Eco-Friendly AI: a framework for Data Centric Green Federated Learning

Mattia Sabella[1,*,†], Monica Vitali[1,†]

[1]DEIB, Politecnico di Milano, Milan, Italy

Abstract

The environmental cost of deep learning is increasingly significant, prompting a shift from performance-driven "Red AI" to sustainability-focused "Green AI". In this paper, we propose a data-centric framework for environmentally sustainable Federated Learning (FL), focused on optimising data quality and node selection to minimise carbon emissions without compromising model performance. Central to our approach is an interactive FL Configuration Selection System, which, given dataset and infrastructure characteristics, assists researchers in configuring greener FL training workflows. Our system integrates data quality metrics and carbon footprint estimates to select environmentally optimal nodes and applies intelligent data reduction through three strategies: Node Selection, Minimal Smart Reduction, and Smart Reduction. We demonstrate the effectiveness of our tool in the context of time series classification, offering a practical solution for sustainable FL research.

Keywords

Green Federated Learning, Data Centric AI, Sustainable Computing

1. Introduction

The growing computational demands of Deep Learning (DL) present serious challenges related to energy consumption, bringing the environmental footprint of Artificial Intelligence (AI) to the forefront. As AI-powered technologies become increasingly widespread, there is a pressing need to move from a performance-centric mindset, often referred to as Red AI, toward a more environmentally sustainable approach, known as Green AI. This shift demands new strategies for optimising AI workflows, particularly in data management, to reduce energy consumption while maintaining model performance.

Machine Learning (ML) efficiency is strictly correlated to data quality. In large-scale, distributed environments, traditional data pre-processing methods must be re-evaluated to account for heterogeneous data sources, privacy constraints, and resource limitations. Federated Learning (FL) has emerged as a promising solution, enabling decentralised model training without requiring raw data to be transferred to a central server. However, FL also introduces new complexities due to variations in data quality, volume, and computational capabilities across participating nodes.

Our research is based on studying the role of data quality measures on FL processes, aiming to reduce the energy consumption and carbon emissions of AI training in federated environments. This paper focuses on presenting an interactive configuration selection system for FL training tasks. It is designed to provide recommendations and predictions on data and node selection to minimise the environmental impact of FL training, based on input related to dataset characteristics and FL architecture.

The paper is organised as follows. After introducing the State of the Art in Section 2, we provide an overview of the methodology applied in our research in Section 3, followed by a detailed explanation of the Eco Federated Learning Framework's capabilities in Section 4. Further clarifications regarding the effectiveness of the proposed solution are presented in Section 5. Section 6 concludes the paper by summarizing the objectives and outlining potential directions for future research.

2nd Workshop on Green-Aware Artificial Intelligence, 28th European Conference on Artificial Intelligence (ECAI 2025), October 25–30, 2025, Bologna, Italy

* Corresponding author.

† These authors contributed equally.

mattia.sabella@polimi.it (M. Sabella); monica.vitali@polimi.it (M. Vitali)

2. State of the Art

AI significantly contributes to IT systems' environmental impact due to rapidly increasing computational demands, mainly driven by Deep Learning (DL) and Deep Neural Networks (DNNs)[1, 2, 3, 4, 5]. Schwartz et al.[6] distinguish between *Red AI* (performance-focused) and *Green AI* (sustainability-focused). Studies highlight AI's carbon footprint concerning infrastructure [7], model tuning [8], and data preparation [9].

Data-Centric AI prioritises enhancing training datasets over model optimisation, significantly affecting model performance and dataset balance [10, 11, 12, 13]. Large datasets often lead to excessive energy use with diminishing returns [14, 15, 16], but reducing dataset size effectively lowers energy consumption without accuracy loss [17]. **Data Quality** (DQ) metrics critically influence data selection, as poor-quality data can bias models [18, 19, 20, 21]. Anselmo et al. [22] proposed optimising data volume and quality to mitigate DL's environmental impact.

Federated Learning, a distributed approach to model training that preserves privacy and reduces data transmission, faces challenges from data heterogeneity affecting model accuracy [23, 24, 25, 26]. Although decentralised FL leverages energy-efficient edge devices, its communication overhead can significantly increase emissions [27]. Carbon-aware methods, like FedZero and other adaptive techniques, effectively lower emissions through strategic location selection and adaptive model sizing [28, 29, 30].

Current research emphasises system-level enhancements or centralised data-centric approaches, leaving the relationship between dataset characteristics and FL's environmental impact understudied. Our research addresses this gap by introducing a federated, data-centric framework that:

- **Analyses the impact of dataset volume and quality on FL emissions**, optimising training configurations.
- **Selects data subsets and participant nodes** based on environmental and computational efficiency.
- **Optimises FL training through an interactive configuration selection system**.

