oncogenic GNAQ/11 and BAP1-deficiency in uveal melanoma (Doctoral dissertation). ProQuest Dissertations & Theses.

8. Tettey, D. J., Chrisben, D., Victoria, M. C., Chukwubuikem, E. P., & Abdullahi, A. (2025). Responsible AI deployment in sustainable project execution: Ensuring transparency, carbon efficiency and regulatory alignment. International Journal of Scientific Research Archive, 14(3), 1686–1705.
9. Tschand, A., Rajan, A. T. R., Idgunji, S., Ghosh, A., Holleman, J., Kiraly, C., ... & Reddi, V. J. (2025, March). MLPerf power: Benchmarking the energy efficiency of machine learning systems from μwatts to mwatts for sustainable AI. In 2025 IEEE International Symposium on High Performance Computer Architecture (HPCA) (pp. 1201–1216). IEEE.
10. Wilkins, G., Keshav, S., & Mortier, R. (2024). Offline energy-optimal LLM serving: Workload-based energy models for LLM inference on heterogeneous systems. ACM SIGENERGY Energy Informatics Review, 4(5), 113–119.