Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Federated learning reconciles data sovereignty with collaborative energy-transition modelling across the Global South: a public-data benchmark anchored in Indonesia

Sumartono Darmanto

ericsdarmanto@gmail.com

Tsinghua University

Article

Keywords:

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10722871/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. The author declares the following interests. The author is a former senior advisor for strategic economic affairs at the Executive Office of the President, Republic of Indonesia, and has advised other Government of Indonesia priority programmes. The author is affiliated with a collaborative research programme that intends to pursue federated AI infrastructure for low-carbon development in the Global South as a policy initiative and potentially funded venture, with Indonesia envisaged as a first pilot. Accordingly, all statements about potential infrastructure in the Discussion are framed as requiring independent economic, technical, governance and environmental evaluation, and no such system is represented as validated by the present study, whose empirical claims are confined to the benchmark reported in Results.

Federated learning reconciles data sovereignty with collaborative energy-transition modelling across the Global South: a public-data benchmark anchored in Indonesia

Sumartono Darmanto (corresponding author)
Department of Earth System Science, Tsinghua University, Beijing, China

Email: ericsdarmanto@gmail.com

Abstract

Meeting Paris-aligned trajectories in the Global South requires rapid, data-intensive coordination of where and how low-carbon energy infrastructure is built, yet the underlying data are fragmented across sovereign jurisdictions reluctant to pool sensitive energy, land and infrastructure information. Wealthy blocs are answering their own version of this problem with shared or sovereign compute at extraordinary cost — a €10-billion European network of publicly funded 'AI Factories' and a US$500-billion private buildout in the United States — scales no developing country can replicate alone. Federated learning (FL), in which institutions jointly train a shared model by exchanging model parameters rather than raw data, offers an alternative technical foundation, but its value for cross-country energy modelling in developing regions has not been quantified. Treating 44 Global South countries (16,404 operating facilities) in the Global Energy Monitor Global Integrated Power Tracker as sovereign data silos, we benchmarked four training regimes on a geospatial renewable-versus-fossil classification task. Federated averaging matched centralized pooling of all raw data to three decimal places (AUC 0.895 versus 0.895) while no record left its country of origin, and personalized FL exceeded isolated national training (0.910 versus 0.905; two-sided Wilcoxon $p = 0.048$), with the advantage larger on average among data-poor silos and widening as training data were artificially restricted. Independent replication on the 2021 WRI Global Power Plant Database (31 silos, 6,545 plants) reproduces the sovereignty result, though there the personalization advantage is statistically indistinguishable from zero. A high-capacity gradient-boosted model trained on pooled data outperforms every federated configuration tested (0.960 versus 0.815 trained locally), a gap driven by training-set volume rather than implicit country identification: the median silo holds 63 training facilities against 11,479 pooled. Aggregation therefore has large and rising value with model capacity, no country's own data suffice for the most capable methods, and extending federation to such models is the decisive next step. These results provide a quantitative foundation for federated architectures — national AI Labs coupled to shared, renewably powered training infrastructure under parity governance — as coordination infrastructure for Global South energy transitions, with Indonesia as the natural first national pilot.

Introduction

The climate arithmetic of this decade is demanding. The Intergovernmental Panel on Climate Change assesses that limiting warming to 1.5°C requires global greenhouse-gas emissions to fall by around 43% by 2030 relative to 2019, and that annual mitigation investment must grow by a factor of three to six this decade[1]; the International Energy Agency's net-zero pathway requires world renewable capacity to triple to roughly 11,000 GW by 2030[2]. Low-carbon technology and capital exist at the required scale — newly commissioned onshore wind and utility-scale solar are now typically cheaper than fossil alternatives[3] — while the adaptation burden falls disproportionately on developing countries, whose needs are projected at US$310–365 billion per year by 2035 against international public adaptation finance of roughly US$26 billion in 2023[4]. What is missing, at multiple scales simultaneously, is trusted infrastructure for connecting resources to the specific places that need them, at the pace the arithmetic requires.

Indonesia — the focal case of this study — has committed to that pace. Its Electricity Supply Business Plan (RUPTL) 2025–2034, ratified in May 2025, plans 69.5 GW of new capacity of which 42.6 GW — some 61% — is new and renewable energy, in support of the national commitment to net-zero emissions by 2060 or earlier[5,6]. Delivery runs through two coordination bottlenecks that motivate this study. The first is financial. The Just Energy Transition Partnership (JETP) paired an initial US$20-billion pledge with an investment plan requiring at least US$97.1 billion of power-sector investment to 2030[7]; after the United States' withdrawal in early 2025 remaining partners raised the pledge to US$21.4 billion, yet only about US$3.1 billion — 14.5% — had been approved by late 2025. Alongside it, the Asian Development Bank's Energy Transition Mechanism had since 2022 pursued the early retirement of the 660 MW Cirebon-1 coal plant as its flagship pilot; towards the end of 2025, three years after the memorandum was signed, the government announced that this retirement plan had been cancelled[8]. Independent analyses attribute this to weak institutional coordination and unresolved policy ambiguity operating alongside a loan-heavy financing structure[8,9]. The second bottleneck is informational, and spans the statistical system rather than any single institution. Official programme reporting in January 2024 recorded a village-electrification ratio of 99.83%, with 140 villages remaining under a definition that counts a village as electrified whatever the source of supply[10]. Rural-electrification monitoring convened by the Directorate General of Electricity later the same year reported that 6,072 villages and 1.3 million households did not yet have access to electricity from the state utility[11]. Both figures are internally valid; they measure different things for different purposes, and the forty-fold spread between them is the signature of a data landscape in which electrification statistics serve programme reporting, utility planning and village-development policy separately, without a shared reconciliation layer. The government has committed to closing the physical gap: the RUPTL's village-electrification programme targets more than 5,700 remote villages and nearly 780,000 households[6]. Delivering that efficiently presupposes exactly what is missing — a single, version-controlled, definitionally reconciled village-level dataset — which is a data-infrastructure task.

Artificial intelligence is the emerging tool for organizing fragmented, multi-source information at this scale, and wealthy blocs are building the infrastructure for it — pointedly, for themselves. The European Union's EuroHPC Joint Undertaking has selected nineteen sites for a network named 'AI Factories', backed by public investment the European Commission puts at roughly €10 billion, with its InvestAI initiative aiming to mobilize €200 billion including a dedicated €20-billion facility for up to five larger 'AI gigafactories'[12,13]. Even the entry ticket is substantial: LUMI, the EuroHPC pre-exascale machine in Finland, cost €144.5 million and required a consortium of ten countries[14]; Denmark's national AI supercomputer Gefion required approximately DKK 700 million (≈US$100 million)[15]. In the United States, the privately led Stargate project intends to invest US$500

billion over four years[16]. Regional blocs are also building the software layer of cooperation: following the green-industry cooperation statement adopted at the SCO's 2025 Tianjin Summit, China established the China–SCO Green Industry Cooperation Platform, whose work includes a shared database of industrial investment projects among members[17]. The pattern is consistent: those who can afford it are building shared or sovereign AI infrastructure inside their own blocs. Most Global South countries belong to none of these arrangements, cannot individually afford LUMI-class — let alone Stargate-class — infrastructure, and face a constraint the wealthy blocs internalize through treaty and market integration: the planning data such systems need are held by individual governments and treated as sensitive for reasons of national security, commercial confidentiality and sovereignty. Each country, and often each ministry, therefore plans in relative isolation.

