International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Can federated learning enable environmentally sustainable AI by reducing centralized data storage and transmission costs?

Asbah Ayub

Abstract:
AI uses a lot of electricity because we store huge datasets in data centres and move many gigabytes across networks. The International Energy Agency (IEA) projects that electricity use by data centres could more than double to roughly 945 TWh by 2030, with AI as a key driver. Data transmission networks also consume large amounts of power—about 260–360 TWh in 2022—so every extra gigabyte moved has an energy and carbon cost.
Federated learning (FL) trains models by keeping data on local devices or sites and sending only model updates to a server. In simple terms, this can cut centralised storage and reduce raw-data transfers. If fewer raw bytes are shipped and stored centrally, we can lower some network and storage energy. However, federated learning(FL) adds new costs: more rounds of communication, on-device computation, and the carbon intensity of where and when devices and servers run. Recent research shows the result is not automatic—depending on settings, FL can even emit much more carbon than centralised training (for example, with many communication rounds, poor device efficiency, or carbon-heavy grids).This paper will tests a simple idea that is federated learning(FL) can support environmentally sustainable AI when designs minimise communication (e.g., compression, fewer rounds), schedule training in low-carbon locations and times, and use energy-efficient devices and servers. Emerging "green FL" methods and carbon-aware schedulers try to do this by adapting model size and timing to local grid carbon intensity.

Keywords: federated learning; sustainable AI; data centres; network energy; carbon footprint; communication efficiency; carbon-aware scheduling.

Introduction

The rapid rise of artificial intelligence (AI) is reshaping how we gather, store, and process data—and with that, the environmental footprint of AI is growing. Global data centres today consume around **415 terawatt-hours (TWh)** of electricity annually, which is about 1.5 % of all global electricity us[1]e.The International Energy Agency (IEA) projects that by 2030 this number could more than double to **about** 945 TWh per year, driven mostly by AI workloads. This data-intensity matters because large amounts of data storage, cooling, networking equipment, and data transmission all consume energy, often from carbon-intensive sources[2].[3]Traditionally, many AI systems rely on a centralised architecture: vast raw data sets are collected in one place (a data centre or cloud server), stored, and then processed centrally. In this model the storage and transmission costs—both in terms of energy and carbon—are high, because every

[1] Rice, S. (2025, April 14). AI, energy and the data centre surge: IEA's global warning. Data Centre Magazine. https://datacentremagazine.com/technology-and-ai/iea-asks-how-is-ai-reshaping-global-energy-demand
[2] **Chen, S., & Nature Magazine. (2025, April 10).** ***Data centers will use twice as much energy by 2030—Driven by AI. Scientific American.*** https://www.scientificamerican.com/article/ai-will-drive-doubling-of-data-center-energy-demand-by-2030
[3] International Energy Agency. (2024, March 7). AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works. IEA. https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works

bit of data must travel from its origin to the central site and be held there, potentially with backups and redundancies. By contrast, a distributed approach keeps data at its origin (for example on individual devices or at local "edge" servers) and only transfers processed summaries or model updates. One such approach is federated learning (FL), which allows multiple devices or sites to collaboratively train an AI model without sending raw data to a central server. Instead, each device computes locally and sends only model adjustments.In the context of environmental sustainability in AI, the core idea is that if less raw data is transmitted and less centralised storage is required, the energy and carbon footprint of AI systems could potentially be reduced. Environmental sustainability here refers to using AI methods in a way that minimises negative impacts on the planet—such as greenhouse gas emissions and resource depletion—while maintaining performance and utility[4]. The question then becomes: Can federated learning enable environmentally sustainable AI by reducing centralised data storage and transmission costs?

I chose this topic because the energy demand of AI is increasing very fast, and few people think about how our digital habits affect the environment. I wanted to explore whether new methods like federated learning can make AI greener by cutting data storage and transmission costs. This research is useful for both individuals and society because it helps us understand how to build smarter and more energy-efficient technologies. If companies and users adopt sustainable AI systems, it can reduce carbon emissions, save energy costs, and create a cleaner digital future for everyone.