This work offers a scalable solution for sustainable FL deployment, covering the intersection of *energy-efficient data systems*, *federated data management*, and *cloud-based AI sustainability.*

3. Green FL Methodology

Our research explores FL within a fog computing environment, where heterogeneous nodes—differing in hardware capabilities, energy efficiency, and carbon footprints—are distributed across the cloud continuum. This setting introduces unique challenges for FL, including increased energy consumption resulting from both system and statistical heterogeneity, as well as carbon variability due to the geographical dispersion of nodes powered by diverse energy sources. As discussed in Section 2, low-quality data and excessive data volumes can negatively impact both performance and sustainability, leading to inefficient resource utilisation. Prior studies [22][21] have shown that reducing training set size can decrease energy consumption in centralised settings; however, the inherent non-IID nature of data and system heterogeneity in FL complicate such reductions. In FL, data volume and quality are intricately linked to node selection strategies.

To address these challenges, we propose a novel approach for sustainable FL training: an interactive configuration selection system designed for researchers and practitioners. This system suggests efficient node and data selection strategies aimed at minimising carbon emissions while preserving a predefined level of model performance.

To achieve the predefined goal, a **FL Configuration Selection System** has been designed. This system focuses on two key actions:

1. **Optimising data volume**: Recommending appropriate reductions in training data volume to lower the environmental cost of the training process while ensuring the resulting model meets predefined accuracy constraints;

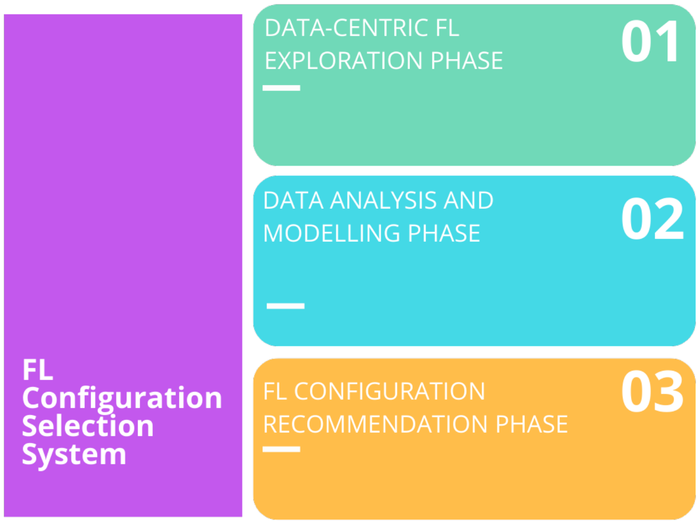

Figure 1: FL Configuration Selection System Architecture.

2. **Selecting efficient client nodes**: Identifying optimal client nodes based on their energy efficiency, environmental impact, and data quality characteristics.

The methodology (Fig. 1) proposes three sequential phases to reach our goal: (i) **Data-Centric FL Exploration**; (ii) **Data Analysis and Modelling**; and (iii) **FL Configuration Selection**.

The **Data-Centric FL Exploration** phase is an investigative stage designed to assess system training across different datasets and data-focused FL configurations. In this phase, an FL simulation environment automates the simulation process to gather experimental data. The main aim is to explore the effects of data volume and quality on the performance and energy efficiency of FL systems during training.

In the **Data Analysis and Modelling** phase, the obtained results are analysed in terms of performance and energy impact. This analysis leads to the development of a predictor capable of recommending optimal dataset size reductions for unseen datasets provided by researchers within an FL setting. Specifically, a machine learning regressor is built to predict the required data volume reduction for an unseen FL configuration while maintaining a predefined accuracy level.

The **FL Configuration Recommendation** phase introduces the mentioned framework that assists researchers by adapting FL setups to their specific contexts. At its core is the FL Configuration Recommender, which selects optimal data subsets and participant nodes based on infrastructure characteristics and task requirements. Using input from the researcher—including dataset properties (volume, accuracy, consistency, completeness), node specifications (e.g., hardware, energy source, location), initial performance estimates, and accuracy target—the system applies volume reductions using the regressor and generates an optimised configuration. It prioritises nodes with low carbon impact and high data quality, applying reduction strategies to improve efficiency and sustainability for the training process.

The proposed approach applies to various types of machine learning tasks. However, for evaluation purposes, we selected a specific FL application: *Time Series Classification.* This task was chosen because it aligns well with FL scenarios, where large volumes of data are generated by multiple devices across different locations.