Federated learning offers a principled way through. Introduced as a method for training a shared model across decentralized clients that communicate only model updates, never raw data[18], FL has been studied extensively as a privacy-preserving paradigm[19] and is the natural candidate architecture for climate cooperation that respects data sovereignty. Its value for cross-country energy modelling in developing regions has, to our knowledge, not been quantified on real sovereign-partitioned data. Here we do so, asking four questions: whether federation sacrifices accuracy relative to pooling raw data; how both behave under the cross-country heterogeneity characteristic of the Global South; whether benefits are distributed equitably between data-rich and data-poor participants; and whether federation's benefit grows as local data become scarce — the operating condition of many Global South agencies. We answer them on the Global Energy Monitor Global Integrated Power Tracker[20], and replicate the whole benchmark independently on the WRI Global Power Plant Database[21], so that the conclusions do not rest on the particularities of either. We then set out what the findings imply for a three-component coordination architecture, and the governance, verification, energy and water requirements that remain before such infrastructure could responsibly be deployed.

Results

A heterogeneous, fragmented generation landscape

We retained all operating facilities in a pre-specified list of 55 Global South countries whose primary fuel is renewable (utility-scale solar, wind, hydropower, bioenergy, geothermal) or fossil (coal, oil/gas), keeping the 44 countries with at least 20 facilities and at least five in each class — 16,404 facilities in total (Fig. 1a; Table 1; Methods). Indonesia sits at the fossil-intensive end, with 299 catalogued facilities of which 48% are renewable; coal (99) and oil or gas (57) dominate installed capacity, with approximately 85.8 GW of fossil capacity against 12.2 GW of renewables, even though utility-scale solar (76) is now the second most numerous facility type (Fig. 1b). Across silos the landscape is highly heterogeneous: the renewable share of facilities ranges from 0.16 (IRQ) to 0.98 (BRA), with a mean of 0.73 (Fig. 2a; Table 1). This is precisely the non-independent-and-identically-distributed (non-IID) condition the federated-learning literature identifies as its central technical challenge[22].

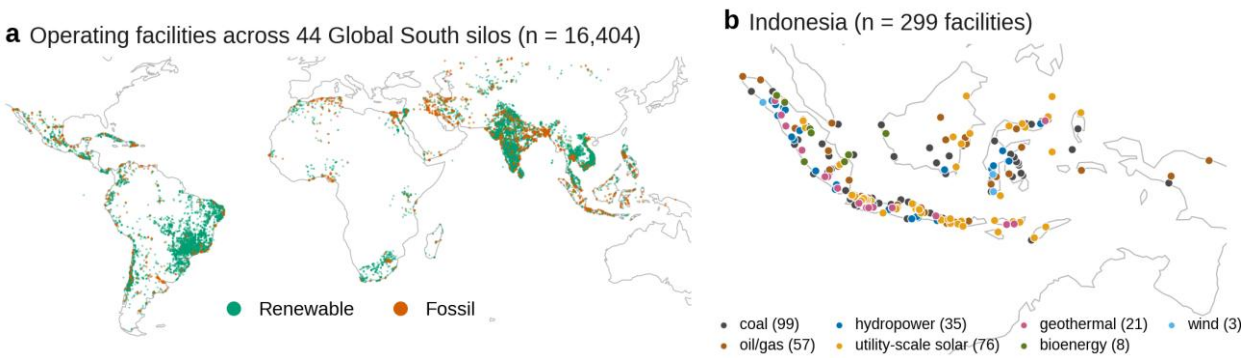


Fig. 1 | Spatial distribution of catalogued operating generation. a, All facilities in the retained Global South country silos, coloured by renewable versus fossil primary fuel. b, Indonesia, coloured by dominant facility type. Both panels are drawn at true geographic aspect ratio. Data: Global Energy Monitor Global Integrated Power Tracker, August 2026 release; coastlines: Natural Earth (public domain).

Country	Facilities	Renew. share	Country	Facilities	Renew. share
Algeria	94	0.41	Kenya	49	0.86
Argentina	204	0.72	Madagascar	22	0.77
Bangladesh	141	0.26	Malaysia	228	0.86
Bolivia	31	0.74	Mexico	368	0.69
Brazil	3876	0.98	Mongolia	23	0.74
Cambodia	39	0.82	Morocco	61	0.85
Chile	857	0.95	Myanmar	75	0.67
Colombia	141	0.86	Nigeria	64	0.55
Cuba	99	0.92	Pakistan	270	0.81
Dominican Rep.	63	0.67	Panama	77	0.94
Ecuador	68	0.79	Peru	81	0.80
Egypt	150	0.60	Philippines	209	0.78
El Salvador	87	0.94	Senegal	28	0.64
Ghana	33	0.64	South Africa	326	0.93
Guatemala	49	0.90	Sri Lanka	118	0.92
Honduras	69	0.86	Tanzania	22	0.68
India	5019	0.93	Thailand	803	0.91
Indonesia	299	0.48	Tunisia	27	0.41
Iran	289	0.51	Uzbekistan	75	0.75
Iraq	96	0.16	Venezuela	23	0.39
Jordan	213	0.94	Viet Nam	1394	0.97
Kazakhstan	109	0.65	Yemen	35	0.46

Table 1 | The 44 country silos of the benchmark: retained operating facilities and renewable share per silo (total n = 16,404). Each country's data remain within its own silo in all federated regimes.

The task, the models, and the sovereignty question

We framed a supervised task representative of energy-planning inference: predicting whether a facility is renewable or fossil from geospatial and physical attributes (location, capacity, distance to nearest populated place and coastline, start year; Methods). The classifier is a transparent L2-regularized logistic regression: the estimated probability that a facility with feature vector x is renewable is

$$p(\text{renewable} \mid x) = 1 / (1 + \exp[-(w \cdot x + b)]) \quad (1)$$

Here x is the facility's eight-dimensional standardized feature vector; w is the vector of eight learned weights, one per feature, encoding how strongly and in which direction each feature pushes the prediction toward 'renewable'; b is a single learned intercept; and w·x is their dot product. Training finds the (w, b) minimizing regularized cross-entropy loss on the training facilities by gradient descent (Methods). Four regimes were compared. **Local-only:** each country trains its own (w, b) on its own facilities. **Centralized:** one model trained on all silos' pooled raw data — the baseline that maximizes information but violates every participant's data sovereignty. **Federated** (FedAvg[18]): over 40 communication rounds each silo k trains briefly from the current shared weights and a coordination server aggregates only the resulting parameters, weighting each country by its number of training facilities n_k:

$$w_shared = \Sigma_k (n_k / n_total) \cdot w_k , \quad b_shared = \Sigma_k (n_k / n_total) \cdot b_k \quad (2)$$

where w_k and b_k are country k's locally trained parameters and $n_total = \Sigma_k n_k$. Equation (2) is the entire content of what crosses a border: eight weights and one intercept per country per round, never a facility record. **Personalized FL:** from the converged shared model each silo fine-tunes locally on its own data[23], ending with a country-specific model that began from the collective one.

