ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

FedHENet: A Frugal Federated Learning Framework for Heterogeneous Environments

Alejandro Dopico-Castro, Oscar Fontenla-Romero, Bertha Guijarro-Berdiñas, Amparo Alonso-Betanzos and Iván Pérez Digón *

Universidade da Coruña, CITIC, Facultade de Informática
Campus de Elviña s/n, A Coruña, Spain

Abstract. Federated Learning (FL) enables collaborative training without centralizing data, essential for privacy compliance in real-world scenarios involving sensitive visual information. Most FL approaches rely on expensive, iterative deep network optimization, which still risks privacy via shared gradients. In this work, we propose FedHENet, extending the FedHEONN framework to image classification. By using a fixed, pre-trained feature extractor and learning only a single output layer, we avoid costly local fine-tuning. This layer is learned by analytically aggregating client knowledge in a single round of communication using homomorphic encryption (HE). Experiments show that FedHENet achieves competitive accuracy compared to iterative FL baselines while demonstrating superior stability performance and up to 70% better energy efficiency. Crucially, our method is hyperparameter-free, removing the carbon footprint associated with hyperparameter tuning in standard FL. Code available in `https://github.com/AlejandroDopico2/FedHENet/`

1 Introduction

Federated learning (FL) has emerged as a crucial paradigm for collaborative model training without sharing sensitive data. However, foundational algorithms like FedAvg [1] or FedProx [2] rely on the iterative optimization of deep neural networks. This leads to three major drawbacks: (1) poor convergence and client drift in heterogeneous scenarios; (2) high communication and computation costs over hundreds of rounds; and (3) vulnerability to privacy attacks via shared gradients.

To address these limitations, we propose FedHENet, a framework designed for "Frugal FL" in computer vision. While FedHEONN [3] validated single-round, homomorphically encrypted (HE) learning for tabular data, extending this to high-dimensional image domains is non-trivial. FedHENet solves this by decoupling feature extraction (via a pre-trained, frozen backbone) from the learning process. Instead of costly local fine-tuning, clients compute a lightweight, analytically solvable output layer (ROLANN). This ensures mathematically exact

*Work funded by Project PID2023-147404OB-I00 (MICIU/AEI/10.13039/501100011033; ERDF/EU; ESF+/EU), Horizon Europe (GA 101070381), and the Ministry for Digital Transformation and Civil Service and Next-GenerationEU/PRTR (TSI-100925-2023-1). CITIC, as a member of the CIGUS Network, receives subsidies from the "Xunta de Galicia" and from the ERDF Operational Programme Galicia 2021-2027 (Grant ED431G 2023/01).

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

global aggregation in a single round, eliminating the need for hyperparameter tuning and drastically reducing the environmental footprint.

Our main contributions are:

- **Frugal adaptation to image data:** A hybrid architecture combining frozen feature extractors with an analytical layer whose transmitted statistics are homomorphically encrypted, enabling secure single-round convergence on image tasks without iterative gradient descent.
- **Superior stability and efficiency:** We demonstrate robustness against extreme non-IID data (achieving high accuracy where iterative baselines fail) with up to 70% energy savings.
- **Scalable deployment:** We integrate MQTT for fault-tolerant communication and provide an open-source library to facilitate encrypted Frugal FL deployment.

2 Method

FedHENet combines a fixed, pre-trained (on ImageNet) feature extractor with a Regularized One-Layer Neural Network (ROLANN) classifier [4]. All clients share the same frozen feature extractor to produce compact embeddings $X_k \in \mathbb{R}^{m_k \times n}$ from their local m_k samples. By not performing any fine-tuning, we avoid high computational and energy costs.

On top of these features, each client trains a ROLANN layer. Unlike traditional FL, this layer computes its optimal weights in a closed form, not via iterative gradient descent. This is achieved by minimizing the MSE measured *before* the activation function on the desired pre-activation outputs d_k, which allows to transform the regularized least-squares problem into a system of linear equations that can be solved analytically, enabling single-round aggregation.

Each client k calculates the required components. The activation derivative f' is calculated on the desired pre-activaction outputs $\bar{d}_k$. Client k computes f', forms a diagonal matrix $F_k = \text{diag}(f'_k)$, and performs:

$$[U_k, S_k, \sim] = \text{SVD}(X_k F_k), \quad M_k = X_k(f'_k \odot f'_k \odot \bar{d}_k), \tag{1}$$

The only component containing sensitive statistical dependencies from the training data is the matrix M_k, defined as a weighted correlation between features and pre-activation terms. To protect M_k, each client encrypts it using the CKKS homomorphic encryption [5], producing $[[M_k]]$. CKKS is used because it enables arithmetic operations on approximate (real) numbers, essential for this analytical aggregation.