For local training, we employ a Deep Learning model based on *ResNet.* The performance of the FL model is assessed using the *accuracy* metric. Energy consumption is measured in kilowatt-hours (kWh), while carbon emissions are computed based on energy usage and the carbon intensity of each node's location, expressed in kilograms of CO_2 equivalents (kg CO_2e).

The framework has been developed using Flower [31], which enables the simulation of an FL setting on a single machine for research purposes. The application will be extended to real-world scenarios involving diverse machines as participants in the FL setting.

For a more detailed discussion of the methodology, readers are referred to the extended paper [32].

4. Framework

The **Eco Federated Learning Framework** is designed to support the methodology introduced in this paper. In particular, it provides to main functionalities:

- the *FL Simulator* feature is used in the initial phases of the methodology. It offers an FL simulation environment, allowing users to explore the effect of data quality on model accuracy by applying data poisoning. This phase automatically collects useful data to find the trade-off between node selection (and data volume reduction) with the FL training task environmental impact;
- the *FL Configuration Recommender* is a service to researchers and practitioners that, given the details of the training data and node infrastructure, recommends node selection and data reduction strategies that allow for reducing the environmental impact while reaching the required model accuracy.

These two features are discussed in more detail in the rest of this section.

The **Eco Federated Learning Framework** is publicly available (https://github.com/POLIMIGreenISE/ecoFL.git). The repository includes all the necessary resources to configure the framework and to test the *FL Simulator* and *FL Configuration Recommender* features. Complete documentation and detailed execution guidelines are provided in the `README.md` file. By making all code, data, and experimental procedures openly accessible, we aim to enhance the reproducibility of our results and foster further research in sustainable Federated Learning.

4.1. FL Simulator

The FL Simulator is a tool that can be leveraged to extract insights on the correlation between data quality dimensions, performance, and energy consumption of the FL training process. The tool allows the execution of multiple experiments by configuring a custom FL training task and initiating a training session. The simulator first requires the user to input the path to the dataset and specify the number of participants in the federated setup (Figure 2). The provided path should point to the training portion of the dataset, which will be distributed among the participants, while the test portion will be used by the central server. The specified number of participants determines how the dataset is partitioned into homogeneous subsets. Each subset is then further divided into training and testing sets in a 4:1 ratio. After these inputs are submitted, the simulator displays the class distribution for each client's dataset partition as well as that of the central server (Figure 3).

For each participant, the simulator provides the option to apply data poisoning on specific data quality aspects, namely data volume, accuracy, consistency and completeness, to the training set (Figure 4). The methodology recommends varying only one data dimension in each sub-experiment within a given experimental setup, although it is also possible to apply data poisoning across all data dimensions simultaneously.

Finally, to run a simulation, it is essential to define both the percentage of participants who actively join the FL training rounds and the mode of their selection. As shown in Figure 5, two modes are available: basic and fixed. In the basic mode, a specified percentage of participants is randomly selected for each training round. In contrast, the fixed mode randomly selects the participants only once during the first round, and the same set of participants is retained for all subsequent rounds. A complete federated learning simulation is executed after specifying the total number of rounds and the file path for storing the simulation results. The results report the accuracy value on the test set of the central server at the last round, emissions, energy consumed and other metadata.

4.2. FL Configuration Recommender

The *FL Configuration Recommender* is a functionality designed to support researchers in the configuration of an FL training task, by selecting an optimal subset of data and participants to achieve a target model accuracy, specified as input, while minimising the environmental impact. The decision-making process

