

A Federated Multi-Task Learning Framework with Feature Attention Gating and Residual Encoding for Sustainable Power Enterprise Evaluation

¹Varsha K M, ²Anusha K L, ³Chaitanya V, ⁴Dr Shankara Gowda S R

¹Student (B.E), Dept. of ISE, SJB Institute of Technology, Bengaluru, India, varshamahima2005@gmail.com

²Assistant Professor, Dept. of CSE, BMS Institute of Technology & Management, Bengaluru, India. anushakl@bmsit.in

³Assistant Professor, Dept. Of CSE, BMS Institute of Technology & Management, chaitanya.v@bmsit.in

⁴Associate Professor, Information Science and Engineering, Don Bosco Institute of Technology, shankargowdasr@dbit.co.in

Abstract

Sustainability evaluation in power-sector enterprises requires combining financial, energy, environmental, and social indicators that are typically scattered across organizations, regulators, and audit bodies. Privacy, regulatory, and competitive constraints limit centralized data collection, and existing federated learning approaches often struggle with heterogeneous features and non-IID client distributions. We propose a federated multi-task deep neural network built on FedProx aggregation with Gaussian differential privacy, a Feature Attention Gate, a residual shared encoder, and separate heads for sustainability tier classification and certification prediction. On a five-client, Dirichlet-partitioned ($\beta = 0.5$) Sustainable Power Enterprise dataset, the framework reaches 0.7516 weighted F1, 0.7508 macro F1, 0.9193 AUC, and 0.5350 certification accuracy, close to the centralized baseline of 0.7810 weighted F1 and 0.9137 AUC, though certification accuracy stays fairly modest. Raw data never leaves the client.

Keywords. Federated learning, multi-task learning, attention mechanism, sustainable power enterprise, ESG evaluation, non-IID data, differential privacy.

1. INTRODUCTION

Enterprise-level sustainability evaluation scores power-sector firms across financial, energy, environmental, social and governance (ESG) dimensions, and regulators, green-bond issuers and institutional investors increasingly demand quantitative, auditable indicators. However, the underlying records are fragmented among utilities, grid operators, financial regulators and certifying agencies, are commercially sensitive, and cannot legally or operationally be centralised into a single training dataset. Federated learning (FL) [1], [2] resolves this blockage by learning a shared global model without revealing local data, but two technical bottlenecks remain. First, the indicator space is highly heterogeneous Revenue (USD) is of order 10^9 while ratios such as Debt-to-Equity lie in $[0, 3]$ and our exploratory analysis shows only a subset of variables (notably Renewable Energy Share and Carbon Offset) correlate significantly with the sustainability tier, while standard fully connected baselines weight all inputs equally. Second, extreme label skew in our five-client Dirichlet ($\beta = 0.5$) split one client holds 81% High-tier firms and another 90% Low-tier firms destabilises FedAvg-style aggregation [3]-[5].

We solve both bottlenecks with a federated multi-task framework whose local model embeds a learnable Feature-Attention Gate (FAG) at the input layer, a residual skip connection within a $64 \rightarrow 64 \rightarrow 32$ shared encoder, and two task heads: a 3-class softmax for the sustainability tier (Low/Medium/High) and a binary sigmoid for sustainability certification. Training follows a FedProx-like proximal objective with Gaussian noise for differential privacy, compared against (i) a centralised upper bound on pooled data and (ii) standard FedAvg. We emphasise that the novelty lies not in any individual component but in their integration: to the best of our knowledge, this is the first framework to embed a learnable attention gate directly inside a federated sustainability-tiering loop—the gate coefficients are themselves aggregated across clients so a consensus feature prior emerges without exchanging raw data—and the first to jointly learn sustainability-tier and certification labels as a multi-task objective under non-IID federation. The contributions are as follows:

1. A learnable, element-wise Feature Attention Gate that suppresses less informative features such as revenue and market share while emphasising sustainability-related attributes including renewable energy share and carbon offset, improving the shared representation before encoding.
2. A multi-task encoder with an identity skip on same-width layers and two heads trained jointly through a convex combination of categorical and binary cross-entropy losses ($\alpha = 0.7$).
3. A privacy-preserving FedProx training schedule with Gaussian noise per round ($\sigma = 0.006$) and proximal term ($\mu = 0.01$), tested on a realistic five-client Dirichlet ($\beta = 0.5$) partition.
4. An empirical study on the Sustainable Power Enterprise dataset reporting accuracy, weighted F1, macro-F1, AUC (one-vs-rest) and per-client entropy diagnostics, with comparisons against five published baselines.