Figure -Data-centre energy growth

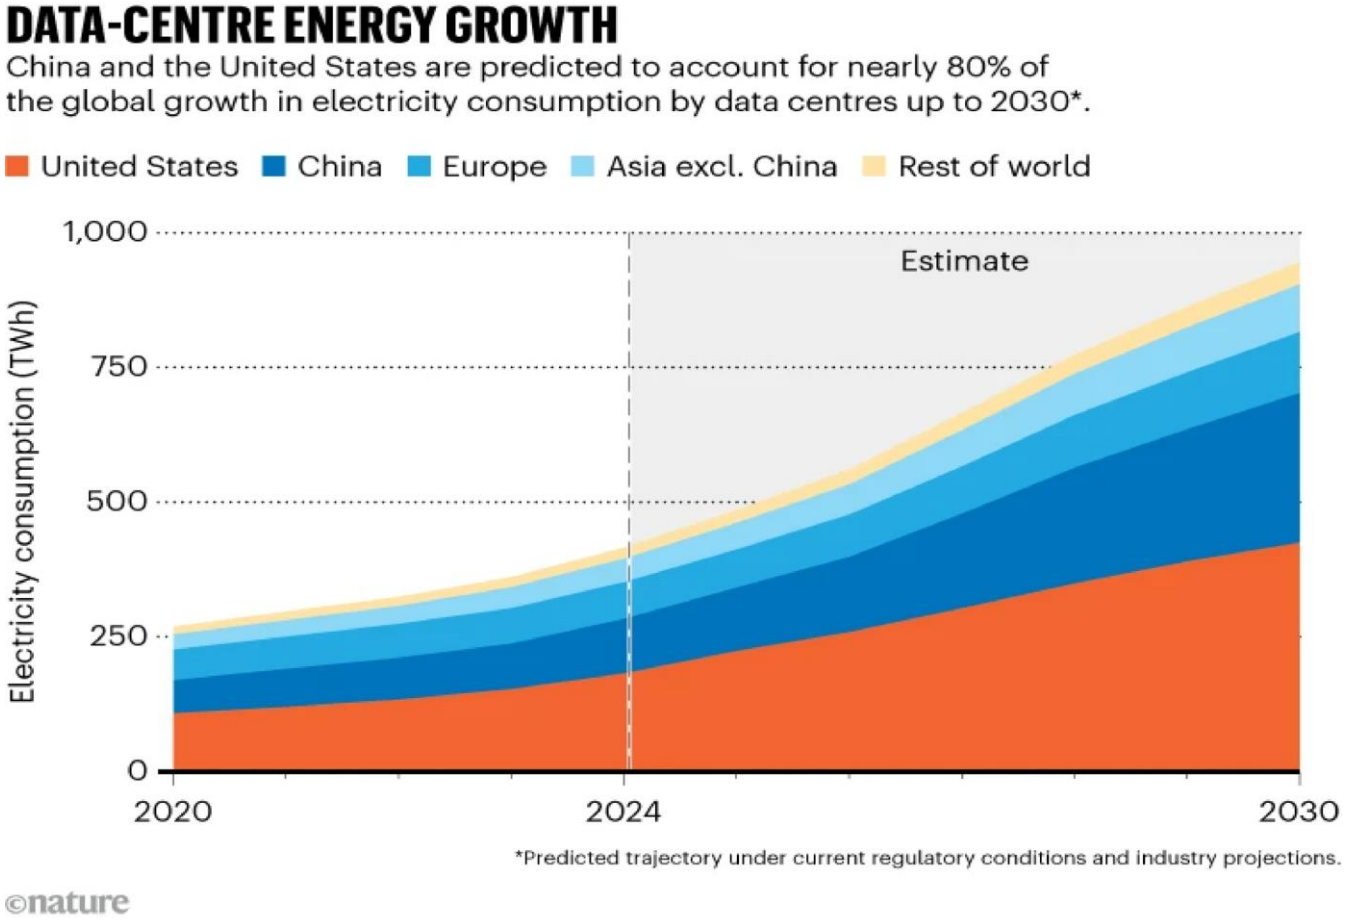

[4] .World Economic Forum. (2023, September 20). How AI can help build a sustainable future. WEF. https://www.weforum.org/stories/2025/07/ai-infrastructure-sustainable-future-energy/

Chen, S., & Nature Magazine. (2025, April 10). ***Data centers will use twice as much energy by 2030—Driven by AI. Scientific American.*** **https://www.scientificamerican.com/article/ai-will-drive-doubling-of-data-center-energy-demand-by-2030/**

Background and Literature Review

In recent years, modern artificial intelligence (AI) systems — especially those using deep learning — have grown increasingly reliant on large-scale, centralized data storage and massive amounts of data transmission. These models often require enormous datasets collected from many users or devices, stored in central data centres, and then processed using powerful servers. The storage of raw data, continuous data flow from devices into clouds, and the heavy computation all add up to significant energy use. For instance, the International Energy Agency (IEA) reports that data centres accounted for about 1.5% of global electricity consumption (415 TWh) in 2024, and this figure is projected to more than double to around 945 TWh by 2030 under current trends[5]. In the United States alone, a report by the Lawrence Berkeley National Laboratory shows data centres consumed about 4.4% of U.S. electricity in 2023, with projections that they could reach 6.7% to 12% by 2028. (U.S. Department of Energy[6].These numbers illustrate how centralised AI training and data-storage practices are stretching energy systems, and by extension, raising concerns about carbon emissions and environmental sustainability.

The concept of reducing data movement and centralised storage has gained traction as a potential way to mitigate these environmental impacts. One key alternative is **federated learning (FL)**: a distributed machine-learning approach in which a global model is trained across many devices or data nodes, while the raw data remains locally stored and only model updates (for example gradients or parameter changes) are sent to a coordinating server. Device-level training occurs locally, the key model updates are aggregated centrally, and then the updated global model is redistributed. This approach is used in two main variations: *cross-device* FL (many user devices, e.g., smartphones) and *cross-silo* FL (fewer, larger entities such as hospitals or enterprises each with their own local dataset). The architecture reduces reliance on shipping large volumes of raw data from each device to the cloud, thereby potentially lowering network transmission volumes and central storage requirements.

From an environmental sustainability perspective, this shift could be meaningful. "Sustainability" in the AI context refers to designing AI systems that deliver value (innovation, insights, services) while minimising negative environmental impacts — such as greenhouse-gas emissions, high energy use, and resource depletion (e.g., cooling water, rare earth materials). When large data transfers and centralised server storage are reduced, the energy associated with transmission infrastructure, data-centre cooling, and memory/hardware usage can also drop. However, the transition is not automatically green: federated learning also has its own costs — more frequent communication rounds among devices, local on-device computation (which may be less efficient), and potential duplication of some processing tasks across nodes.

Prior research has begun to explore this trade-off. A landmark study titled *"A First Look into the Carbon Footprint of Federated Learning"* found that depending on configuration, Federated learning (FL) can emit up to **two orders of magnitude more carbon** than a centralized training setup — especially when communication rounds are many, when many local devices are involved, and when the devices or networks are inefficient[7]. At the same time, the paper notes that in certain settings — for example when

[5] .(International Energy Agency. (2024). *Energy and AI: Executive summary – The transformative potential of AI depends on energy*. *IEA*. https://www.iea.org/reports/energy-and-ai/executive-summary)

[6] (2024, December 20). DOE releases new report evaluating increase in electricity demand from data centers. Energy.gov. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers)

[7] (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). A first look into the carbon footprint of federated learning. Journal of Machine Learning Research, 24, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf)

local devices are energy-efficient, communication volumes are low, and the local grid uses cleaner energy — the carbon footprint of Federated learning (FL) can be comparable to centralised approaches. Building on that, newer work focuses on *green federated learning* or *carbon-aware Federated learning (FL)*, for example a framework named "EcoLearn" achieved a reduction of carbon emissions by up to ~93% compared to a naïve Federated learning (FL) baseline by selecting low-carbon devices and optimising communication schedules[8]. Other research proposes systems like "CAFE: Carbon-Aware Federated Learning" which tailor node selection and learning rounds according to geographic carbon intensity. These studies highlight that to make federated learning environmentally beneficial, one cannot simply adopt it blindly — the design, hardware efficiency, communication overhead, and regional energy-mix all matter significantly[9].