How every number in Table 2 is produced

Two metrics are reported. **Accuracy** is the share of held-out test facilities classified correctly when every facility with p ≥ 0.5 is called renewable. **AUC** measures ranking quality independent of that threshold: over every pair of one renewable and one fossil test facility, the fraction in which the model assigns the renewable one the higher probability:

$$AUC = (\text{renewable–fossil pairs ranked correctly}) / (\text{total such pairs}) \quad (3)$$

A worked miniature makes both concrete. Suppose a country's test set holds five facilities and the model assigns probabilities 0.9, 0.7, 0.4 to its three renewable facilities and 0.6, 0.2 to its two fossil ones. There are 3 × 2 = 6 renewable–fossil pairs; the renewable outranks the fossil in five of them (only the 0.4-renewable versus 0.6-fossil pair is ordered wrongly), so AUC = 5/6 ≈ 0.83. Accuracy at the 0.5 threshold differs: the 0.4-renewable is missed and the 0.6-fossil is a false alarm, so 3 of 5 are correct and accuracy = 0.60. A model can therefore rank well while being miscalibrated at a fixed threshold — which is why, in Table 2, the centralized and federated models show high AUC but markedly lower accuracy: one global threshold cannot fit forty-four different renewable–fossil balances at once. Every entry in Table 2 is then produced in three steps: compute the metric on each country's own held-out facilities; average the country values so every country counts equally regardless of size; and repeat the whole procedure over ten random 70/30 train/test splits, reporting mean ± standard deviation across the ten repeats.

Training regime	Test AUC	Test accuracy	Raw data shared?
Local-only (per country)	0.905 ± 0.013	0.879 ± 0.007	No
Centralized (pooled)	0.895 ± 0.013	0.813 ± 0.008	Yes — all silos
Federated (FedAvg)	0.895 ± 0.013	0.806 ± 0.006	No
Personalized FL	0.910 ± 0.012	0.877 ± 0.008	No

Table 2 | Performance of four training regimes on renewable-versus-fossil classification across the Global South silos (macro-averaged; mean ± s.d. over 10 random splits). See Eq. (3) and the worked example above for how AUC and accuracy are computed and why they can diverge.

Parameter sharing forfeits nothing relative to pooling raw data

The first result is an equality rather than a difference. Federated averaging matched centralized pooling to three decimal places (AUC 0.895 versus 0.895) despite never pooling a single record: whatever a shared global model can learn from pooled raw data, it learns as well from the parameter exchange of Eq. (2) alone. Both single-model regimes trail countries' own local models on AUC by a small margin (0.905) and on accuracy by a wide one (0.879 against 0.813), the latter reflecting threshold miscalibration under heterogeneous class balance rather than weaker discrimination. The practical implication is direct: a government asked to surrender raw data to a central pool would be buying nothing that parameter sharing does not already provide.

Personalization exceeds isolated national training

Fine-tuning the shared model within each silo did more than recover local performance: personalized FL exceeded isolated local training (AUC 0.910 versus 0.905), and across the 44 silos the per-country difference is statistically significant (two-sided Wilcoxon signed-rank test on seed-averaged per-silo AUC differences, $p = 0.048$). A country joining the federation therefore ends with a model better than one trained in isolation, plus membership of a shared model and common technical infrastructure, at no sovereignty cost whatsoever (Fig. 2b). This is the paper's central quantitative result: on current data, sovereignty-preserving collaboration is not a compromise but an improvement.

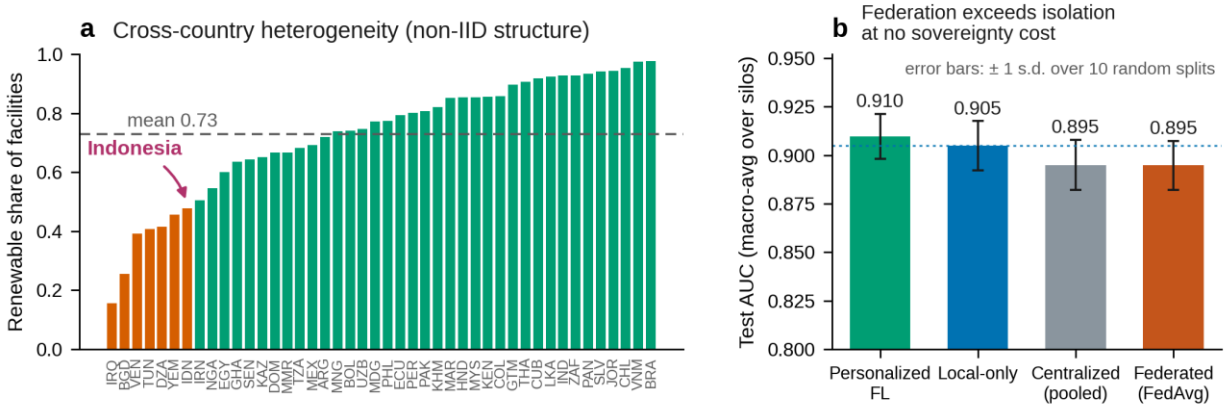


Fig. 2 | Cross-country heterogeneity and model performance. a, Renewable share of facilities by country silo, fossil-majority silos in orange; dashed line marks the mean; Indonesia highlighted. b, Test AUC of the four regimes (mean ± s.d., 10 splits); dotted line marks local-only AUC. Federation matches pooling exactly; personalization exceeds isolated training.

The dividend concentrates where data are scarce — and no one is left worse off

Federation's benefit is not evenly distributed, and its distribution is the point. Descriptively, silos below the median size gained more on average than those above it (0.0069 against 0.0030 AUC), though per-silo gains are widely spread and the association with silo size is weak (Fig. 3a); the controlled scarcity experiment below provides the direct test. Crucially, the larger data-holders are not subsidising the smaller: all five of the largest silos show positive mean gains, so federation is weakly positive for data-rich and data-poor participants alike, with the larger benefit accruing to those with least data. To test the mechanism directly we restricted every silo's training set to at most 40, then at most 16, class-balanced records — the scale of a thin national asset register — while keeping test sets unchanged. The advantage of personalization over isolation widens as data thin, from +0.006 at full data to +0.009 at 16 records per silo (Fig. 3b). The scarcer a country's own data, the more it gains from a federation in which everyone's parameters — and no one's records — are shared.