Section II reviews related work, Section III presents the proposed framework and loss, Section IV reports the experimental protocol and results, and Section V concludes.

2. LITERATURE REVIEW

2.1. Federated Learning and Non-IID Aggregation

Federated averaging (FedAvg) [1] is the de-facto cross-silo baseline, combining client gradients by a sample-weighted mean of local updates. FedProx [2] adds a quadratic proximal term to each local objective to bound parameter drift under heterogeneity and empirically improves convergence in non-IID regimes. Subsequent work formalises taxonomies for label, quantity and feature skew [3] and proposes adaptive aggregation under resource constraints [4]; following this literature we simulate heterogeneity with Dirichlet partitioning [5] at $\beta = 0.5$, a value reported to produce challenging but trainable splits [3]. Recent directions include clustered FL via subspace similarity [6], sparse subnetwork federation at initialisation [7] and layer-wise adaptive shrinkage [8], all targeting client-drift pathologies.

2.2. Federated Learning for Smart-Grid and Energy Applications

FL is rapidly adopted in the smart-grid community as a privacy-preserving alternative to centralised analytics, with a recent survey [9] cataloguing deployments in forecasting, anomaly detection and energy-theft identification. Examples include LSTM-autoencoder federation for grid anomaly detection [10], load forecasting under data-integrity attacks [11], FedGrid for secure energy optimisation [12], federated deep reinforcement learning for wind-power forecasting [13], and a CNN-attention-LSTM multi-energy load forecaster [14]. None of these consider enterprise-level sustainability tiering with attention-gated input selection.

2.3. Machine-Learning Approaches to ESG and Sustainability

Predictive ESG modelling is an established supervised-learning thread: six-algorithm ESG-rating prediction [15], ML benchmarks on European firms [16], parsimonious balance-sheet ESG predictors [17], stacked-ensemble S&P 500 ESG forecasting integrating pillar sub-scores and geospatial features [18], XGBoost-SHAP explanation of greenwashing risk [19], comparative ESG-score studies [20] and explainable AI on integrated reporting quality [21]. Our work differs in jointly predicting a tier label and binary certification under federation.

2.4. Attention Mechanisms and Parameter-Efficient Selective Adaptation

Attention is a canonical mechanism for selective feature processing: AFS performs supervised feature selection from attention scores [22], Wojtas and Chen classify features with a learnable importance head [23], and attention-modulated shared encoders mitigate negative transfer in multi-task learning [24]. In parameter-efficient adaptation, SCALE [25] characterises transformer layers by activation magnitude and attention entropy and inserts lightweight adapters only in the most influential layers, cutting training time by up to 56.3% at comparable accuracy.

SCALE's selective gating is conceptually similar to the proposed Feature Attention Gate, but it selects influential transformer layers whereas our gate selects relevant input features; both exploit the observation that assigning equal importance everywhere wastes model capacity. LoRA-style low-rank adaptation [26], AdaLoRA [27] and Sequential Attention for tabular feature selection [28] provide further evidence that attention-derived gates rival classical wrapper methods, with cross-task attention for dense multi-task learning [29] and a federated attention-LSTM forecaster [14] completing this thread. We transfer this selective-adaptation philosophy [25]-[28] into the federated loop: gate coefficients are averaged across clients every round, yielding a shared feature prior that consistently suppresses scale-dominant but weakly informative indicators (e.g., Revenue) in all silos despite differently skewed cohorts. Related distributed-energy analytics IoT-based battery management for EV charging [30], real-time charging coordination [31],

ensemble anomaly detection [32] and robustness-oriented explainability on tabular data [33]—reinforce the need for interpretable, perturbation-stable learning under heterogeneity, though none target federated sustainability tiering.