Global electricity growth

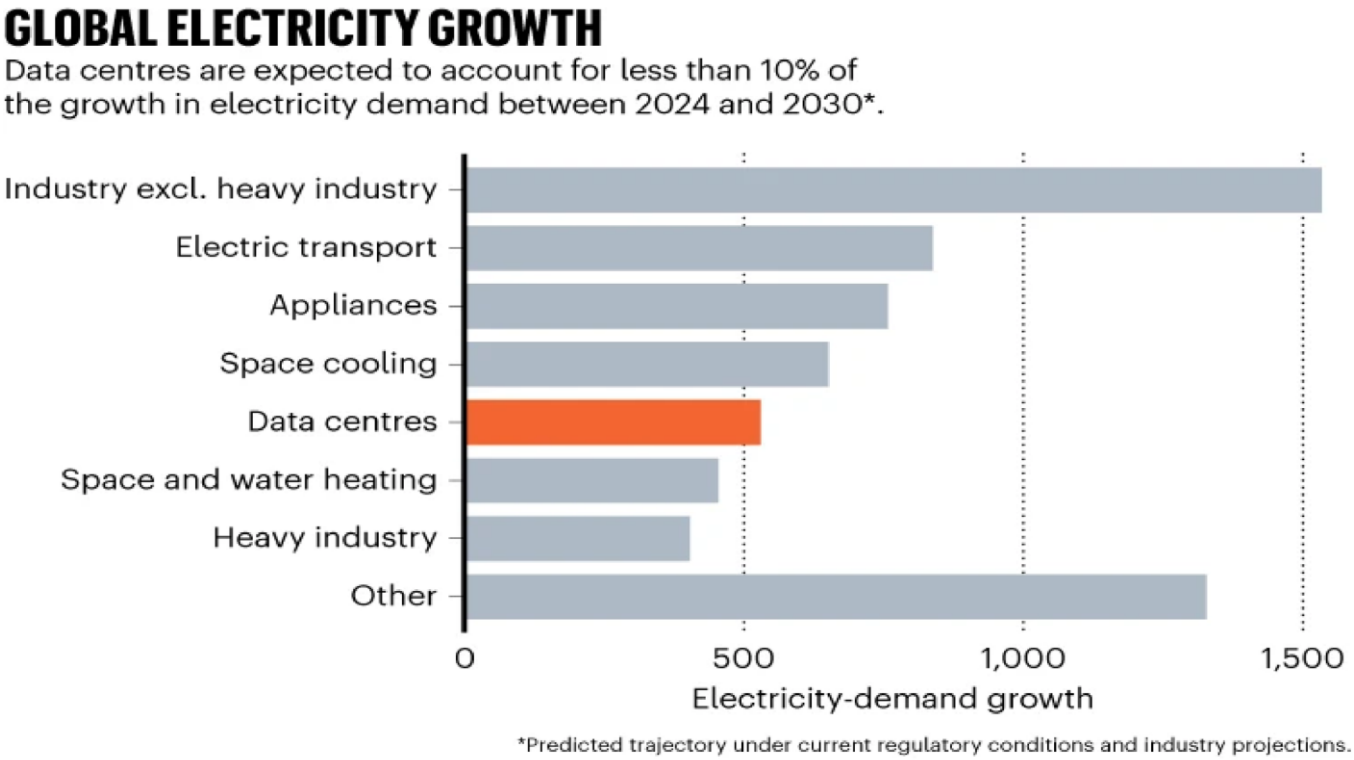

Altogether, the literature indicates that while centralised AI training and data storage are clearly energy-intensive and raise sustainability concerns, federated learning offers a promising alternative **provided** system design is optimized for low transmission volumes, efficient local computation, and low-carbon energy. Researchers consistently highlight the need for standardised metrics of carbon footprint in ML, for transparency about energy consumption in AI workloads, and for joint optimization of model

[8] (Bian, J., Wang, L., Ren, S., & Xu, J. (2024, February 6). CAFE: Carbon-aware federated learning in geographically distributed data centers (Version 2). arXiv preprint arXiv:2311.03615. https://arxiv.org/abs/2311.03615)

[9] .(Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024). 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2001637.pdf eta-publications.lbl.gov)

performance and environmental cost[10]. For example, the "Green Algorithms" framework provides a tool to estimate carbon emissions of computation tasks by accounting for location, hardware type, energy source, and runtime[11]. This body of work establishes the foundation for asking: under what specific conditions can federated learning truly enable environmentally sustainable AI by reducing centralised data storage and transmission costs?

3. Analytical Framework

This research paper uses a comparative analytical framework to examine whether federated learning (FL) can make artificial intelligence (AI) more environmentally sustainable by reducing centralized data storage and transmission costs. The approach compares two main system models—centralized learning and federated learning—and studies how their design affects energy use and carbon emissions. In centralized learning, large volumes of raw data are collected and stored in data centers for model training, leading to high storage and data transmission costs. In contrast, federated learning keeps the data at its original source and sends only model updates, which can reduce the amount of information transmitted and stored centrally. The framework therefore links three key stages: design choice, energy consumption, and environmental impact. The assumption is that if federated learning reduces the amount of data transferred and stored, then total energy use and carbon emissions will also decline.

To study this, the research defines measurable indicators such as data volume transferred, energy used for storage and transmission, and total carbon emissions. Energy use is measured in kilowatt-hours (kWh) and converted into carbon dioxide equivalent (CO_2e) using emission factors from the International Energy Agency (IEA, 2023). For transmission, the study estimates the total gigabytes of data moved across networks, multiplied by the average energy intensity of data transfer (kWh per GB). For storage, it measures the number of terabytes of data stored per month in centralized data centers, multiplied by energy intensity per terabyte-month. For computation, energy use is measured by the time required for device or server processing and the power consumed per device-hour. These three forms of energy use—network, storage, and computation—are then converted into CO_2e based on regional electricity emission factors. The total environmental cost of each method is therefore estimated as the sum of all three categories.