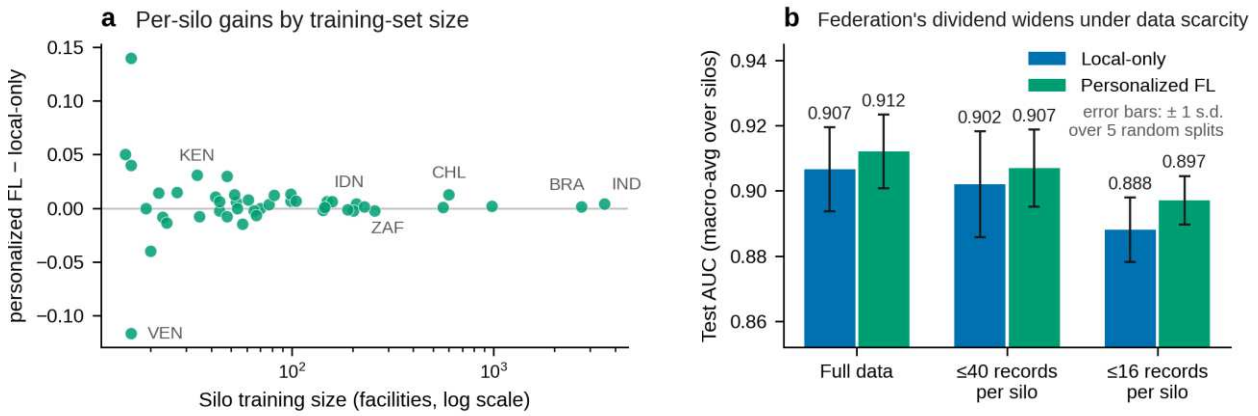


Fig. 3 | Distribution of federation benefits. a, Change in test AUC from personalized FL relative to isolated local training against silo training size (log scale; each point one country, averaged over 10 splits). b, Macro-averaged AUC of local-only versus personalized FL as per-silo training data are restricted to at most 40 and 16 records (mean ± s.d., 5 splits).

Robustness: the pattern holds — with one instructive exception

We tested the main result against four alternatives (Methods). First, replacing the logistic classifier with a single-hidden-layer neural network — a non-linear model to which federated averaging applies unchanged — reproduces the qualitative pattern — personalized FL is again strongest (0.933) and exceeds local training (0.931), while both single global models trail — though federated averaging here concedes a small margin to centralized training (0.910 versus 0.919; Table 3). Second, varying the silo-inclusion threshold leaves the ordering unchanged at every setting. Third, adding the clipped Gaussian noise required for formal privacy protection leaves personalized FL essentially untouched across the whole range tested (AUC 0.912 with no noise, 0.910 at the heaviest level), indicating that in this setting privacy engineering is close to free; we report this as a noise-sensitivity analysis rather than a certified guarantee, having not computed a formal privacy budget. Fourth, because commissioning-year metadata could in principle let a model learn from documentation habits rather than physical geography, we re-ran the comparison without the record-completeness indicator and again without any start-year information: neither ablation disturbs the pattern (Table 3), and in this dataset the issue is minor in any case, with only 5.9% of facilities lacking a start year.

The instructive exception concerns model capacity. A gradient-boosted tree ensemble — a high-capacity model to which federated averaging does not directly apply — reverses the comparison: trained locally it performs poorly (0.815), and trained on pooled data it outperforms every federated configuration tested (0.960). A diagnostic rules out the obvious explanation that the pooled model implicitly learns country identity from coordinates: removing latitude and longitude leaves the pooled result nearly unchanged (0.953). The driver is training-set volume. The median silo contributes only 63 training facilities and the smallest 15, against 11,479 pooled — far too few for a high-capacity learner locally, ample in aggregate. Read correctly this sharpens rather than undermines the paper's argument: the value of aggregation rises steeply with model capacity, and isolated national modelling fails worst exactly where the most capable methods are involved. What it does not show is that centralization is required to capture that value, because we did not test a federated high-capacity model. Federated gradient-boosted trees are an established and actively developed method[24,25]; benchmarking them on this task is, on our evidence, the single most consequential extension of this work, and until it is done the accuracy claims here should be read as applying to the model classes tested rather than to all model capacities.

Robustness variant	Local-only	Centralized	Federated	Personalized FL
Logistic regression (main result, Table 2)	0.905	0.895	0.895	0.910

Robustness variant	Local-only	Centralized	Federated	Personalized FL
Neural network (1 hidden layer)	0.931	0.919	0.910	0.933
Gradient-boosted trees	0.815	0.960	not tested	not tested
Threshold ≥ 40 facilities (34 silos)	0.909	0.902	0.902	0.912
Personalized FL with privacy noise (σ = 0.5)	—	—	0.882	0.910
Without record-completeness flag (7 features)	0.909	0.903	0.902	0.914
Without any start-year data (6 features)	0.909	0.905	0.905	0.911

Table 3 | Robustness of the four-regime comparison (test AUC, macro-averaged across silos). Robustness rows use 5 random splits (3 for the threshold row); the main result uses 10. Federated averaging applies directly to the logistic and neural models; gradient-boosted trees require a different federated construction, which we did not test.

Independent-dataset replication

Because dataset vintage and coverage could in principle drive these results, we repeated the entire benchmark on an independent source: the WRI Global Power Plant Database v1.3.0[21], which catalogues the fleet as of its 2021 release. Applying the same inclusion rule yields a smaller and older sample — 31 silos and 6,545 plants — with a mean renewable share of 0.58 against 0.73 here, reflecting the solar and wind capacity built since 2021 (Table 4; Methods). The core findings replicate. Federated averaging again tracks centralized pooling closely (0.711 versus 0.724), personalized FL again recovers the level of isolated local training within noise (0.838 versus 0.842), the scarcity dividend again widens as data thin, and the threshold miscalibration of single global models is again pronounced. Two quantities differ. On the older catalogue the personalization advantage is not statistically distinguishable from zero (Wilcoxon p = 0.15) where on current data it is, so the strength of that claim rests on the more recent evidence. Conversely, the older catalogue shows centralized pooling trailing local models by a wide margin (0.118 AUC) that is not reproduced here (0.010). We therefore state the claim at the strength both datasets support: pooling raw data never buys an advantage that federation does not also deliver, and costs sovereignty to obtain what parameter sharing obtains for free.

Quantity (test AUC unless noted)	GEM GIPT (Aug 2026, primary)	WRI GPPD v1.3.0 (2021)
Silos / facilities	44 / 16,404	31 / 6,545
Mean renewable share	0.73	0.58
Local-only	0.905	0.842
Centralized (pooled)	0.895	0.724
Federated (FedAvg)	0.895	0.711
Personalized FL	0.910	0.838
Personalized FL vs local (Wilcoxon p)	0.048	0.15
Scarcity gain at ≤16 records/silo	+0.009	+0.006

Table 4 | The same benchmark on two independent datasets, using the identical inclusion rule, features, split protocol and regimes. Federation matches pooling on both; personalization exceeds isolated training on the current dataset and matches it within noise on the older one.

Discussion

Our central finding inverts the usual framing of the data-sharing debate. Global South governments are routinely asked to trade sovereignty for the analytical benefits of pooling; on this evidence there is nothing to trade. Federation matches centralized pooling to three decimal places, personalization exceeds isolated

national training, and the advantage widens as data grow scarce. These results are a proof of mechanism, not a finished planning system: a deployable system would integrate the resource, terrain, land-cover and grid layers that genuine multi-criteria siting analysis requires[26], within the same federated discipline. The methodological result reported here is the precondition for that work.