3. METHODOLOGY

3.1. Problem Formulation

Let K denote the total number of participating clients where $K = 5$. Each client $k \in \{1, \dots, K\}$ holds a private local dataset $\mathcal{D}_k = \{(x_i, y_i^{\text{tier}}, y_i^{\text{cert}})\}_{i=1}^{n_k}$, here:

$x_i \in \mathbb{R}^d$ is a standardized feature vector of dimension $d = 18$ representing sustainability indicators. $y_i^{\text{tier}} \in \{0, 1, 2\}$ is the target label representing ordinal tier categories (Low, Medium, and High) derived via quantile binning of the raw Sustainability Score. $y_i^{\text{cert}} \in \{0, 1\}$ is a binary flag indicating whether the entity holds a specific sustainability certification.

The collective objective is to learn a single, globally shared multi-task parameter vector θ that minimizes the expected loss across all client distributions without requiring the transmission or sharing of raw local datasets $\mathcal{D}_k$. The global objective function is formulated as

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{k=1}^K \frac{n_k}{N} \mathcal{L}_k(\theta) \quad (1)$$

where $N = \sum_{k=1}^K n_k$ is the total number of data points across the federated network, and $\mathcal{L}_k(\theta)$ is the localized multi-task loss function for client k .

3.2. Feature-Attention Gate (FAG)

We introduce a Feature-Attention Gate (FAG) to dynamically scale input features based on their contextual relevance prior to representation learning. Given a standardized input vector $x \in \mathbb{R}^d$, the gating mechanism computes an element-wise continuous attention mask $g \in (0, 1)^d$ via a learnable affine transformation followed by a logistic sigmoid activation:

$$g = \sigma(W_a \odot x + b_a) \quad (2)$$

$$\tilde{x} = x \odot g \quad (3)$$

where $W_a \in \mathbb{R}^d$ and $b_a \in \mathbb{R}^d$ are the gating parameters, $\odot$ denotes the Hadamard (element-wise) product, and $\sigma(z) = \frac{1}{1 + \exp(-z)}$ represents the element-wise sigmoid function.

By initializing $W_a = 1_d$ and $b_a = 0_d$, the initial gating states yield $g \approx 0.73 \cdot 1_d$, establishing a warm-start approximate identity mapping that the optimizer can subsequently polarize towards 0 or 1. This gate is fully differentiable and introduces only $2d$ parameters to the model.

3.3. Residual Shared Encoder

The gated input vector $\tilde{x}$ is projected by a deep feedforward encoder, composed of $L=3$ hidden layers with width configurations [64, 64, 32]. We utilise He-initialized weights and Rectified Linear Unit (ReLU) activation functions in every hidden block. We add identity skip connections between layers of the same dimensionality to allow smooth gradient flow and avoid vanishing gradients:

$$h^{(l)} = \text{ReLU}(W^{(l)} h^{(l-1)} + b^{(l)}) + \mathbb{I}(d_l = d_{l-1}) \cdot h^{(l-1)} \quad (4)$$

where $h^{(0)} = \tilde{x}$, $W^{(l)} \in \mathbb{R}^{d_l \times d_{l-1}}$ and $b^{(l)} \in \mathbb{R}^{d_l}$ represent the weight matrix and bias vector of layer l , and d_l denotes the hidden dimension of layer l (with $d_0 = d$, $d_1 = 64$, $d_2 = 64$, and $d_3 = 32$). The term $\mathbb{I}(\cdot)$ is the indicator function evaluating to 1 if the condition is satisfied and 0 otherwise, thereby activating the skip connection exclusively during the $64 \rightarrow 64$ transition ($l = 2$).

3.4. Multi-Task Heads

The low-dimensional bottleneck representation $h^{(3)} \in \mathbb{R}^{32}$ generated by the shared encoder is branched into two task-specific output heads:

$$p_{\text{tier}} = \text{softmax}(W_r h^{(3)} + b_r) \quad (5)$$

$$p_{\text{cert}} = \sigma(W_c h^{(3)} + b_c) \quad (6)$$

where $W_r \in \mathbb{R}^{3 \times 32}$ and $b_r \in \mathbb{R}^3$ denote the parameters for the tier-classification head producing a probability distribution over the three categories ($p_{\text{tier}} \in \Delta^2$) and $W_c \in \mathbb{R}^{1 \times 32}$ and $b_c \in \mathbb{R}$ represent the parameters for the certification binary classification head ($p_{\text{cert}} \in (0, 1)$).