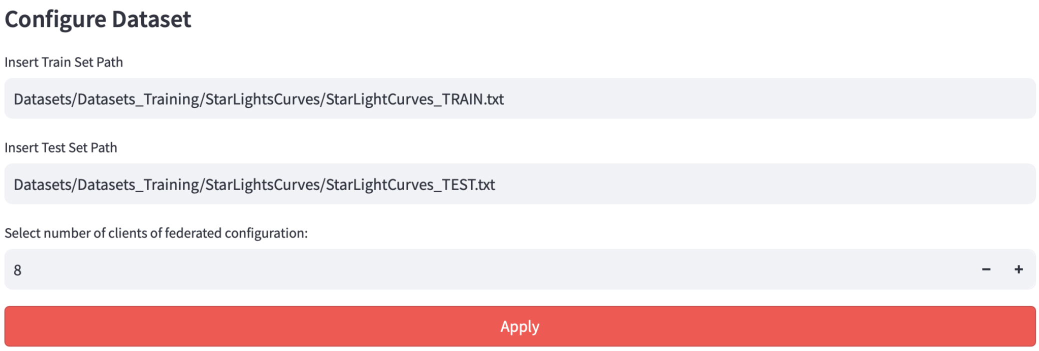

Figure 2: Dataset Input

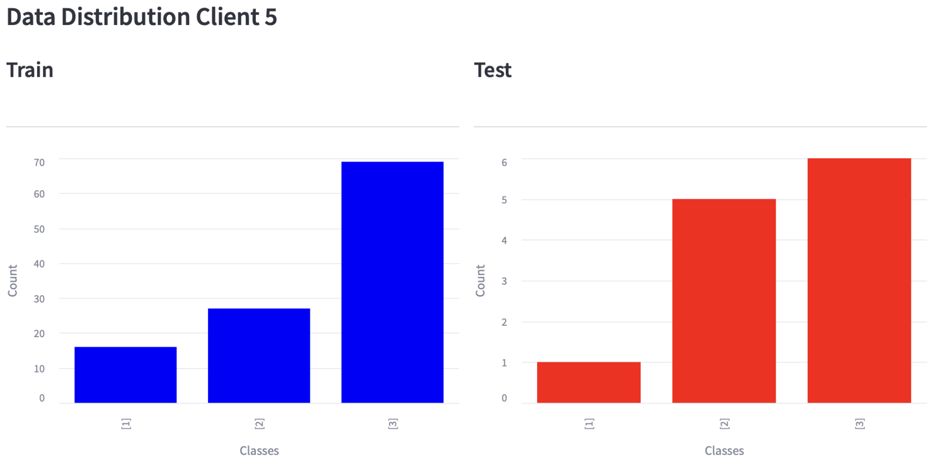

Figure 3: Class Distribution on Participant Node

considers energy mix data, infrastructure specification, and the characteristics of the dataset (i.e., volume and quality) held by each participant.

The *FL Configuration Recommender* implements the *FL Configuration Recommendation* phase of the methodology by configuring and executing a complete federated training. The system exploits the *FL Reduction System*, a regression model that suggests the necessary number of nodes $\hat{N}$ for training. This model has been generated by collecting data from a limited, yet representative, set of experiments following the Data Analysis and Modelling phase. This model and its creation are not described in detail in the current paper because it is out of scope.

The *FL Configuration Recommender* requires the user to provide the path to the dataset to be used for FL training and to specify the participants (Figure 2). In this phase, each participant is characterised by specific attributes, including data volume, data quality dimensions, energy consumption (derived from hardware specifications), geographic location, and associated carbon intensity (Figure 6). Carbon intensity can be inferred from the location using available services (e.g., Electricity Maps[1]). Data volume refers to the percentage of the dataset owned by each participant; therefore, the sum of the data volumes of all participants in the FL configuration must equal the data volume of the training set.

Based on the inputs, the dataset is partitioned between the nodes and the class distribution for both the training and test sets is visualised (Figure 3). It is worth noticing that while this step is essential for the simulation, the datasets will already be distributed among participants in a real-world FL architecture.

The user also specifies an accuracy goal to be achieved by the model generated after the FL training process. Higher accuracy targets may require greater data volume, which in turn necessitates more participants and larger datasets. The *FL Configuration Recommender* ranks participants according to