Why participation is rational — and for whom

Three incentives follow from the results, and it is worth being precise about which the evidence supports. The first is cost. Independently replicating even the smallest documented national AI system (Gefion, ≈US$100 million[15]) competes directly with electrification and adaptation budgets; the largest (Stargate, US$500 billion[16]) exceeds Indonesia's entire JETP pledge roughly 23-fold and the JETP financing actually approved to date roughly 160-fold[8]. Under the three-component design (Fig. 4) large-scale compute is provisioned only once, at the shared facility: the AI Factory supplies it, and the coordination model's own demands — training, fine-tuning and continuously serving an LLM-scale system — run on that shared hardware rather than on any national budget; a National AI Lab needs only the capacity to fine-tune and query models — the personalization step benchmarked here — which is orders of magnitude cheaper than training from scratch. A second calculation makes the pace explicit. Let I_req be the investment an agreed plan requires, I_appr the amount approved so far, r_obs the observed approval rate and T_rem the years remaining. The acceleration factor needed to deliver the plan on time is

$$A = (I_req - I_appr) / (T_rem \times r_obs) \qquad (4)$$

For Indonesia's JETP, I_req = US$97.1 billion by 2030, I_appr ≈ US$3.1 billion approved over roughly three years (r_obs ≈ US$1.0 billion per year), and T_rem = 5 years from the start of 2026, giving A ≈ (97.1 − 3.1) / (5 × 1.0) ≈ 19: mobilization must accelerate roughly nineteen-fold for the plan's own target to be met. This is illustrative arithmetic on cited figures rather than a forecast; its purpose is to make the scale of the coordination challenge explicit, in the spirit in which project-level climate consequences have been made explicit through transparent calculation elsewhere in this journal[27]. The second is data. Even where capital exists, national datasets are thin: the median silo in this study holds 63 training facilities and the smallest 15. The gradient-boosting result quantifies what that costs — 0.815 trained alone against 0.960 on aggregated data, a gap roughly thirty times the margin federation currently delivers with linear models. For the most capable methods, no participant's own data suffice, and that is true of the largest silos in our sample as well as the smallest. The third concerns who gains. Federation here is not a transfer from data-rich to data-poor participants: below-median silos gain 0.0069 AUC on average and above-median silos 0.0030, with all five of the largest showing positive mean gains. Participation is therefore weakly positive for every class of participant, with the larger benefit flowing to those with least data — which is the distributional property a cooperative arrangement needs if it is to be joined voluntarily and sustained. That such federations do hold together among parties with strong commercial reasons for secrecy is no longer conjecture: in the MELLODDY consortium ten of the world's largest pharmaceutical companies jointly trained models across their proprietary chemical libraries for three years using cryptographically secured aggregation, each improving its own models, with no compound data exposed to any competitor[28]. Two limits should be stated plainly. Our sample is Global South only, so while the mechanism is symmetric — any silo's parameters inform every other — we have not measured what technologically advanced non-participants such as China or the European Union would gain, and no claim to that effect follows from these results. And the accuracy margins involved are small in absolute terms; the case for participation rests on their direction and distribution, on the cost asymmetry, and on the capacity gap the gradient-boosting result exposes, not on a large measured performance jump.

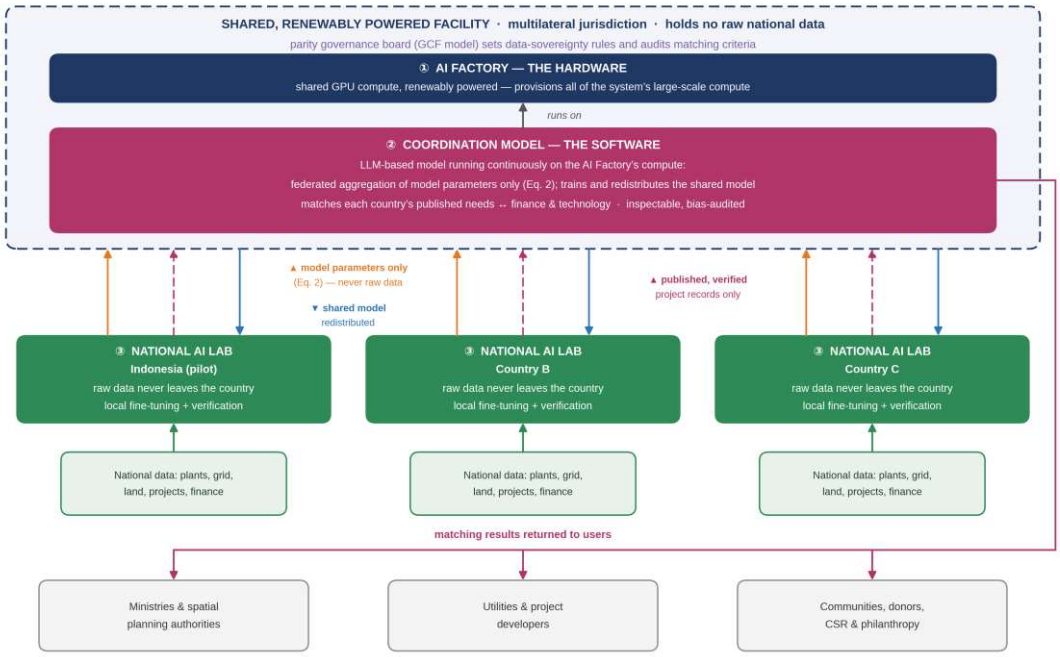


Fig. 4 | Three components, one system — the architecture for sovereignty-preserving AI coordination. Two of the three components constitute the shared, renewably powered facility, which operates under multilateral jurisdiction and holds no raw national data: the AI Factory (component 1) is the facility's hardware — shared GPU compute, provisioning all of the system's large-scale demand, including that of the coordination model — and the coordination model (component 2) is its software, an LLM-based model running continuously on that compute, which performs the federated aggregation of Eq. (2), trains and redistributes the shared model, and matches each country's published needs to finance and technology. The third component is distributed: each National AI Lab (component 3) holds all its country's raw national data and performs the lightweight local fine-tuning benchmarked in this study, together with in-country verification of published project records. Only the model parameters of Eq. (2) and deliberately published, verified project records travel to the shared facility — never raw sensitive data — and the shared model is redistributed in return. A parity governance board modelled on the Green Climate Fund sets data-sovereignty rules and audits matching criteria but does not decide individual matches.