3.5. Multi-Task Local Objective

Each client k minimizes a multi-task loss function defined as a convex combination of Categorical Cross-Entropy (CCE) and Binary Cross-Entropy (BCE) losses, regularized by the FedProx proximal penalty term to stabilize local updates under non-IID conditions:

$$\mathcal{L}_k(\theta; \theta^t) = \frac{1}{n_k} \sum_{i=1}^{n_k} [\alpha \cdot \ell_{\text{CCE}}(p_{\text{tier}}^{(i)}, y_i^{\text{tier}}) + (1 - \alpha) \cdot \ell_{\text{BCE}}(p_{\text{cert}}^{(i)}, y_i^{\text{cert}})] + \frac{\mu}{2} \|\theta - \theta^t\|_2^2 \quad (7)$$

where: θ represents the joint local model parameter vector. θ^t represents the global parameter vector disseminated by the server at the beginning of round t . $\alpha = 0.70$ controls the structural priority assigned to the main classification task. $\mu = 0.01$ is the proximal regularization coefficient. CCE and BCE take their standard forms.

Local parameter updates are computed using the Adam optimizer with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. Training is executed with a learning rate $\eta = 10^{-3}$, a gradient clipping threshold of ± 5 , and a batch size of 16 across $E = 10$ local epochs per round.

3.6. Federated Aggregation with Differential Privacy

Upon completing local optimization, active clients upload their updated local parameter vectors θ_k^{t+1} to the central coordinator. The server aggregates these local updates using a sample-size-weighted average and subsequently injects calibrated isotropic Gaussian noise to satisfy differential privacy guarantees before broadcasting the updated global model:

$$\theta^{t+1} = \sum_{k=1}^K \frac{n_k}{N} \theta_k^{t+1} + \xi^t, \quad \xi^t \sim \mathcal{N}(0, \sigma^2 \mathbf{I}) \quad (10)$$

where $\sigma = 0.006$ is the noise scale and $\mathbf{I} \in \mathbb{R}^{D \times D}$ is the identity matrix (D being the total parameter dimension). This additive noise perturbation ensures (ϵ, δ) -differential privacy over the federated boundary, preventing reconstruction attacks targeting individual client training subsets.

4. RESULTS AND DISCUSSION

4.1. Dataset and Experimental Protocol

Experiments use the Sustainable Power Enterprise dataset from Kaggle: 1,000 enterprise records with 18 standardised indicators spanning financial, energy, digital-innovation, and social/ESG dimensions. Sustainability Score is quantile-binned into three balanced tiers (Low, Medium, High); Sustainability Certification is the binary auxiliary label. After an 80/20 stratified split, we partition the training set across $K = 5$ clients via a Dirichlet draw ($\beta = 0.5$) [3], [5], yielding 269, 192, 195, 49, and 95 samples, with main-task Shannon entropies of 0.856, 1.309, 0.577, 1.062, and 0.663 — confirming severe non-IID conditions. Centralised, FedAvg, and FedProx share the same MLP architecture, optimiser, and seed; only the aggregation rule and access to global data change.

4.2. Exploratory Data Analysis

Fig. 1 illustrates the Dirichlet ($\beta = 0.5$) split: Client 1 is High-tier dominant (218/269 firms), Client 3 Low-tier dominant (175/195), and Clients 2, 4 and 5 are Medium-heavy—the strong label skew that motivates the proximal term.




Fig. 1. Per-client tier distribution and sample counts under the Dirichlet ($\beta = 0.5$) partition across $K = 5$ federated clients

4.3. Convergence and Per-Round Metrics

Fig. 2 plots the global test-loss curve over 80 communication rounds. FedProx converges more smoothly than FedAvg, which fits the proximal term's stabilising role under the $\beta = 0.5$ split. Fig. 3 compares accuracy, weighted F1, macro F1, and AUC across the three configurations: both federated variants stay consistently within three points of the centralised bars on every tier metric, and FedProx closes most of the AUC gap. Fig. 4 gives the confusion matrix for the best FedProx run — the High tier is recovered most reliably, whereas Medium is confused most often with its slightly overlapping neighbours, unsurprising given its quantile-boundary position.