[1]https://www.electricitymaps.com

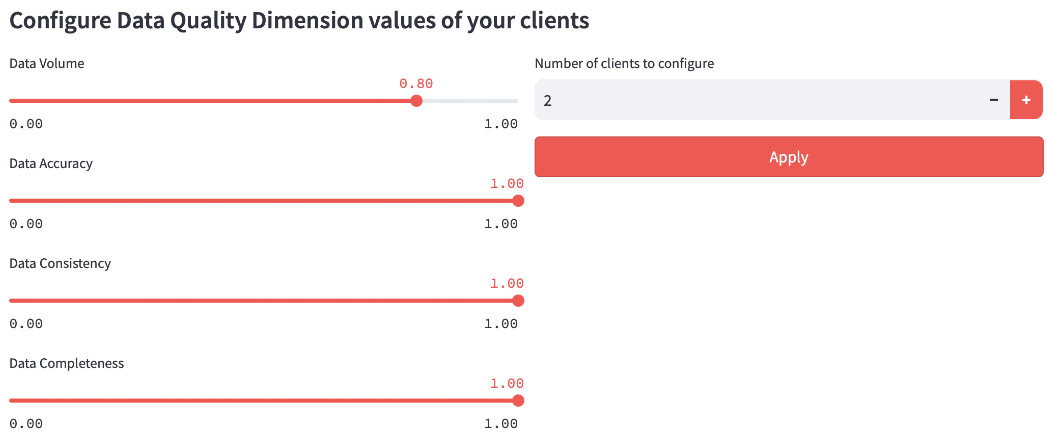

Figure 4: Data Poisoning

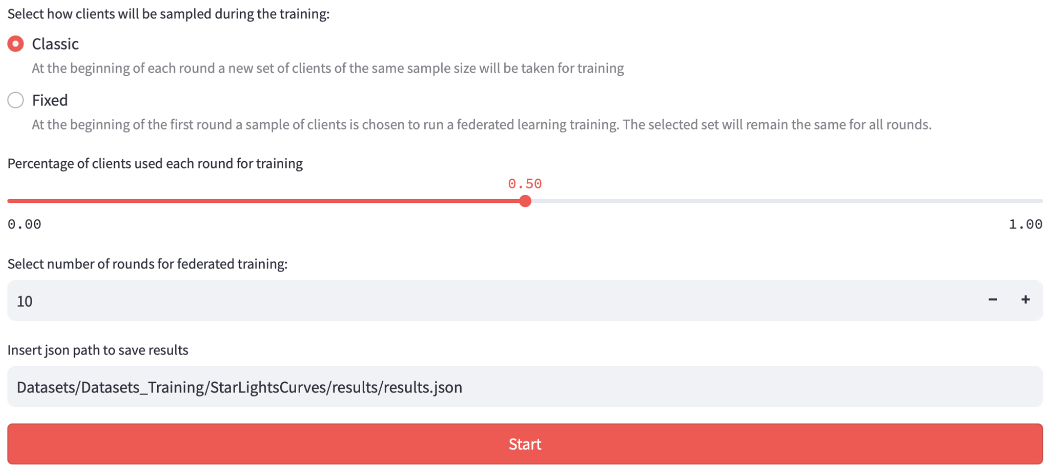

Figure 5: Participant Selection Inputs

their environmental impact and data quality. The system selects a specific set of participants and the subset of data in each node according to the nodes' characteristics, the dataset characteristics, and the accuracy goal. Querying the *FL Reduction System*, the system obtains the predicted number of nodes $\hat{N}$ required to fulfil the accuracy target. This number is subsequently adapted to fit the FL configuration specified by the researcher by translating this number into a *Required Data Volume Percentage* for each node V_n, and a *Target Data Volume Percentage*, that represents the percentage of the dataset that must be retained across all the $\hat{N}$ nodes to guarantee the performance target. The framework supports and implements three different node selection approaches (Figure 7) that can be selected by the user. Each of them applies a different algorithm for selecting the optimal subset of data volume and participants:

- **Node Selection (NS)**: This method selects the first $\hat{N}$ nodes with the highest score. The selected nodes must satisfy the required data volume percentage V_n for each node. If a node n has a dataset larger than the required V_n, its data volume is randomly reduced to match V_n. Conversely, if a node's data volume does not meet the requirement, it is excluded in favour of the next candidate in the ranking.
- **Minimal Smart Reduction (MSR)**: This method follows the same strategy as NS but incorporates data cleaning by removing low-quality (dirty) data. Only clean data are used for training. As a consequence, the actual data volume E, which is the sum of the clean data volumes from the selected nodes, may be lower than the target volume V, potentially impacting model accuracy.
- **Smart Reduction (SR)**: This method extends MSR by ensuring that the target data volume V is