Governance, verification, and privacy engineering

Technical feasibility does not settle governance. The architecture's oversight is modelled on an operating precedent — the Green Climate Fund's 24-member board, evenly split between developed- and developing-country representatives, has operated parity governance at multi-billion-dollar scale since its initial resource mobilization in 2014, under a governing instrument adopted in 2011[29] — but is restricted here to setting data-sovereignty rules and auditing matching criteria. Bloc-level best-practice platforms are already normal practice — the EU's AI Factories network explicitly interconnects national hubs[12], and the China–SCO platform is building a shared project and investment database[17] — underscoring the motivating gap: most Global South countries belong to neither. Three requirements are preconditions for deployment rather than solved problems. First, privacy engineering: federated learning is not automatically privacy-preserving in the strong sense, because parameter updates can leak information through gradient-inversion attacks[30]. The countermeasures are established and complementary — differential privacy adds calibrated noise so no individual record can be inferred from updates[31]; secure aggregation encrypts individual updates so the server decrypts only their combined average[32]; bounded update clipping limits what any single record can imprint on the model[33]. These protect all participants symmetrically, which matters because a natural extension of this benchmark is federated knowledge transfer in which technologically advanced partners participate as donor silos whose accumulated engineering, financing and proven policy experience informs local models without any party exposing raw proprietary data[23]. Second, in-country verification: a coordination system that faithfully organizes fabricated project claims would launder them into credibility, so each National AI Lab requires a verification function confirming that claimed projects, contractors and outcomes are real before its

data enter any matching pipeline. Third, auditable matching: matching logic must be inspectable and bias-audited across countries, institution types and project scales before operational use.

Energy, water, and siting requirements

Any shared AI infrastructure would itself consume energy and water amid rapidly rising compute demand — global data-centre electricity use is projected to more than double to around 945 TWh by 2030[34] — and no low-carbon coordination infrastructure can credibly add to that burden with fossil power. We therefore treat 100% renewable operation, audited efficiency and transparent water accounting as design requirements for every component. The precedents are practical: LUMI runs entirely on certified hydropower and its waste heat supplies up to a fifth of its host city's district heating[14]; Gefion is hosted in a facility running on 100% renewable electricity[15]. Where freshwater or land is the binding constraint, subsea siting is a demonstrated option: Microsoft's Project Natick ran 864 servers underwater for over two years with high reliability before the programme concluded[35], and a commercial wind-powered underwater data centre off Shanghai draws over 95% of its electricity from offshore wind, eliminates freshwater cooling entirely and reduces land occupation by more than 90%[36]. Subsea siting is not environmentally free — thermal discharge requires credible marine-impact assessment, particularly under marine-heatwave conditions[35] — but for archipelagic states it converts a constraint into an option, provided ecosystem safeguards are demonstrated before scale-up.

Indonesia as first national pilot, and limitations

Indonesia is the natural first host for a National AI Lab, for reasons this paper's own evidence supplies: a national electricity plan committing to 42.6 GW of new and renewable capacity additions, more than 5,700 newly electrified villages and net-zero emissions by 2060[5,6]; a documented coordination challenge at both flagship-finance and village-data scale, resting entirely on official published statistics[8,10,11]; a fleet still 52% fossil by facility count and far more so by capacity (85.8 GW against 12.2 GW), despite utility-scale solar now being its second most numerous facility type (Fig. 1b); and existing public data sufficient — as this study demonstrates end to end — to begin without any privileged government cooperation. A first Lab's opening workloads are well defined and public-data-feasible: constructing the reconciled, version-controlled village-electrification dataset that the national programme presupposes, with transparent definitional mapping between programme, utility and village-development statistics; and site–technology matching for renewable deployment in data-sparse provinces, for which open-source geospatial least-cost electrification methods provide direct precedent[37] and the AI-for-solar forecasting literature provides the modelling toolkit[38]. Four limitations bound our claims. The benchmark task is a stand-in for full siting inference, not a planning system. Both databases are incomplete, and neither is an official national registry. The equity and scarcity effects, while directionally consistent across datasets, are small in magnitude and should not be generalized without replication on further tasks. And, most consequentially for the argument advanced here, we have not benchmarked a federated high-capacity model, so whether federation can capture the aggregation advantage that pooled gradient boosting demonstrates remains open. Within those bounds the conclusion stands: the largest barrier to pooling climate-relevant data across the Global South has never been purely technical, but the legitimate reluctance of states to cede control over sensitive information. By showing that this control can be retained at no cost — and, on current data, with a measurable gain concentrated among those with least data — federated learning offers a concrete technical basis for cooperation that is compatible with, rather than contingent on the surrender of, national sovereignty.

Methods

Data and dataset selection

The primary dataset is the Global Energy Monitor Global Integrated Power Tracker, August 2026 release[20] (CC BY 4.0). Two openly licensed datasets meet the requirements of a multi-country federated benchmark — facility-level coordinates, capacity, start year and technology for all generation types in one consistent schema — and we use both, taking GEM as primary for currency and coverage and the WRI Global Power Plant Database v1.3.0[21] (CC BY 4.0) as independent replication. GEM records generating units rather than facilities, so units were aggregated on GEM's location identifier, summing capacity, taking the earliest start year, and assigning each facility the fuel class holding the majority of its capacity; only facilities with operating status were retained. Renewable comprises utility-scale solar, wind, hydropower, bioenergy and geothermal; fossil comprises coal and oil/gas; nuclear is excluded. We pre-specified a list of 55 Global South countries across Asia, the Middle East and North Africa, sub-Saharan Africa, and Latin America and the Caribbean (full list in the analysis code) and kept the countries with ≥ 20 facilities and ≥ 5 per class (Table 1). Because GEM is distributed through a data-request form rather than a direct download, the derived facility-level table used here is deposited with the analysis code so that the primary analysis remains fully reproducible without registration; the WRI database is retrievable programmatically. Open Power System Data, the other well-known open alternative, is European in scope. No United Nations or intergovernmental agency publishes a comparable georeferenced, facility-level global asset registry; official energy statistics are compiled as national aggregates with plant-level registries held nationally — a fact that is itself part of this paper's motivation. Populated places and coastlines are from Natural Earth 110m products (public domain). Policy and infrastructure figures cited in the Introduction and Discussion are drawn from the primary and institutional sources listed in the references, accessed and verified during July–August 2026.

Features and preprocessing

For each facility we computed eight features: log10 capacity (MW); absolute latitude; raw latitude and longitude; great-circle distance to the nearest Natural Earth populated place and to the nearest coastline vertex (haversine, BallTree nearest-neighbour queries); start year (median-imputed) and a missingness indicator. Features were standardized to zero mean and unit variance using pooled training-set statistics, applied identically in all regimes as a shared preprocessing specification.

Training regimes

The classifier is the L2-regularized logistic regression of Eq. (1) ($\lambda = 10^{-3}$), trained by full-batch gradient descent (learning rate 0.5), chosen for transparency and exact compatibility with federated averaging[18]. Each country was split 70/30 into train/test, stratified by class within country; reported figures are means over ten random seeds. Local-only: an independent model per country (200 epochs). Centralized: one model on the pooled training data (200 epochs). Federated: FedAvg (Eq. 2) for 40 communication rounds of 5 local epochs each; the federation is simulated in the standard manner of the FL benchmarking literature — silo partitions are kept strictly separate and only parameters cross partition boundaries — which validates the learning mechanics a physically distributed deployment would use. Personalized FL: from the converged federated model each silo fine-tunes locally for 50 epochs[23]. In both federated regimes raw data never leave the silo.