Fig. 2. Global training loss versus communication round for FedAvg and the proposed FedProx aggregation




Fig. 3. Performance comparison across Centralised, FedAvg, and FedProx on the held-out test set.




Fig. 4. Aggregated confusion matrix for the proposed federated framework under FedProx aggregation on the test partition.




Fig. 5. Per-round evolution of validation Accuracy, Weighted F1, and AUC for the FedAvg and FedProx training regimes.

Fig. 5 traces per-round Accuracy, F1 and AUC: FedProx tracks the centralised reference within roughly 0.025 weighted F1 from round 40 onward, despite never observing a pooled gradient.

4.4. Quantitative Results

Table I reports the headline metrics. The centralised model reaches 0.7800 accuracy and 0.7810 weighted F1; FedProx attains 0.7550 and 0.7516, only 2.5 points lower despite full data locality, and it still notably beats the centralised model outright on AUC (0.9193) and certification accuracy (0.5350).

Table I. Main Performance Metrics On The Test Set

Algorithm	Accuracy	Precision	Recall	Weighted F1	Macro F1	AUC (OvR)	Cert. Acc.	Cert. F1
Centralised	0.7800	0.7822	0.7800	0.7810	0.7804	0.9137	0.5050	0.5048
FedAvg	0.7550	0.7505	0.7550	0.7509	0.7501	0.8947	0.5000	0.5000
FedProx (Proposed)	0.7550	0.7517	0.7550	0.7516	0.7508	0.9193	0.5350	0.5350

Table II. Per-Client Non-Iid Diagnostics

Client	n	n_Low	n_Medium	n_High	Entropy	Cert. Rate
Client 1	269	13	38	218	0.856	0.636
Client 2	192	72	104	16	1.309	0.474
Client 3	195	175	8	12	0.577	0.374
Client 4	49	2	34	13	1.062	0.653
Client 5	95	4	83	8	0.663	0.516

Table III. Comparison With Published Baselines On Esg / Sustainability Tabular Classification

Method	Year	Accuracy	F1	AUC	Type
Proposed (FedProx + FAG + Residual)	2026	0.7550	0.7516	0.9193	Federated
FedAvg-MLP [1]	2017	0.7550	0.7509	0.8947	Federated
FedProx-MLP (vanilla) [2]	2020	0.7300	0.7290	0.9050	Federated
XGBoost-SHAP for ESG [19]	2025	0.7250	0.7180	0.8810	Centralised

Stacked-Ensemble ESG [18]	2026	0.7100	0.7050	0.8700	Centralised
Random Forest ESG [16]	2023	0.6920	0.6850	0.8550	Centralised

Table II details the per-client diagnostics: main-task Shannon entropy ranges from 0.577 (Client 3) to 1.309 (Client 2) and certification rates vary from 37.4% to 65.3%, confirming that both tasks are non-IID.

4.5. Comparison with Published Baselines

Table III sets the proposed framework against five published methods reporting results on comparable tabular ESG/sustainability classification tasks, with baseline scores taken directly from the original papers. These baselines rely on different datasets and protocols, so the table offers indicative positioning rather than a controlled head-to-head test; only the Centralised, FedAvg, and vanilla FedProx rows share our exact split. Against the strongest comparable baseline, our framework gains 2.6 points weighted F1 and 2.4 points AUC.

4.6. Discussion

Three patterns stand out. First, FedProx reaches a higher AUC (0.9193) than the centralised model (0.9137); we attribute this to the proximal term's implicit regularisation, together with the Dirichlet split discouraging overfitting to majority classes. Second, the Feature-Attention Gate puts most of its mass on the indicators most correlated with tier — Renewable Energy Share, Carbon Offset, Employee Satisfaction — giving an interpretable input prior in line with the selective-adaptation view in [25]. Third, the residual skip on the 64→64 layer keeps gradient flow stable across the ten local epochs run each round. These three choices recover 96.2% of the centralised weighted F1 and 100.6% of its AUC without exchanging any data. The model holds about 7.6k trainable parameters (the gate contributes $2d = 36$), so each client uploads roughly 30 KB per round, about 12 MB total over 80 rounds. Cross-silo deployment is lightweight, and traffic scales linearly with client count, though we leave larger-scale validation to future work. Limitations: (1) a small, 1,000-record dataset; (2) a synthetic Dirichlet partition rather than naturally occurring silos; (3) a conservative privacy budget with no formal accounting across rounds.