The analysis focuses on key trade-offs between centralized and federated systems. While federated learning reduces the need for central data storage and raw data transfers, it also creates new costs such as repeated communication rounds, device-level computation, and aggregation at the central server. If model updates are large and require frequent synchronization, the network energy use of FL may approach or even exceed that of centralized systems. On the other hand, if compression methods, selective participation, and fewer communication rounds are used, Federated learning (FL) can significantly lower total transmission energy. Similarly, while centralized data centers are usually more energy efficient, edge devices in Federated learning (FL) systems can benefit from lower-carbon electricity in local grids. Therefore, the balance between computation and communication is a key variable that determines whether Federated learning (FL) actually reduces total energy and carbon costs.

The research adopts a mixed-method and literature-based design. It reviews recent scientific and technical studies from peer-reviewed sources such as the *Journal of Machine Learning Research*[12] and the

[10] (Mytton, D. (2025, February 9). Data centre energy consumption data-centre energy and AI in 2025. dev/sustainability. https://www.devsustainability.com/p/data-center-energy-and-ai-in-2025 devsustainability.com)
[11] (Lannelongue, L., Grealey, J., & Inouye, M. (2020, December 17). Green algorithms: Quantifying the carbon footprint of computation (Version 5). arXiv preprint arXiv:2007.07610. https://arxiv.org/abs/2007.07610)
[12] (Qiu et al., 2023) (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). A first look into the carbon footprint of federated learning. Journal of Machine Learning Research, 24, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf)

Foundations and Trends in Machine Learning[13] , as well as government and industry reports from the IEA and the Lawrence Berkeley National Laboratory[14]. These studies provide data on energy intensity for data storage, data transmission, and computation in AI systems. The inclusion criteria are that sources must provide quantitative estimates of energy or carbon costs, focus on federated or distributed learning systems, and use transparent methodologies. Non-technical or speculative sources are excluded. In addition, emission factors and energy intensity benchmarks are taken from credible institutions such as the IEA and U.S. Department of Energy to ensure consistency and comparability.

Three main scenarios are used for comparison. The first is **Centralized Learning (CEN)**, where all data is uploaded once, stored in large data centers, and used for model training in centralized servers. The second is **Naive Federated Learning (FL-N)**, where many clients participate in every communication round without compression or optimization. The third is **Green Federated Learning (FL-G)**, which includes strategies like communication compression, partial client selection, adaptive training rounds, and carbon-aware scheduling—meaning training is performed when and where electricity is cleaner. For each scenario, energy and emissions are estimated separately for storage, transmission, and computation. These are compared to evaluate how design and optimization choices affect total environmental performance.Finally, this framework also considers uncertainty and sensitivity. Energy intensities vary across studies, regions, and hardware types, so the analysis presents results in ranges rather than single figures. Sensitivity analysis is used to test how changing parameters such as the number of communication rounds, size of model updates, or carbon intensity of the grid affects the total outcome. The methodology also uses standardized carbon accounting tools such as the Green Algorithms model (Lannelongue et al., 2021), which helps to calculate emissions based on hardware power, runtime, and energy source. This transparent, step-by-step framework allows for an objective comparison of centralized and federated systems and helps identify under what conditions federated learning can truly enable environmentally sustainable AI.

Evidence & Case Analysis

Modern AI relies on moving and storing very large amounts of data, and this has clear energy and carbon costs. Centralised training pipelines collect raw data from many sources, ship it to cloud data centres, and keep multiple replicas for reliability. The International Energy Agency (IEA) estimates that global electricity use by data centres could double to about **945 TWh by 2030**, approaching 3% of world electricity consumption, with AI a major driver. This is in addition to the energy used by data-transmission networks themselves, which consumed **~260–360 TWh in 2022**. When we aggregate these loads, each extra terabyte stored or gigabyte transmitted has a non-trivial footprint, especially in regions with carbon-intensive grids[15].

In a centralised design, every data source sends raw bytes to a small number of large facilities. The storage layer must hold the full dataset (often with replication), while the training layer reads this data repeatedly during model development. Two consequences follow. First, **network energy scales with ingress volume**: when billions of images, sensor readings, or text samples are uploaded, the cumulative energy of transport increases in proportion to total GB moved over fixed networks and mobile infrastructure.

[13] (Kairouz et al., 2021)(Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). *Advances and open problems in federated learning*. Foundations and Trends® in Machine Learning, 14(1–2), 1–210. https://arxiv.org/abs/1912.04977)

[14] Shehabi et al., 2024)(Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., … Sartor, D. (2024). *United States data center energy usage report*. Lawrence Berkeley National Laboratory. https://eta-publications.lbl.gov/sites/default/files/lbnl-2001637.pdf)

[15] International Energy Agency. (2024, March 7). AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works. IEA. https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works

Second, storage energy scales with TB-months: long retention periods, multiple replicas, and hot storage tiers raise the continuous power draw for disks, memory, and cooling. Because AI workloads intensify read/write cycles and require high-performance storage, they can push the data-centre power envelope even further. This is why the IEA's projections grow faster for data centres than for most other sectors through 2030, and why network energy remains material despite efficiency gains[16].Federated learning (FL) changes the flow: data stays local and only model updates (parameters/gradients) travel. In cross-device FL (e.g., phones), training runs on devices; in cross-silo Federated learning (FL) (e.g., hospitals, banks), training runs on organisational servers. Either way, the central server aggregates updates rather than ingesting raw data. This has two potential environmental gains. First, central storage load can shrink because the global system does not warehouse the raw dataset—only the model, metadata, and temporary aggregates. Second, network traffic can shift from large one-off raw uploads to smaller, repeated update exchanges. If the total size of all updates is less than the raw data that would have been centralised—and if updates are compressed or sparsified—then total network energy can fall. A University of Cambridge overview explains the basic mechanism and argues that, in some settings, distributing training over many devices can reduce emissions relative to training everything in data centres[17].
In practical deployments, the savings depend on update size and number of rounds. For example, if a 50-MB model is updated over 50 rounds with 1,000 clients per round but only 10% participate each time, the total traffic might still be large; however, careful client sampling, update compression/quantisation, and fewer rounds with more local epochs can drastically reduce total bytes on the wire. Because raw datasets can be many hundreds of gigabytes to terabytes, there are realistic cases (e.g., on-device next-word prediction, keyboard suggestion models) where keeping data local and exchanging only compact updates leads to **lower cumulative transmission** than centralising everything.