Data-scarcity experiment

To simulate data-poor national agencies, the local-only and personalized-FL comparison was repeated with each silo's training set restricted to a stratified, class-balanced subsample of at most 40, and separately at most 16, records (minimum two per class), with the federation itself trained on the restricted sets and test sets unchanged.

Robustness analyses

Four robustness analyses accompany the main benchmark. (i) Non-linear replication: the logistic classifier was replaced by a fully connected network with one hidden layer of 12 tanh units under the identical protocol, with federated averaging applied to all weight matrices and biases; a gradient-boosted tree ensemble (histogram-based, 100–150 boosting iterations) was additionally trained locally and centrally, and re-run with latitude and longitude removed as a diagnostic. (ii) Threshold sensitivity: the silo-inclusion rule was varied to ≥ 10 and ≥ 40 facilities per country. (iii) Privacy-noise sensitivity: each silo's per-round update was clipped to unit norm and perturbed with zero-mean Gaussian noise at $\sigma \in \{0, 0.01, 0.05, 0.1, 0.25, 0.5\}$ before aggregation — the Gaussian mechanism used in differentially private training — without computing a formal privacy budget, so results are a sensitivity curve rather than a certified ε. (iv) Feature ablation: the comparison was repeated with the start-year missingness indicator removed, and again with both the indicator and start year removed, to test whether performance depends on record-completeness patterns rather than physical and geographical attributes. Robustness runs use 5 random splits (3 for threshold sensitivity).

Independent-dataset replication

The benchmark was repeated in full on the WRI Global Power Plant Database v1.3.0, which is plant-level rather than unit-level and therefore needs no aggregation step. The same 55-country Global South list, the same inclusion rule (≥ 20 plants and ≥ 5 per class), the same eight features, the same stratified 70/30 protocol, the same four regimes and the same ten random seeds were applied, retaining 31 silos and 6,545 plants. Renewable comprises hydro, solar, wind, geothermal, biomass and waste; fossil comprises coal, gas and oil; nuclear is excluded. The data-scarcity experiment and the Wilcoxon test were repeated identically. Because this database is retrievable programmatically without registration, the accompanying self-contained script downloads it and reproduces the replication column of Table 4 end to end.

Evaluation and statistics

AUC (Eq. 3) and accuracy are computed per silo on held-out test data, macro-averaged across silos with a valid two-class test set, and reported as mean ± s.d. over ten seeds. For the distributional analysis we computed, per silo, the seed-averaged difference in AUC between personalized FL and local-only training, and tested the null hypothesis of zero median difference with a two-sided Wilcoxon signed-rank test across silos ($p = 0.048$ on the primary dataset; $p = 0.15$ on the replication dataset).

Reproducibility

The pipeline uses only openly licensed data and a standard scientific Python environment (NumPy, pandas, scikit-learn) and is deterministic given the seeds. A self-contained script reproduces the replication-dataset analysis end to end including data download; the primary analysis runs from the deposited derived dataset. A plain-language guide enables non-programmers to execute and verify the results. Analyses were executed during July–August 2026.

Data Availability

All data are openly licensed and publicly available. The primary dataset is the Global Energy Monitor Global Integrated Power Tracker, August 2026 release (CC BY 4.0), obtained through the tracker's public data-request form; the derived facility-level table analysed here is deposited with the code. The replication dataset is the WRI Global Power Plant Database v1.3.0 (CC BY 4.0). Coastlines and populated places are from Natural Earth (public domain). No new primary data were generated.

Code availability

The complete analysis code, including a self-contained reproduction script and the derived primary dataset, is archived in a public repository with a permanent DOI [https://doi.org/10.5281/zenodo.21973074].

Acknowledgements

The author thanks Dr. Zhu Liu (Department of Earth System Science, Tsinghua University) for supervision and guidance throughout this research. The author gratefully acknowledges the Ministry of Agrarian Affairs and Spatial Planning / National Land Agency (ATR/BPN) of the Republic of Indonesia, Directorate General of Spatial Planning, for its recommendation of the research project of which this study forms part, and thanks Ir. Reny Windyawati, S.T., M.Sc. and Mr. Dedi Noor Cahyanto for facilitating that recommendation. The Ministry's recommendation relates to the research project and does not constitute endorsement of the findings reported here, for which the author is solely responsible.

Author contributions

S.D. conceived and designed the study, directed the analysis, executed and verified the reproduction of all results, and wrote the manuscript.

Competing interests

The author declares the following interests. The author is a former senior advisor for strategic economic affairs at the Executive Office of the President, Republic of Indonesia, and has advised other Government of Indonesia priority programmes. The author is affiliated with a collaborative research programme that intends to pursue federated AI infrastructure for low-carbon development in the Global South as a policy initiative and potentially funded venture, with Indonesia envisaged as a first pilot. Accordingly, all statements about potential infrastructure in the Discussion are framed as requiring independent economic, technical, governance and environmental evaluation, and no such system is represented as validated by the present study, whose empirical claims are confined to the benchmark reported in Results.

Use of AI-assisted technologies

AI assistance (Anthropic Claude) was used for literature identification, drafting and development of the analysis code under the author's direction. The analyses were executed as described in Methods and their outputs independently re-run and verified by the author. All content was reviewed and approved by the author, who takes full responsibility for it. AI tools are not authors, consistent with Nature Portfolio policy.

Funding Declaration

No funding was received for conducting this study.

References

1. IPCC. Climate Change 2023: Synthesis Report. Summary for Policymakers (IPCC, Geneva, 2023); IPCC. Climate Change 2022: Mitigation of Climate Change. Summary for Policymakers (Cambridge Univ. Press, 2022).

2. International Energy Agency. Net Zero Roadmap: A Global Pathway to Keep the 1.5 °C Goal in Reach — 2023 Update (IEA, Paris, 2023).

3. International Renewable Energy Agency. Renewable Power Generation Costs in 2023 (IRENA, Abu Dhabi, 2024).

4. United Nations Environment Programme. Adaptation Gap Report 2025: Running on Empty (UNEP, Nairobi, 2025).

5. Kementerian ESDM & PT PLN (Persero). Rencana Usaha Penyediaan Tenaga Listrik (RUPTL) PLN 2025–2034 (ratified 26 May 2025); summarised in IESR, Government needs to ensure strategy to achieve renewable energy target in RUPTL 2025–2034 (27 May 2025). https://iesr.or.id/en/government-needs-to-ensure-strategy-to-achieve-renewable-energy-target-in-ruptl-2025-2034/

6. ECADIN. Indonesia Electricity Supply Business Plan 2025–2034 (2025). https://ecadin.org/indonesia-electricity-plan-ruptl-pln-2025-2034/

7. JETP Secretariat. Comprehensive Investment and Policy Plan 2023: Just Energy Transition Partnership Indonesia (Jakarta, 21 November 2023). https://jetp-id.org

8. Institute for Energy Economics and Financial Analysis. Building credibility in Indonesia's energy transition: insights from the ETM and JETP Indonesia (IEEFA, February 2026). https://ieefa.org/resources/building-credibility-indonesias-energy-transition-insights-etm-and-jetp-indonesia

9. Selvaraju, S., Indah Pratiwi, A., Sabogal, L. & Ahlgren, V. Just Energy Transition Partnership grants and country platforms: lessons from Indonesia and South Africa (Grantham Research Institute, LSE, 2025).