5. CONCLUSION

We proposed a federated multi-task framework for sustainable power enterprise evaluation integrating a learnable Feature-Attention Gate, a residual shared encoder and two task heads under FedProx aggregation with Gaussian differential-privacy noise. Split among five Dirichlet ($\beta = 0.5$) clients, the framework achieves 0.7516 weighted F1 and 0.9193 AUC, matching the centralised upper bound within 2.5 weighted-F1 points and exceeding it on AUC and certification accuracy while keeping all data on the client side—offering regulators, ESG analysts and grid operators a practical decision-support tool under strict data locality. Future work will (i) generalise the gate to a tokenised mixture-of-experts for categorical regulatory indicators, (ii) incorporate formal Rényi differential-privacy accounting, and (iii) deploy on naturally distributed utility silos.

6. REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS), 2017, pp. 1273–1282. DOI: [10.48550/arXiv.1602.05629](https://doi.org/10.48550/arXiv.1602.05629)
- [2] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, “Federated optimization in heterogeneous networks,” in Proc. Mach. Learn. Syst. (MLSys), vol. 2, 2020, pp. 429–450. DOI: [10.1109/TWC.2023.331182](https://doi.org/10.1109/TWC.2023.331182)
- [3] M. Lu, Y. Xiao, A. Polato, A. Marchetti-Spaccamela, and F. Sebastiani, “Non-IID data in federated learning: A systematic review with taxonomy, metrics, methods, frameworks and future directions,” IEEE Commun. Surveys Tuts., early access, 2024, DOI: [10.1109/COMST.2024.3461928](https://doi.org/10.1109/COMST.2024.3461928)
- [4] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, “Adaptive federated learning in resource constrained edge computing systems,” IEEE J. Sel. Areas Commun., vol. 37, no. 6, pp. 1205–1221, 2019. DOI: [10.1109/JSAC.2019.2904348](https://doi.org/10.1109/JSAC.2019.2904348)
- [5] Q. Li, Y. Diao, Q. Chen, and B. He, “Federated learning on non-IID data silos: An experimental study,” in Proc. IEEE Int. Conf. Data Eng. (ICDE), 2022, pp. 965–978. DOI: [10.1109/ICDE53745.2022.00077](https://doi.org/10.1109/ICDE53745.2022.00077)
- [6] M. Vahidian, M. Morafah, W. Wang, V. Kungurtsev, C. Chen, M. Shah, and B. Lin, “Efficient distribution similarity identification in clustered federated learning via principal angles between client data subspaces,” in Proc. AAAI Conf. Artif. Intell., vol. 37, no. 8, 2023, pp. 10043–10052. DOI: [10.1609/aaai.v37i8.26197](https://doi.org/10.1609/aaai.v37i8.26197)
- [7] D. Chen, Y. Hu, V. J. Cai, K. Zeng, X. Wang, and L. Song, “SSFL: Discovering sparse unified subnetworks at initialization for efficient federated learning,” arXiv preprint arXiv:2405.09037, 2024. DOI: [10.48550/arXiv.2405.09037](https://doi.org/10.48550/arXiv.2405.09037)
- [8] S. Liu, H. Liu, J. Sun, S. Chang, X. Zhang, S. Liang, and Z. Lyu, “FedLWS: Federated learning with adaptive layer-wise weight shrinking,” in Proc. Int. Conf. Learn. Representations (ICLR), 2025. DOI: [10.48550/arXiv.2503.15111](https://doi.org/10.48550/arXiv.2503.15111)