when Federated learning (FL) can be worse

The literature also warns that Federated learning (FL) is not automatically greener. The most cited early study on this topic is *A First Look into the Carbon Footprint of Federated Learning* (Qiu et al., 2023). It offers a quantitative model and experiments showing that depending on configuration, Federated learning (FL) can emit up to two orders of magnitude (≈100×) more carbon than centralised training. Why? There are several reasons: (1) many communication rounds inflate aggregate traffic; (2) inefficient or battery-powered edge devices may draw more energy per unit of useful compute than modern data-centre GPUs/TPUs; (3) poor network links increase the energy per GB (especially over mobile networks); (4) carbon-intensive local grids amplify emissions from both device compute and communication. The same study also notes that, in some settings—particularly when embedded devices are efficient and when communication is restrained—FL's footprint can be comparable to centralised learning. In short: design choices and context determine outcomes[18].These findings match operational experience: naïvely configured Federated learning (FL) (large, frequent updates; all clients every round; no compression) can easily overshoot network energy compared to the one-time upload of raw data to the cloud. Moreover, edge heterogeneity (different chips, batteries, throttling) can slow convergence, requiring more rounds and thus more traffic. Finally, centralised data centres sometimes run on lower-carbon power or use smart procurement of renewables, while phones or branch servers may charge or operate on grids with higher emission factors, erasing expected gains.

[16] (International Energy Agency. (2024). **Energy demand from AI**. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai**)**
[17] (International Energy Agency. (2023). **Data centres and data transmission networks**. https://www.iea.org/reports/data-centres-and-data-transmission-networks)
[18] . (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). **A first look into the carbon footprint of federated learning**. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf)

Recent developments: towards "Green Federated learning (FL)"

Responding to these limits, recent work proposes Green Federated Learning—designs that explicitly minimise emissions while maintaining model quality. Thakur et al. (2024) survey more than a hundred Federated learning (FL) papers and synthesise strategies such as energy-aware client selection, update sparsification/quantisation, adaptive local epochs, early stopping, and data/participant pruning to lower total communication and computation. The survey frames sustainability as a first-class objective alongside accuracy and privacy, especially in IoT contexts where devices and networks are resource-constrained. [19]

A second strand is carbon-aware Federated learning (FL) orchestration. Bian et al. (2023) introduce CAFE (Carbon-Aware Federated Learning) for geo-distributed data centres: the controller selects where and when to train based on real-time carbon intensity of regional grids, subject to a carbon budget. Using tools like coreset selection (to preserve learning quality with fewer, better-chosen data contributions) and a control-theoretic scheduler (Lyapunov drift-plus-penalty), CAFE cuts emissions while maintaining accuracy. This directly targets two drivers from Qiu et al. (2023): number of rounds and carbon intensity by location/time[20]. Related efforts from the University of Cambridge include a carbon calculator and interactive tools to estimate Federated learning (FL) emissions under different hardware profiles, countries, numbers of rounds, and client selections. Such tools make it easier to plan low-carbon training **schedules** and to **quantify the benefit** of compression and participation strategies before deployment[21].

When does Federated learning (FL) reduce environmental impact?

The overall evidence from current research suggests that federated learning (FL) can reduce environmental impact, but only under specific and well-controlled conditions. Federated learning (FL) has the potential to lower centralized data storage and raw data transmission needs by keeping data local and only sharing model updates. However, total emissions depend on five major interacting factors that determine whether the system becomes more sustainable or not.

The first factor is communication volume, which is determined by the size of model updates, the number of communication rounds, and the number of participating clients. When the total size of all updates remains smaller than the raw data that would have been uploaded to a central server, and when compression techniques such as 8-bit quantization, sparsity, and sketching are used, Federated learning (FL) can significantly reduce network energy use. Additionally, using partial client participation—where only a subset of clients contribute per round—further reduces total communication costs. Conversely, large and frequent updates can quickly erase any energy savings achieved[22].

The second factor is **compute efficiency**, which refers to the energy effectiveness of devices performing the training. While data centres use highly optimized hardware such as GPUs and TPUs, edge devices can still perform competitively if models are properly sized and if more local training (local epochs) reduces the number of total communication rounds. When local computations become too heavy for edge devices, the energy used per unit of training progress increases sharply, leading to higher emissions. To counter this, hybrid models that combine light local computation with efficient central aggregation can maintain balance and sustainability[23].

---

[19] (Thakur, D., Guzzo, A., Fortino, G., & Piccialli, F. (2024). **Green federated learning: A new era of green-aware AI**. *arXiv*. https://arxiv.org/abs/2409.12626)

[20] (Bian, J., Wang, L., Ren, S., & Xu, J. (2023). **CAFE: Carbon-Aware Federated Learning in geographically distributed data centres**. *arXiv*. https://arxiv.org/abs/2311.03615)

[21] (University of Cambridge, Machine Learning Systems Lab. (n.d.). **Carbon dioxide in federated learning (interactive tool)**. https://mlsys.cst.cam.ac.uk/carbon_fl/)

[22] (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf).[22]

[23] (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf).

The third factor is carbon intensity by location and time. The environmental benefit of Federated learning (FL) improves when training is scheduled during low-carbon energy periods, such as times of high renewable output (solar or wind), or in regions with cleaner energy grids. Carbon-aware scheduling systems, such as the CAFE (Carbon-Aware Federated Learning) framework, can dynamically choose clients and locations based on the carbon intensity of electricity in real time. These adaptive strategies allow organizations to stay within a predefined carbon budget while maintaining model accuracy[24].
The fourth factor concerns storage patterns and retention policies. In cases where a centralized pipeline would store raw data for extended periods, Federated learning (FL) can save energy by preventing continuous storage and cooling demands. However, if organizations still upload partial data or logs for auditing and debugging, the storage advantage becomes smaller. This means that institutional policies and governance decisions—such as rules about data retention and redundancy—play a major role in determining actual sustainability benefits.