10. Kementerian ESDM. Kementerian ESDM Tuntaskan Program Elektrifikasi di Tahun 2024 (media release, 15 January 2024). https://www.esdm.go.id/id/media-center/arsip-berita/kementerian-esdm-tuntaskan-program-elektrifikasi-di-tahun-2024

11. Direktorat Jenderal Ketenagalistrikan, Kementerian ESDM. Meningkat, Rasio Desa Berlistrik dan Rasio Elektrifikasi Triwulan III 2024 Ditetapkan. Press release, 6 November 2024, reporting the Rural Electrification (Lisdes) Monitoring and Evaluation Meeting for Sulawesi, Maluku, Papua and Nusa Tenggara, Makassar, 6 November 2024; figures stated by Nur Hidayanto, Coordinator for Electricity Distribution Planning. Directorate General news portal: https://gatrik.esdm.go.id/berita (search by title). Reported in: Kontan, Pemerintah Tingkatkan Rasio Elektrifikasi, 6.072 Desa Masih Belum Diterangi Listrik (8 November 2024). https://industri.kontan.co.id/news/pemerintah-tingkatkan-rasio-elektrifikasi-6072-desa-masih-belum-diterangi-listrik

12. EuroHPC Joint Undertaking. The EuroHPC JU Selects Six Additional AI Factories to Expand Europe's AI Capabilities (10 October 2025). https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-six-additional-ai-factories-expand-europes-ai-capabilities-2025-10-10_en; European Commission. EU expands network of AI Factories, strengthening its AI continent ambition (10 October 2025). https://digital-strategy.ec.europa.eu/en/news/eu-expands-network-ai-factories-strengthening-its-ai-continent-ambition

13. European Investment Bank. EIB Group and European Commission join forces to finance AI gigafactories (2025). https://www.eib.org/en/press/all/2025-491-eib-group-and-european-commission-join-forces-to-finance-ai-gigafactories

14. EuroHPC Joint Undertaking. Inauguration of LUMI: the fastest and greenest supercomputer in Europe (13 June 2022). https://www.eurohpc-ju.europa.eu/inauguration-lumi-fastest-greenest-supercomputer-europe-2022-06-13_en; LUMI consortium. LUMI data center will be powered with hydroelectricity supplied by Vattenfall (2020). https://lumi-supercomputer.eu/lumi-data-center-will-be-powered-with-hydroelectricity-supplied-by-vattenfall/

15. Novo Nordisk Foundation. Denmark's first AI supercomputer is now operational (23 October 2024). https://novonordiskfonden.dk/en/news/denmarks-first-ai-supercomputer-is-now-operational/

16. OpenAI. Announcing The Stargate Project (21 January 2025). https://openai.com/index/announcing-the-stargate-project/

17. Shanghai Cooperation Organisation Secretariat. A delegation of the International Cooperation Center under the NDRC of the People's Republic of China visits the SCO Secretariat (19 May 2026). https://eng.sectsco.org/20260519/2314424.html

18. McMahan, H. B., Moore, E., Ramage, D., Hampson, S. & Agüera y Arcas, B. Communication-efficient learning of deep networks from decentralized data. Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS) 54, 1273–1282 (2017).

19. Kairouz, P. et al. Advances and open problems in federated learning. Found. Trends Mach. Learn. 14, 1–210 (2021).

20. Global Energy Monitor. Global Integrated Power Tracker, August 2026 release. https://globalenergymonitor.org/projects/global-integrated-power-tracker/

21. Global Energy Observatory, Google, KTH Royal Institute of Technology, Enipedia, World Resources Institute. Global Power Plant Database v1.3.0 (Resource Watch, 2021); Byers, L. et al. A Global Database of Power Plants (WRI Technical Note, 2018).

22. Li, Q., Diao, Y., Chen, Q. & He, B. Federated learning on non-IID data silos: an experimental study. Proc. IEEE 38th Int. Conf. Data Engineering (ICDE) 965–978 (2022).

23. Tan, A. Z., Yu, H., Cui, L. & Yang, Q. Towards personalized federated learning. IEEE Trans. Neural Netw. Learn. Syst. 34, 9587–9603 (2023).

24. Cheng, K. et al. SecureBoost: a lossless federated learning framework. IEEE Intell. Syst. 36, 87–98 (2021).

25. Li, Q. et al. FedTree: a federated learning system for trees. Proc. Machine Learning and Systems (MLSys) 5 (2023).

26. Al Garni, H. Z. & Awasthi, A. Solar PV power plant site selection using a GIS-AHP based approach with application in Saudi Arabia. Appl. Energy 206, 1225–1240 (2017).

27. Abram, N. J. et al. Quantifying the regional to global climate impacts of individual fossil fuel projects to inform decision-making. npj Clim. Action 4, 92 (2025). https://doi.org/10.1038/s44168-025-00296-5

28. Heyndrickx, W. et al. MELLODDY: cross-pharma federated learning at unprecedented scale unlocks benefits in QSAR without compromising proprietary information. ChemRxiv (2022). https://doi.org/10.26434/chemrxiv-2022-ntd3r

29. Green Climate Fund. About GCF: governance and Board (accessed July 2026). https://www.greenclimate.fund

30. Zhu, L., Liu, Z. & Han, S. Deep leakage from gradients. Adv. Neural Inf. Process. Syst. 32 (2019).

31. Dwork, C. & Roth, A. The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci. 9, 211–407 (2014).

32. Bonawitz, K. et al. Practical secure aggregation for privacy-preserving machine learning. Proc. 2017 ACM SIGSAC Conf. Computer and Communications Security (CCS) 1175–1191 (2017).

33. Abadi, M. et al. Deep learning with differential privacy. Proc. 2016 ACM SIGSAC Conf. Computer and Communications Security (CCS) 308–318 (2016).

34. International Energy Agency. Energy and AI (IEA, Paris, 2025).

35. Microsoft. Project Natick: Microsoft finds underwater datacenters are reliable, practical and use energy sustainably (2020). https://news.microsoft.com/source/features/sustainability/project-natick-underwater-datacenter/; Scientific American. China powers AI boom with undersea data centers (2025). https://www.scientificamerican.com/article/china-powers-ai-boom-with-undersea-data-centers/

36. Sustainability Magazine. How sustainable is China's wind-powered, subsea data centre? (2025). https://sustainabilitymag.com/news/how-sustainable-is-chinas-wind-powered-subsea-data-centre

37. ESMAP / World Bank. Global Electrification Platform and the Open Source Spatial Electrification Tool (OnSSET) (accessed July 2026). https://electrifynow.energydata.info

38. Husein, M., Gago, E. J., Hasan, B. & Pegalajar, M. C. Towards energy efficiency: a comprehensive review of deep learning-based photovoltaic power forecasting strategies. Heliyon 10, e33419 (2024). https://doi.org/10.1016/j.heliyon.2024.e33419