Clients send $(U_k, S_k, [[M_k]])$ to the coordinator using MQTT, which provides lightweight, asynchronous, and fault-tolerant communication. U_k and S_k are transmitted in plaintext and concatenated at the server to compute the global singular values (U, S). The matrices M_k are sent encrypted and their statistics are aggregated homomorphically:

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

$$[U, S, \sim] = \text{SVD}([U_1, S_1 \| U_2, S_2 \| \ldots \| U_K, S_K), \quad [[M]] = \sum_{k=1}^{K} [[M_k]] \tag{2}$$

The coordinator computes the global ROLANN weights W from U, S and $[[M]]$ in a single analytical step and publishes W to all clients using MQTT. This avoids redundant local computation, since all clients would obtain the same weights.

$$W = U \cdot (S \cdot S + \lambda I)^{-1} U^T [[M]] \tag{3}$$

This closed-form expression yields the globally optimal classifier without iterative refinement and ensures exact aggregation across heterogeneous clients.

3 Experimental Results and Discussion

Experimental Setup We implemented the framework[1] in PyTorch and executed on a workstation (Intel Core i7-10700KF, NVIDIA RTX 3080). Communication relied on MQTT (Eclipse Mosquitto) and CKKS Homomorphic encryption. Experiments used CIFAR-10 and CIFAR-100 datasets, containing 10 and 100 classes of 32×32 color images following the standard train and test splits, respectively. Data heterogeneity was simulated using a Dirichlet distribution (α) where lower α signifies non-IID data. Additionally, a "single-class" scenario (where each client only gets data for one class) was used to represent extreme heterogeneity. We compare FedHENet against FedAvg and FedProx using the same frozen ResNet-18 backbone. Baselines were trained for 10 global rounds on CIFAR-10 and 50 global rounds on CIFAR-100 with 1 local epoch per round and full client participation. These round counts were empirically determined to ensure the iterative algorithms reached a stable, near-optimal convergence for a fair performance comparison. Metrics include test accuracy, energy consumption (measured via CodeCarbon[2]), training time, and cumulative communication volume (the total data transmitted by all participants across the entire process).

Accuracy and Robustness Analysis Table 1 highlights the test results. While iterative methods (FedAvg and FedProx) perform well on CIFAR-10 in near-IID settings ($\alpha = 1.0$) with few client ($N = 10$), their performance deteriorates as data heterogeneity increases. In the extreme "single class" scenario ($N = 10$), baselines drop to $33 - 40\%$ accuracy due to severe client drift. FedHENet, relying on exact analytical aggregation, remains immune to drift, maintaining $\sim 83.68\%$ accuracy. In the large-scale experiment with $N = 500$ clients, our approach outperforms FedAvg by over 8% in the non-IID setting ($\alpha = 0.1$).

On the complex CIFAR-100 dataset, FedHENet achieves competitive results ($\sim 56\%$) against baselines that needed 50 rounds to achieve stable convergence.

[1] `https://github.com/AlejandroDopico2/FedHENet/`

[2] `https://github.com/mlco2/codecarbon`

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Dataset	N	α	FedAvg	FedProx	**FedHENet**
CIFAR-10	10	1.0	**85.34**	85.28	83.69
		0.1	81.25	78.52	**83.68**
		single	33.26	40.75	**83.65**
	100	1.0	79.82	79.84	**83.64**
		0.1	76.44	76.89	**83.66**
		single	48.12	51.19	**83.63**
	500	1.0	73.78	73.78	**83.79**
		0.1	75.18	75.18	**83.59**
		single	73.26	73.28	**83.70**
CIFAR-100	100	1.0	**59.11**	59.08	56.63
		0.1	**58.2**	58.15	56.13
		single	52.61	52.6	**56.91**

Table 1: Top-1 test accuracy (%) comparison between FedAvg, FedProx, and FedHENet, under a varying number of clients (N) and data heterogeneity (α).

In highly heterogeneous settings, FedHENet matches or beats the baselines, which suffer from divergence despite the extended training.

Frugal Efficiency and Sustainability Table 2 confirms massive resource reductions. For CIFAR-10 ($N = 10$), FedHENet reduces energy by $\sim 70\%$ (11.5 Wh vs 36.5 Wh). Although FedHENet's single payload is larger than a standard gradient update, the overall byte count is significantly lower than the baselines. This gap widens for CIFAR-100 where baselines consumed ~ 240 Wh (due to the 50-round requirement), while FedHENet required only 118.9 Wh, halving the cost.

A crucial, often overlooked aspect of sustainability is the cost of hyperparameter tuning. Iterative methods require energy-intensive grid searches (LR, decay schedules, and local epochs) to prevent divergence. This search phase multiplies the true energy cost of deployment. FedHENet is hyperparameter-free, eliminating this search phase and drastically reducing the carbon footprint of the entire development lifecycle, not just the final training run.