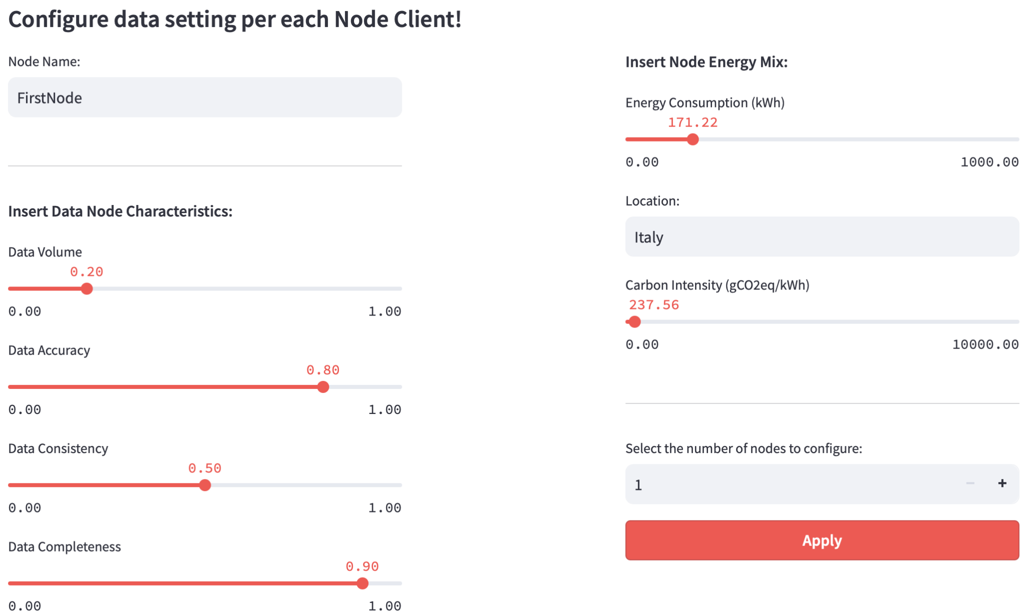

Figure 6: FL Participant Configuration

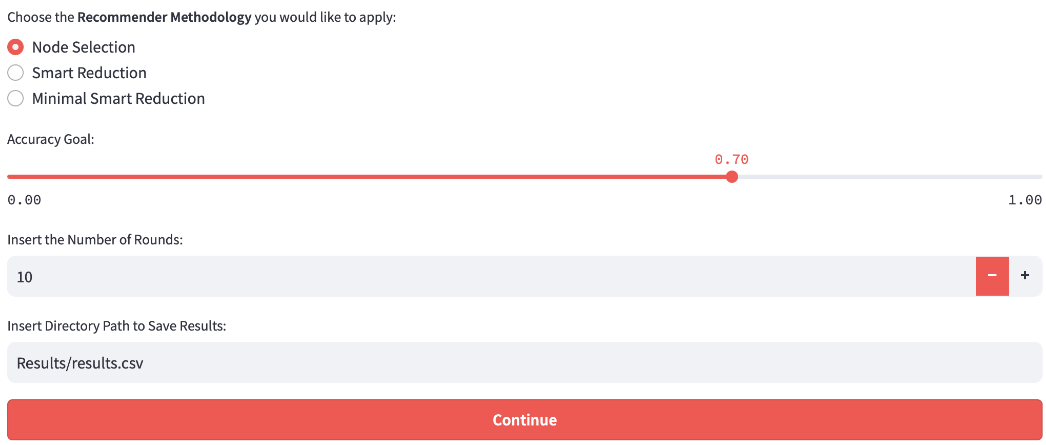

Figure 7: Recommender Methodologies

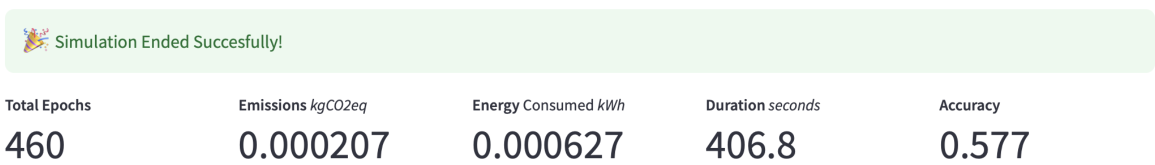

Figure 8: Simulation Results

met. Since removing low-quality data may reduce the total data volume ($E < V$), SR compensates by selecting additional nodes until E reaches or exceeds V.

The three approaches implement different trade-offs between model accuracy and environmental impact (see Section 5).

At the end of the FL training, results are shown reporting the cumulative total epochs across all the participants, the time spent (in seconds) for the FL training process, the carbon emissions produced and energy consumed by the overall system during both pre-processing and training phases, and the final accuracy achieved in predicting a disjoint test set by the trained FL configuration (Figure. 8).

5. Validation

The proposed FL Configuration Recommender assumes the effectiveness of the proposed methodology in reducing environmental impact by lowering carbon emissions while maintaining the desired performance level during FL training. Our research involved conducting experiments to obtain results that support this claim and ensure the satisfaction of performance and sustainability requirements.

For evaluation purposes, three distinct initial FL configurations were tested using three different datasets. Each configuration comprises heterogeneous nodes distributed across various global regions, resulting in differing carbon intensity levels, hardware capabilities—thus varying energy consumption—and data quality. The three recommender methodologies presented were evaluated and benchmarked against a baseline. The *Baseline* approach is defined as training without applying any optimisation techniques; hence, it corresponds to the original FL configuration utilising all available resources and data.