The fifth factor involves **model and data characteristics**, such as model size, data heterogeneity, and non-independent and identically distributed (non-IID) data. Large models with frequent update cycles are difficult to make environmentally efficient unless strong compression and sub-sampling are applied. Moreover, heterogeneous or non-IID data can slow convergence, requiring additional communication rounds and hence more energy. New approaches like coreset selection, knowledge distillation, and personalized Federated learning (FL) models can help reduce communication overhead while maintaining model performance.[25]
When we combine all these findings, a clear pattern emerges. Federated learning is likely to reduce net emissions compared to centralized AI systems when several conditions are met. The dataset must be extremely large—hundreds of gigabytes or terabytes—that would otherwise be transmitted and stored centrally for long periods. Model updates should be small due to compression and sparsity, and the total number of communication rounds must be limited using local epochs or selective participation. Devices performing training must be energy-efficient, and the orchestration should use carbon-aware scheduling that moves computation to cleaner energy grids. Finally, organizations must avoid "silent centralization," where they continue uploading data copies to the cloud for safety or backup, as that negates the main environmental benefit of Federated learning (FL).
On the other hand, Federated learning (FL) can increase emissions when updates are large and frequent, when devices are inefficient, when networks are poor (e.g., mobile uplinks), or when local grids are carbon-intensive and training is not scheduled wisely. These results explain why published findings range from "Federated learning (FL) can be greener" (under careful design) to "Federated learning (FL) can be much worse" (under naïve settings). The key point is that Federated learning (FL) is a lever, not a guarantee: it enables sustainability if we optimise for it.
For completeness, a table comparing representative studies—IEA for macro energy baselines (data centres, networks), Cambridge's explanatory note and tools (mechanism and planning), Qiu et al. (risk of high Federated learning (FL) emissions), Thakur et al. (surveyed strategies), and Bian et al. (carbon-aware scheduling)—can be placed in the appendix. The table can list each study's context, method, main result,

[24] (Bian, J., Wang, L., Ren, S., & Xu, J. (2024, February 6). *CAFE: Carbon-aware federated learning in geographically distributed data centers* (Version 2). *arXiv preprint arXiv:2311.03615*. https://arxiv.org/abs/2311.03615)
[25] (Bian, J., Wang, L., Ren, S., & Xu, J. (2024, February 6). *CAFE: Carbon-aware federated learning in geographically distributed data centers* (Version 2). *arXiv preprint arXiv:2311.03615*. https://arxiv.org/abs/2311.03615)

and levers (e.g., compression, client selection, carbon-aware timing). This visual will help readers quickly see why results differ and which design choices matter most[26].[27]

Discussion

The evidence from recent studies suggests that **federated learning (FL)** can support environmentally sustainable AI—but only under certain technical and organizational conditions. Federated learning changes the way AI models are trained by keeping data where it is created and sending only small model updates to a central server. This structure can help lower energy costs tied to centralized data storage and reduce large data transfers across global networks. However, the actual environmental benefit depends on how efficiently the system is designed and managed. If implemented poorly, Federated learning (FL) can consume more energy than traditional centralized systems because of repeated communication, inefficient local devices, and longer convergence times.From the analysis in previous sections, we can see a pattern: Federated learning (FL) enables sustainable AI when data is large, networks are efficient, and updates are optimized. The *International Energy Agency (IEA, 2024)* estimates that data centers already consume around 460 terawatt-hours (TWh) annually, and this could double by 2030 due to AI demand[28].. When data volumes are in the petabyte range, moving all information to a single data center is extremely energy-intensive. In such scenarios, using Federated learning (FL)—where only small updates are transmitted—can significantly cut total data flow and central storage energy. The *University of Cambridge (2021)* notes that federated learning can lower emissions by keeping data local and sharing only model updates, which can reduce transmission energy if updates are well compressed[29]

However, as shown by *Qiu et al. (2023)* in *The Journal of Machine Learning Research*, the opposite can happen when communication rounds are too many or devices are inefficient. In such cases, Federated learning (FL) can produce up to 100 times more carbon emissions than centralized learning because of redundant updates and high local computation costs[30].[31]These results highlight that design choices decide sustainability. For example, reducing the number of communication rounds, using update compression (such as sparsification or quantization), and training during low-carbon energy hours can all improve efficiency. On the other hand, if model updates are large or sent frequently from thousands of devices, energy savings from avoiding central storage may be cancelled out by the extra computation and network usage.

In simpler terms, federated learning reduces data movement but increases local processing. Centralized learning has high data-transfer and storage costs but efficient computing hardware and power management. Federated learning (FL) distributes the workload to many devices that may not be energy efficient. Therefore, the environmental balance between the two depends on which side of the trade-off dominates. A well-designed Federated learning (FL) system can outperform centralized setups when the data is massive, the communication is optimized, and the hardware is energy efficient. But when local

[26] University of Cambridge. (2021, May 10). Can federated learning save the world? https://www.cam.ac.uk/research/news/can-federated-learning-save-the-world)

[27] Thakur, D., Guzzo, A., Fortino, G., & Piccialli, F. (2024, September 19). *Green federated learning: A new era of green-aware AI*. *arXiv preprint arXiv:2409.12626*. https://arxiv.org/abs/2409.12626.

[28] International Energy Agency. (2024, March 7). *AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works*. https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works

[29] (University of Cambridge. (2021, May 10). *Can federated learning save the world?* https://www.cam.ac.uk/research/news/can-federated-learning-save-the-world.

[30] Bian, J., Wang, L., Ren, S., & Xu, J. (2023). *CAFE: Carbon-Aware Federated Learning in geographically distributed data centres*. *arXiv*. https://arxiv.org/abs/2311.03615

[31]).(Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf)

devices are slow, unoptimized, or located in carbon-heavy regions, Federated learning (FL) can have a higher carbon footprint[32].

For industries and governments, this trade-off has real policy implications. As AI becomes a core part of digital economies, regulators and organizations must set energy and carbon transparency standards for AI training. Companies deploying Federated learning (FL) at scale—such as in healthcare, finance, or smart cities—should perform carbon audits of their systems, similar to financial audits. This includes measuring transmission volume, local computation time, and the carbon intensity of electricity used. Governments can encourage this by providing green-compute credits or low-carbon data zones, similar to renewable energy certificates. Large technology firms can also design carbon-aware scheduling systems that shift training to low-emission times or regions, as proposed by *Bian et al. (2023)* in their CAFE (Carbon-Aware Federated Learning) framework[33].