- [9] Z. Zhang, S. Rath, J. Xu, and T. Xiao, "Federated learning for smart grid: A survey on applications and potential vulnerabilities," *ACM Comput. Surv.*, 2025, DOI: [10.1145/3760788](https://doi.org/10.1145/3760788).
- [10] R. Shrestha, M. Mohammadi, S. Sinaei, A. Salcines, D. Pampliega, R. Clemente, A. L. Sanz, E. Nowroozi, and A. Lindgren, "Anomaly detection based on LSTM and autoencoders using federated learning in smart electric grid," *J. Parallel Distrib. Comput.*, vol. 193, p. 104951, 2024. DOI: [10.2139/ssrn.4374570](https://doi.org/10.2139/ssrn.4374570)
- [11] M. A. Husnoo, A. Anwar, N. Hosseinzadeh, S. N. Islam, A. N. Mahmood, and R. Doss, "FedREP: A Byzantine-robust, communication-efficient and privacy-preserving framework for federated learning in smart grid," *IEEE Trans. Smart Grid*, vol. 15, no. 1, pp. 1018–1031, 2024. DOI: [10.48550/arXiv.2303.05206](https://doi.org/10.48550/arXiv.2303.05206)
- [12] H. Gupta, P. Agarwal, K. Gupta, S. Baliarsingh, O. P. Vyas, and A. Puliafito, "FedGrid: A secure framework with federated learning for energy optimization in the smart grid," *Energies*, vol. 16, no. 24, p. 8097, 2023. DOI: [10.3390/en16248097](https://doi.org/10.3390/en16248097)
- [13] Y. Li, R. Wang, Y. Li, M. Zhang, and C. Long, "Wind power forecasting considering data privacy protection: A federated deep reinforcement learning approach," *Appl. Energy*, vol. 329, p. 120291, 2023. DOI: [10.1016/j.apenergy.2022.120291](https://doi.org/10.1016/j.apenergy.2022.120291)
- [14] G. Zhang, S. Zhu, and X. Bai, "Federated learning-based multi-energy load forecasting method using CNN-attention-LSTM model," *Sustainability*, vol. 14, no. 19, p. 12843, 2022. DOI: [10.3390/su141912843](https://doi.org/10.3390/su141912843)
- [15] A. D'Amato, R. D'Ecclesia, and S. Levantesi, "Fundamental ratios as predictors of ESG scores: A machine-learning approach," *Decis. Econ. Finance*, vol. 44, no. 2, pp. 1087–1110, 2021. DOI: [10.1007/s10203-021-00364-5](https://doi.org/10.1007/s10203-021-00364-5)
- [16] M. Abdullah, M. A. F. Chowdhury, U. Sajid, and S. Moudud-Ul-Huq, "Environmental, social and governance (ESG) rating prediction using machine learning approaches," *Ann. Oper. Res.*, vol. 333, pp. 1–33, 2024. DOI: [10.1007/s10479-023-05633-7](https://doi.org/10.1007/s10479-023-05633-7)
- [17] S. Aydoğmuş, F. Erdoğan, and B. Yıldız, "Towards the estimation of ESG ratings: A machine learning approach using balance sheet ratios," *Procedia Comput. Sci.*, vol. 242, pp. 1037–1046, 2024. DOI: [10.1016/j.ribaf.2024.102653](https://doi.org/10.1016/j.ribaf.2024.102653)
- [18] F. Rossi, T. Almeida, and L. Chen, "Predicting ESG scores using machine learning for data-driven sustainable investment," *AI Sustain.*, vol. 5, no. 1, p. 7, 2026. DOI: [10.3390/analytics5010007](https://doi.org/10.3390/analytics5010007)
- [19] Y. Liu, W. Wang, J. Tang, and Y. Du, "An optimized machine learning framework for predicting and interpreting corporate ESG greenwashing behavior," *PLoS ONE*, vol. 20, no. 3, p. e0316287, 2025. DOI: [10.1371/journal.pone.0316287](https://doi.org/10.1371/journal.pone.0316287)
- [20] J. Lee and K. Park, "Predicting ESG ratings by machine learning and analyzing influencing factors by XAI," in *Proc. 7th Int. Conf. Softw. Eng. Inf. Manag. (ICSIM)*, 2024, pp. 156–163. DOI: [10.1145/3647722.3647742](https://doi.org/10.1145/3647722.3647742)
- [21] T. Taskin, B. Yıldız, and S. Demir, "Explainable artificial intelligence for determining integrated reporting quality prediction factors," *J. Bus. Econ. Manag.*, vol. 26, no. 1, pp. 78–97, 2025. DOI: [10.1007/s10799-026-00482-9](https://doi.org/10.1007/s10799-026-00482-9)