The results were analysed by evaluating each configuration individually, comparing final accuracy and carbon emissions. Validation experiments involved repeating each experiment eight times. Considering three FL configurations, three proposed methodologies, and the Baseline approach applied to each configuration, a total of 96 simulations were conducted for the validation.

All three methods (NS, MSR, and SR) outperform the Baseline in terms of accuracy, with SR achieving the highest accuracy improvement (12%) and meeting the accuracy threshold in 92% of experiments due to selecting more nodes with clean data. NS, prioritising nodes based on data volume without extensive filtering, reaches a 10% improvement, while MSR, selecting identical nodes as NS but using cleaner data, achieves an 8% improvement; both meet the accuracy threshold in 87% of cases.

Regarding carbon emissions, NS is the most energy-efficient, reducing emissions by an average of 56% (up to 90% peak) compared to the Baseline. MSR, requiring longer training with clean data, achieves a 45% reduction, while SR, selecting more nodes, attains only a 25% reduction. Ultimately, NS offers the best balance of accuracy and carbon efficiency.

6. Conclusions

This work introduces a data-centric approach to improve the efficiency of FL systems, supported by a user-friendly framework for simulation and resource optimisation while maintaining desired performance during FL training. The framework offers three methodologies to reduce the environmental impact of FL and underscores the importance of federated and cloud-based data management in optimising distributed machine learning workflows.

Future research can adopt the *FL Simulator* feature to collect additional simulation results and deepen the evaluation of the proposed methods through the *FL Configuration Recommender* usage. As the framework is still under development, several capabilities are possible: adapting it to real-world fog computing scenarios involving multiple machines, generalising it to support a broader range of machine learning tasks, exploring new data selection or reduction strategies, automating data quality assessment, and incorporating additional data quality dimensions to further improve distributed learning efficiency.

Acknowledgments

This work was supported by the project FREEDA, funded by the frameworks PRIN (MUR, Italy) and by the European Union (TEADAL, 101070186).

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT-4 and Grammarly to perform grammar and spelling checks. After using these tools, the authors reviewed and edited the content as needed, and they take full responsibility for the publication's content.

References

[1] D. Rolnick, et al., Tackling Climate Change with Machine Learning, ACM Computing Surveys (CSUR) 55 (2022) 1–96.

[2] W. Knight, AI can do great things - if it doesn't burn the planet, Wired Magazine (2020).

[3] T.-Y. Hsiao, et al., Filter-based Deep-compression with Global Average Pooling for Convolutional Networks, Journal of Systems Architecture 95 (2019) 9–18.

[4] K. He, et al., Deep Residual Learning for Image Recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.

[5] E. Strubell, A. Ganesh, A. McCallum, Energy and policy considerations for modern deep learning research, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 13693–13696.

[6] R. Schwartz, et al., Green AI, Communications of the ACM 63 (2020) 54–63.

[7] S. Georgiou, et al., Green ai: Do deep learning frameworks have different costs?, in: Proceedings of the 44th International Conference on Software Engineering, 2022, pp. 1082–1094.

[8] N. C. Frey, D. Zhao, S. Axelrod, M. Jones, D. Bestor, V. Gadepally, R. Gómez-Bombarelli, S. Samsi, Energy-aware Neural Architecture Selection and Hyperparameter Optimization, in: 2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), IEEE, 2022, pp. 732–741.

[9] R. C. Castanyer, S. Martínez-Fernández, X. Franch, Which design decisions in ai-enabled mobile applications contribute to greener ai?, Empirical Software Engineering 29 (2024) 2.

[10] D. Zha, Z. P. Bhat, K.-H. Lai, F. Yang, Z. Jiang, S. Zhong, X. Hu, Data-centric artificial intelligence: A survey, ACM Computing Surveys 57 (2025) 1–42.

[11] J. Jakubik, M. Vössing, N. Kühl, J. Walk, G. Satzger, Data-centric artificial intelligence, Business & Information Systems Engineering 66 (2024) 507–515.

[12] Y. Shin, et al., Practical Methods of Image Data Preprocessing for Enhancing the Performance of Deep Learning Based Road Crack Detection, ICIC Express Letters, Part B: Applications 11 (2020) 373–379.

[13] V. Werner de Vargas, et al., Imbalanced Data Preprocessing Techniques for Machine Learning: a Systematic Mapping Study, Knowledge and Information Systems 65 (2023) 31–57.

[14] C. Chai, J. Wang, Y. Luo, Z. Niu, G. Li, Data management for machine learning: A survey, IEEE Transactions on Knowledge and Data Engineering 35 (2023) 4646–4667. doi:10.1109/TKDE.2022.3148237.