The results also suggest that policies supporting renewable-powered data centers and energy-efficient edge devices can strengthen Federated learning (FL)'s sustainability. For example, if 5G or IoT networks used renewable energy and supported efficient data transfer, the benefits of federated learning would become larger. On the technical side, investment in hardware optimization (low-power chips, co-processors for Federated learning (FL)) and software innovations (smarter aggregation algorithms, adaptive compression) are critical to realizing Federated learning (FL)'s potential.

At the same time, it is important to note the limitations of this paper. Much of the evidence on Federated learning (FL)'s carbon impact comes from simulations or lab-scale studies, not real-world systems. There are still data gaps in measuring energy use at the device level, particularly for mobile phones, embedded IoT systems, or rural networks. Emission factors also vary across regions depending on the local electricity mix, which limits generalizability. Furthermore, different models have very different energy behaviors—training a large natural language model is not comparable to a small medical model—so results cannot be uniformly applied across all domains. The current literature also lacks standard methods to measure and report the carbon footprint of federated systems, which makes it difficult to compare studies objectively.

Future research should therefore focus on three main areas. First, there is a need for real-world longitudinal studies that measure Federated learning (FL) energy consumption over time, across devices and regions, rather than relying solely on simulations. Second, researchers should develop standardized metrics and tools for calculating the carbon footprint of Federated learning (FL), similar to the "Green Algorithms" model (Lannelongue et al., 2021) that estimates emissions from computation using power, runtime, and grid intensity[34] Third, there is potential for combining Federated learning (FL) with renewable-energy integration—for example, scheduling Federated learning (FL) rounds when solar or wind energy availability is high. This can be supported by cloud platforms offering APIs that provide real-time carbon intensity data.

[32] (International Energy Agency. (2023). *Data centres and data transmission networks*. https://www.iea.org/reports/data-centres-and-data-transmission-networks).

[33] (Bian, J., Wang, L., Ren, S., & Xu, J. (2023). *CAFE: Carbon-Aware Federated Learning in geographically distributed data centres*. *arXiv*. https://arxiv.org/abs/2311.03615)

[34] (Lannelongue, L., Grealey, J., & Inouye, M. (2021). *Green Algorithms: Quantifying the carbon footprint of computation*. *Methods in Ecology and Evolution, 12*(1), 108–118. https://besjournals.onlinelibrary.wiley.com/doi/10.1111/2041-210X.13534).

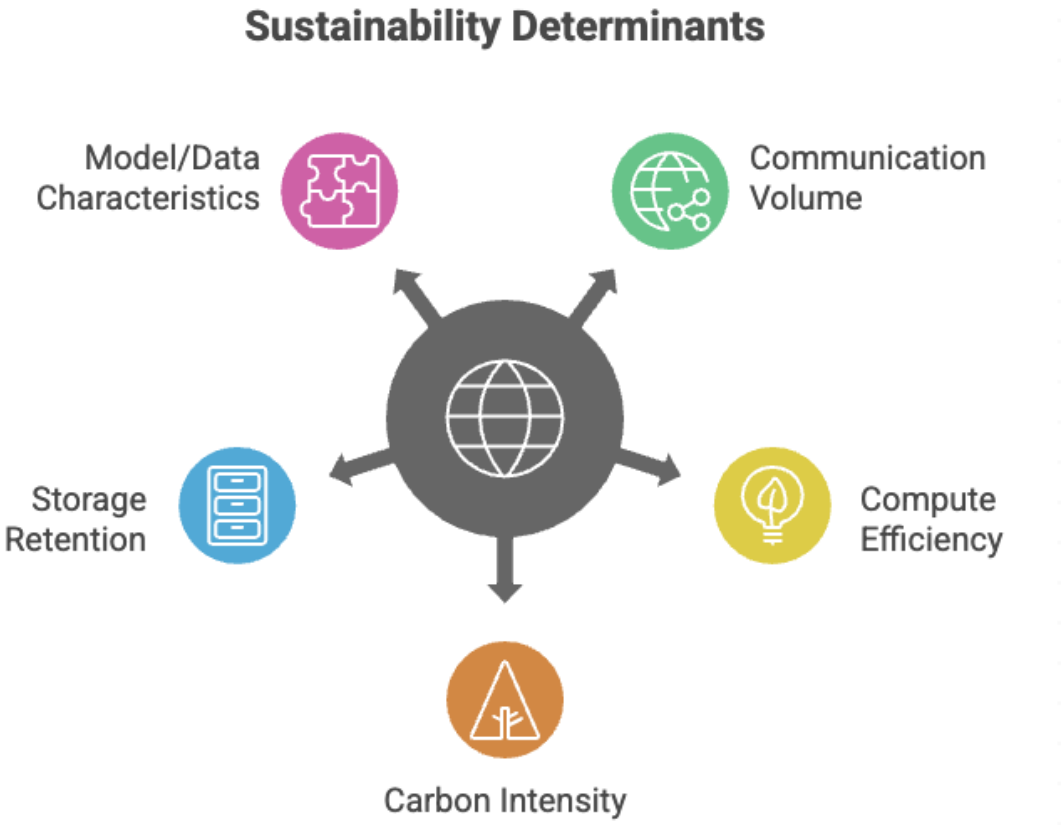

CONCLUSION

This paper asked a simple question: **Can federated learning (FL) enable environmentally sustainable AI by reducing centralized data storage and transmission costs?** The evidence suggests a **conditional yes**. FL can lower the need to move and store large volumes of raw data in data centres by keeping data local and sending only model updates. When updates are well-compressed, communication rounds are limited, and training is scheduled in low-carbon places and hours, total energy use and emissions can fall compared with a fully centralized workflow[35]. At the same time, today's AI expansion is putting strong pressure on electricity systems—data-centre demand is projected to more than double to ~945 TWh by 2030—so any method that reduces network traffic and storage load can have real impact if engineered well[36].