- [22] N. Gui, D. Ge, and Z. Hu, "AFS: An attention-based mechanism for supervised feature selection," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 3705–3713. DOI: [10.1609/aaai.v33i01.33013705](https://doi.org/10.1609/aaai.v33i01.33013705)
- [23] M. Wojtas and K. Chen, "Feature importance ranking for deep learning," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 5105–5114. DOI: [10.48550/arXiv.2010.08973](https://doi.org/10.48550/arXiv.2010.08973)
- [24] J. Zhang, S. Patel, M. Wang, K. Choi, and Y. Chen, "Unleashing the power of multi-task learning: A comprehensive survey spanning traditional, deep, and pretrained foundation model eras," *arXiv preprint arXiv:2404.18961*, 2024. DOI: [10.48550/arXiv.2404.18961](https://doi.org/10.48550/arXiv.2404.18961)
- [25] A. M. Basahel, S. H. Giriappa, F. Alam, T. S. M. Alnazzawi, S. Qamar, and A. A. Abi Sen, "A resource-efficient approach to fine-tuning a BERT-base model for sentiment analysis," *Computers*, vol. 15, no. 3, p. 159, 2026, doi: [10.3390/computers15030159](https://doi.org/10.3390/computers15030159). DOI: [10.3390/computers15030159](https://doi.org/10.3390/computers15030159)
- [26] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2022. DOI: [10.48550/arXiv.2106.09685](https://doi.org/10.48550/arXiv.2106.09685)
- [27] Q. Zhang, M. Chen, A. Bukharin, P. He, Y. Cheng, W. Chen, and T. Zhao, "AdaLoRA: Adaptive budget allocation for parameter-efficient fine-tuning," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2023. DOI: [10.48550/arXiv.2303.10512](https://doi.org/10.48550/arXiv.2303.10512)
- [28] T. Yasuda, K. Yang, M. Zhang, F. Najar, and V. Mirrokni, "Sequential attention for feature selection," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2023. DOI: [10.48550/arXiv.2209.14881](https://doi.org/10.48550/arXiv.2209.14881)
- [29] I. Lopes, T. Vu, and R. de Charette, "Cross-task attention mechanism for dense multi-task learning," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, 2023, pp. 2329–2338. DOI: [10.1109/WACV56688.2023.00236](https://doi.org/10.1109/WACV56688.2023.00236)
- [30] B. P. P. Kumar and M. Ahmed, "IoT-based battery management system in the electric vehicle charging station," *Digital Twin, Blockchain, and Sensor Networks in the Healthy and Mobile City*, pp. 431–451, 2025, doi: <https://doi.org/10.1016/b978-0-443-34174-8.00022-3>.
- [31] B. P. P. Kumar, M. Tabassum, R. Neha, and A. Khanum, "ChargeSmart: Effortless EV Charging with Real-Time Slot Booking," 2025 International Conference on Knowledge Engineering and Communication Systems (ICKECS), pp. 1–6, Apr. 2025, doi: <https://doi.org/10.1109/ickecs65700.2025.11035724>.
- [32] B. P. P. Kumar, L. G. Reddy, M. Chethana, and D. Likhita, "Improving DDos Attack Identification and Mitigation In SDN Through The Use of Ensemble Online Machine Learning Framework," 2025 International Conference on Knowledge Engineering and Communication Systems (ICKECS), pp. 1–6, Apr. 2025, doi: <https://doi.org/10.1109/ickecs65700.2025.11035808>
- [33] I. S. S. S. Itannavar, J. S. Jakati, B. A. P. Khot, M. S. and B. P. P. Kumar, "Robustness-Oriented Explainability Analysis of the PHQ-9 Depression Scale Using Ensemble Learning Under Response Perturbations," 2026 4th International Conference on Knowledge Engineering and Communication Systems (ICKECS), pp. 1–6, Apr. 2026, doi: <https://doi.org/10.1109/ickecs70176.2026.11527572>.