[15] F. Lucivero, Big data, big waste? A reflection on the environmental sustainability of big data initiatives, Science and Engineering Ethics 26 (2020) 1009–1030.

[16] C. Sun, et al., Revisiting Unreasonable Effectiveness of Data in Deep Learning Era, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 843–852.

[17] R. Verdecchia, L. Cruz, J. Sallou, M. Lin, J. Wickenden, E. Hotellier, Data-centric green ai an exploratory empirical study, in: 2022 International Conference on ICT for Sustainability (ICT4S), IEEE, 2022, p. 35–45. URL: http://dx.doi.org/10.1109/ICT4S55073.2022.00015. doi:10.1109/ict4s55073.2022.00015.

[18] L. Berti-Equille, Learn2clean: Optimizing the Sequence of Tasks for Web Data Preparation, in: The World Wide Web Conference, 2019, pp. 2580–2586.

[19] A. Jain, H. Patel, L. Nagalapatti, N. Gupta, S. Mehta, S. Guttula, S. Mujumdar, S. Afzal, R. Sharma Mittal, V. Munigala, Overview and importance of data quality for machine learning tasks, in: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining, 2020, pp. 3561–3562.

[20] N. Gupta, S. Mujumdar, H. Patel, S. Masuda, N. Panwar, S. Bandyopadhyay, S. Mehta, S. Guttula, S. Afzal, R. Sharma Mittal, et al., Data quality for machine learning tasks, in: Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining, 2021, pp. 4040–4041.

[21] L. Budach, M. Feuerpfeil, N. Ihde, A. Nathansen, N. Noack, H. Patzlaff, F. Naumann, H. Harmouch, The effects of data quality on machine learning performance, arXiv preprint arXiv:2207.14529

(2022).
[22] M. Anselmo, M. Vitali, A data-centric approach for reducing carbon emissions in deep learning, in: International Conference on Advanced Information Systems Engineering, Springer, 2023, pp. 123–138.
[23] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečnỳ, S. Mazzocchi, B. McMahan, et al., Towards federated learning at scale: System design, Proceedings of machine learning and systems 1 (2019) 374–388.
[24] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, B. He, A survey on federated learning systems: Vision, hype and reality for data privacy and protection, IEEE Transactions on Knowledge and Data Engineering 35 (2023) 3347–3366. doi:10.1109/TKDE.2021.3124599.
[25] Y. Liu, Y. Kang, T. Zou, Y. Pu, Y. He, X. Ye, Y. Ouyang, Y.-Q. Zhang, Q. Yang, Vertical federated learning: Concepts, advances, and challenges, IEEE Transactions on Knowledge and Data Engineering 36 (2024) 3615–3634. doi:10.1109/TKDE.2024.3352628.
[26] Y. Wu, N. Xing, G. Chen, T. T. A. Dinh, Z. Luo, B. C. Ooi, X. Xiao, M. Zhang, Falcon: A privacy-preserving and interpretable vertical federated learning system, Proceedings of the VLDB Endowment 16 (2023) 2471–2484.
[27] X. Qiu, T. Parcollet, J. Fernandez-Marques, P. P. Gusmao, Y. Gao, D. J. Beutel, T. Topal, A. Mathur, N. D. Lane, A first look into the carbon footprint of federated learning, Journal of Machine Learning Research 24 (2023) 1–23.
[28] P. Wiesner, R. Khalili, D. Grinwald, P. Agrawal, L. Thamsen, O. Kao, Fedzero: Leveraging renewable excess energy in federated learning, in: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems, 2024, pp. 373–385.
[29] A. Abbasi, F. Dong, X. Wang, H. Leung, J. Zhou, S. Drew, Fedgreen: Carbon-aware federated learning with model size adaptation, in: 2024 IEEE International Conference on Communications Workshops (ICC Workshops), IEEE, 2024, pp. 1352–1358.
[30] Y. Mao, X. Yu, K. Huang, Y.-J. A. Zhang, J. Zhang, Green edge ai: A contemporary survey, Proceedings of the IEEE (2024).
[31] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, J. Fernandez-Marques, Y. Gao, L. Sani, H. L. Kwing, T. Parcollet, P. P. d. Gusmão, N. D. Lane, Flower: A friendly federated learning research framework, arXiv preprint arXiv:2007.14390 (2020).
[32] M. Sabella, M. Vitali, Eco-Friendly AI: Unleashing Data Power for Green Federated Learning, 2025. URL: https://arxiv.org/abs/2507.17241. arXiv:2507.17241.