However, Federated learning (FL) is not a guaranteed win. It reduces raw-data movement but adds local computation and repeated communications. Poor design—large uncompressed updates, too many rounds, inefficient devices, or high-carbon grids—can erase savings and even make Federated learning (FL) worse than centralized training. The most comprehensive study to date finds Federated learning (FL) can emit up to 100× more carbon in unfavourable settings, though results can be comparable or better when

[35] (University of Cambridge. (2021, May 10). *Can federated learning save the world?* https://www.cam.ac.uk/research/news/can-federated-learning-save-the-world)
[36] International Energy Agency. (2024). *Energy and AI: Energy demand from AI*. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

configurations are improved. In short: architecture and implementation decide the outcome[37].Practically, this means organisations should treat Federated learning (FL) as a sustainability lever rather than a panacea. Before deployment, teams should measure expected GB transferred, TB-months stored, device/server kWh, and convert these into CO_2e using location-specific grid factors. Tools and methods such as Green Algorithms help standardise this accounting. During operation, firms should apply carbon-aware orchestration—selecting clients in cleaner regions and running rounds at low-carbon times—to keep emissions within a budget while maintaining model quality[38]. Policymakers can support this by encouraging energy and carbon transparency for AI, creating incentives for low-carbon computing zones, and promoting efficient networks and renewable-powered data centres, given their growing role in electricity demand[39]. FL can enable more sustainable AI when designed for minimal communication, efficient local training, and carbon-aware scheduling; it cannot guarantee sustainability without these controls. The key is to optimise the full pipeline—communication, storage, and compute—rather than assuming decentralisation alone is greener.Researchers should publish standardised carbon metrics with Federated learning (FL) results and run real-world, longitudinal measurements across devices and grids. Industry should build **default compression and client-selection** into Federated learning (FL) stacks, expose **carbon-intensity APIs**, and report **audited CO_2e** alongside accuracy. With these steps, federated learning can move from promising idea to proven practice for lowering the environmental footprint of AI[40].

REFERENCES:

1. Bian, J., Wang, L., Ren, S., & Xu, J. (2023). *CAFE: Carbon-aware federated learning in geographically distributed data centres*. *arXiv preprint arXiv:2311.03615*. https://arxiv.org/abs/2311.03615
2. Bian, J., Wang, L., Ren, S., & Xu, J. (2024, February 6). *CAFE: Carbon-aware federated learning in geographically distributed data centers* (Version 2). *arXiv preprint arXiv:2311.03615*. https://arxiv.org/abs/2311.03615
3. Chen, S., & Nature Magazine. (2025, April 10). *Data centers will use twice as much energy by 2030—Driven by AI*. *Scientific American*. https://www.scientificamerican.com/article/ai-will-drive-doubling-of-data-center-energy-demand-by-2030
4. International Energy Agency. (2023). *Data centres and data transmission networks*. https://www.iea.org/reports/data-centres-and-data-transmission-networks
5. International Energy Agency. (2024). *Artificial intelligence: Topics page*. https://www.iea.org/topics/artificial-intelligence
6. International Energy Agency. (2024). *Energy and AI: Executive summary – The transformative potential of AI depends on energy*. *IEA*. https://www.iea.org/reports/energy-and-ai/executive-summary
7. International Energy Agency. (2024). *Energy and AI: Energy demand from AI*. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

[37] Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf

[38] Lannelongue, L., Grealey, J., & Inouye, M. (2021). *Green Algorithms: Quantifying the carbon footprint of computation*. *Methods in Ecology and Evolution, 12*(1), 108–118. https://arxiv.org/abs/2007.07610

[39] International Energy Agency. (2024). *Artificial intelligence: Topics page*. https://www.iea.org/topics/artificial-intelligence

[40] (Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf)

8. International Energy Agency. (2024, March 7). *AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works*. *IEA*. https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works
9. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). *Advances and open problems in federated learning*. *Foundations and Trends® in Machine Learning*, *14*(1–2), 1–210. https://arxiv.org/abs/1912.04977
10. Lannelongue, L., Grealey, J., & Inouye, M. (2020, December 17). *Green algorithms: Quantifying the carbon footprint of computation* (Version 5). *arXiv preprint arXiv:2007.07610*. https://arxiv.org/abs/2007.07610
11. Lannelongue, L., Grealey, J., & Inouye, M. (2021). *Green algorithms: Quantifying the carbon footprint of computation*. *Methods in Ecology and Evolution*, *12*(1), 108–118. https://besjournals.onlinelibrary.wiley.com/doi/10.1111/2041-210X.13534
12. Mytton, D. (2025, February 9). *Data centre energy consumption data-centre energy and AI in 2025*. *dev/sustainability*. https://www.devsustainability.com/p/data-center-energy-and-ai-in-2025
13. Qiu, X., Parcollet, T., Fernandez-Marques, J., Gusmão, P. P. B., Gao, Y., Beutel, D. J., Topal, T., Mathur, A., & Lane, N. D. (2023). *A first look into the carbon footprint of federated learning*. *Journal of Machine Learning Research, 24*, 1–23. https://www.jmlr.org/papers/volume24/21-0445/21-0445.pdf
14. Rice, S. (2025, April 14). *AI, energy and the data centre surge: IEA's global warning*. *Data Centre Magazine*. https://datacentremagazine.com/technology-and-ai/iea-asks-how-is-ai-reshaping-global-energy-demand
15. Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024). *2024 United States data center energy usage report*. *Lawrence Berkeley National Laboratory*. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2001637.pdf
16. Thakur, D., Guzzo, A., Fortino, G., & Piccialli, F. (2024, September 19). *Green federated learning: A new era of green-aware AI*. *arXiv preprint arXiv:2409.12626*. https://arxiv.org/abs/2409.12626
17. University of Cambridge. (2021, May 10). *Can federated learning save the world?* https://www.cam.ac.uk/research/news/can-federated-learning-save-the-world
18. University of Cambridge, Machine Learning Systems Lab. (n.d.). *Carbon dioxide in federated learning (interactive tool)*. https://mlsys.cst.cam.ac.uk/carbon_fl/
19. U.S. Department of Energy. (2024, December 20). *DOE releases new report evaluating increase in electricity demand from data centers*. *Energy.gov*. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers
20. World Economic Forum. (2023, September 20). *How AI can help build a sustainable future*. *WEF*. https://www.weforum.org/stories/2025/07/ai-infrastructure-sustainable-future-energy/