Accuracy-Energy Trade-off The plot in Figure 1, illustrating Accuracy vs. Energy consumption per round, confirms FedHENet's frugality. The iterative baselines are forced to continue consuming energy over multiple rounds to asymptotically approach FedHENet's single-round accuracy. For example, in CIFAR-10, FedHENet achieves $\sim 83.6\%$ accuracy using only ~ 11.5 Wh of energy, while iterative methods consume over three times the energy (~ 40 Wh) to converge. Furthermore, their performance in the non-IID setting ($\alpha = 0.1$) is highly unstable, oscillating wildly between 45% and 77% accuracy across rounds, demonstrating low robustness despite high energy expenditure. This frugality is magnified on

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Dataset	N	Metric	FedAvg	FedProx	**FedHENet**
CIFAR-10	10	Time (min.)	8.55	10.11	**3.07**
		Energy (Wh)	30.2	36.5	**11.5**
		Bytes (MB)	358.03	358.04	**331.33**
	100	Time (min.)	12.15	14.37	**3.77**
		Energy (Wh)	40.6	47.0	**11.2**
		Bytes (MB)	3369.84	3369.85	**2855.51**
	500	Time (min.)	32.80	32.72	**12.32**
		Energy (Wh)	96.85	95.07	**36.01**
		Bytes (MB)	16755.52	16755.50	**9694.33**
CIFAR-100	100	Time (min.)	77.71	81.98	**40.32**
		Energy (Wh)	236.01	248.93	**118.9**
		Bytes (MB)	42121.21	42121.21	**29615.90**

Table 2: Energy, communication and time efficiency comparison ('Green metrics') of FedAvg, FedProx, and FedHENet for CIFAR datasets.

CIFAR-100, where the single-round execution provided an energy reduction of nearly half compared to baselines that required extensive, multi-round training.

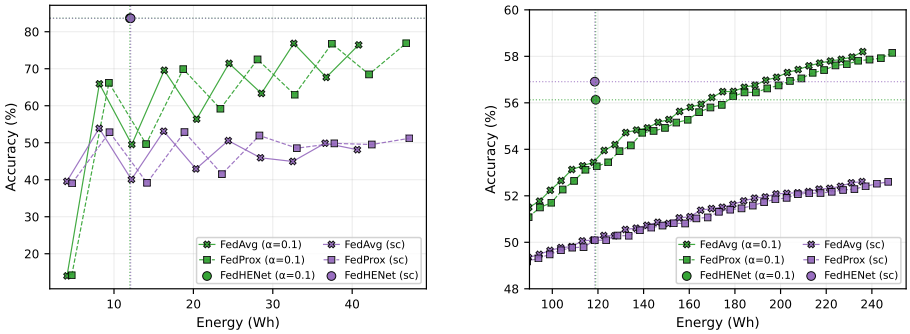

Fig. 1: Accuracy vs. energy consumption per training round for CIFAR-10 (left) and CIFAR-100 (right) for $\alpha = 0.1$ and single-class (sc) scenarios with 100 clients. FedHENet achieves peak accuracy in a single round, while baselines require multiple rounds. In CIFAR-10, the two circles overlap. For clarity, in CIFAR-100 only the last 30 baseline rounds are shown.

Homomorphic Encryption Overhead Encryption increases the payload size (M_k matrix) by approximately 2.25×, increasing transmitted bytes (Table 3). However, the computational time overhead is negligible ($< 5\%$). This confirms that FedHENet provides robust privacy via without compromising its frugal nature. While this overhead remains negligible for a moderate number of clients, it could

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

FedHENet	Acc. (%)	Bytes (MB)	CT size inflation	Energy (Wh)	Time (min.)	Overhead (ms/client)
w/o HE	83.69	147.6	1.0×	10.60	2.92	–
w/ HE	83.69	331.33	2.25×	10.93	3.06	840

Table 3: Impact of applying Homomorphic Encryption (HE) to FedHENet on CIFAR-10 (α=1.0, N=10).

become more noticeable in massive-scale federations, where encryption and communication costs scale linearly with the client count.

4 Conclusions

We introduced FedHENet, a framework that addresses resource constraints and heterogeneity via a computationally frugal, single-round design. By decoupling feature extraction from a closed-form learning layer, it achieves up to 70% energy savings compared to iterative baselines. Our results demonstrate that, unlike gradient-based methods, FedHENet avoids client drift in extreme non-IID settings. Furthermore, by being hyperparameter-free, it offers a sustainable path for edge AI, eliminating the energy-intensive tuning phase required by standard approaches. The production-ready core implementation of this framework, which is designed for seamless deployment in a client network, is available as part of a general-purpose federated learning library repository.

For future work, we plan to validate FedHENet on heterogeneous edge hardware (e.g., Raspberry Pi clusters) to precisely measure real-world performance metrics. We will also explore the use of other foundational backbones, such as Vision Transformers.

References

[1] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.

[2] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.

[3] Oscar Fontenla-Romero, Bertha Guijarro-Berdiñas, Elena Hernández-Pereira, and Beatriz Pérez-Sánchez. Fedheonn: Federated and homomorphically encrypted learning method for one-layer neural networks. *Future Generation Computer Systems*, 149:200–211, 2023.

[4] Oscar Fontenla-Romero, Bertha Guijarro-Berdiñas, and Beatriz Pérez-Sánchez. Regularized one-layer neural networks for distributed and incremental environments. In *Advances in Computational Intelligence: 16th International Work-Conference on Artificial Neural Networks, Part II*, page 343–355. Springer-Verlag, 2021.

[5] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic encryption for arithmetic of approximate numbers. In *International conference on the theory and application of cryptology and information security*, pages 409–437. Springer, 2017.