EA-DFL: Balancing Consensus and Energy in Decentralized Federated Learning for e-Health Wearables

Hadi Mekdad
MIS Laboratory
UPJV
Amiens, France
hadi.mekdad@u-picardie.fr

Wafa Badreddine
MIS Laboratory
UPJV
Amiens, France
wafa.badreddine@u-picardie.fr

Jérôme Bosche
MIS Laboratory
UPJV
Amiens, France
jerome.bosche@u-picardie.fr

***Abstract*—Decentralized Federated Learning removes the central aggregation server from collaborative training, which is crucial in clinical settings. Deploying it on battery-constrained wearable sink nodes is limited, however, by the radio energy spent broadcasting model parameters over low-bandwidth links. We propose EA-DFL, an energy-aware decentralized federated learning algorithm built on probabilistic gossip. Each node adjusts its broadcast probability from a localized urgency score combining residual battery, local learning progress and topological node density, throttling transmissions once the update is not worth spending energy on. Coupling this rule with an energy model covering computation and radio costs, we simulate EA-DFL on three non-IID medical datasets against unthrottled, probabilistic and periodic decentralized stochastic gradient descent. EA-DFL reaches the reference accuracy levels of 80%, 85% and 90% used for these tasks, cuts the mean reception burden on MIT-BIH from 432 to 193 messages per node, and is the only configuration in which all nodes survive on UCI-HAR.**

***Index Terms*—Decentralized federated learning, wearable devices, energy efficiency, probabilistic gossip, WBAN, TinyML, e-Health**

I. Introduction

Wearable edge devices generate continuous physiological data that machine learning can turn into predictive healthcare. Federated Learning trains such models without sending sensitive patient records to the cloud, sharing only model updates, but it depends on an aggregation server that is both a single point of failure and a high-value target. Decentralized Federated Learning (DFL) removes that orchestrator and lets edge devices aggregate directly with their neighbors [1]. Although DFL may still be susceptible to attacks such as gradient leakage and model inversion by neighbors, it provides better privacy by eliminating the central server.

Ubiquitous monitoring requires pushing this intelligence onto Wireless Body Area Networks (WBANs), where a sink node such as a smartwatch aggregates the physiological data gathered by sensors worn by one patient [2]. Sink nodes carry small batteries while DFL mandates continuous model broadcasting, and the resulting radio drain limits how long decentralized algorithms can run. When a node depletes its battery mid-training the topology fractures, survivor nodes lose the skewed distributions of the dead nodes, and convergence degrades.

Prior work has addressed this communication burden from several directions. Roy *et al.* [3] introduced BrainTorrent, a peer-to-peer architecture for medical DFL, and Zhu *et al.* [4] studied DFL over device-to-device wireless links with scheduling protocols that manage channel interference. Hegedűs *et al.* [5] treated gossip learning through peer-to-peer random walks as a scalable alternative to federated averaging. Others shrink the payload rather than the schedule: Hu *et al.* [6] proposed Combo, a segmented gossip scheme splitting models into distributable fragments, and Koloskova *et al.* [7] introduced Choco-SGD, which tolerates arbitrary compression of the exchanged updates. Closest to our setting, Han *et al.* [8] considered DFL on energy-harvesting devices. These approaches either assume a stable energy supply or reduce how much is sent; none conditions the decision to transmit on the real-time hardware state of the node, in particular its residual battery. EA-DFL addresses this gap with probabilistic gossip driven by an urgency score over remaining energy, learning progress and topological density, so each node decides locally when a broadcast is worthwhile.

II. Proposed Algorithm: Energy-Aware DFL

A. Preliminaries and Energy Constraints

We consider N wearable sink devices $\mathcal{V} = \{1, \dots, N\}$ forming an undirected communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ under a finite transmission range R_{TX}. Device i holds a private non-IID dataset $\mathcal{D}_i$ of D_i samples and a local empirical loss $f_i(\boldsymbol{\theta}) = \frac{1}{D_i} \sum_{\xi \in \mathcal{D}_i} l(\boldsymbol{\theta}; \xi)$. The network learns shared parameters $\boldsymbol{\theta} \in \mathbb{R}^d$ minimizing the data-weighted global objective

$$F(\boldsymbol{\theta}) = \frac{1}{D} \sum_{i=1}^{N} D_i f_i(\boldsymbol{\theta}), \qquad D = \sum_{i=1}^{N} D_i. \tag{1}$$

Under Decentralized Stochastic Gradient Descent (DSGD), at round r node i aggregates the parameters of its active

neighbors $\mathcal{N}_i^r$ through a doubly stochastic matrix $\mathbf{W}^r$ and applies a local gradient step,

$$\boldsymbol{\theta}_i^{r+1} = \sum_{j \in \mathcal{N}_i^r} \mathbf{W}_{ij}^r \boldsymbol{\theta}_j^r - \eta \nabla f_i(\boldsymbol{\theta}_i^r). \tag{2}$$

Run continuously under a finite energy budget E_{max}, this drains a wearable quickly. Computation costs $E_{comp}(M) = M \cdot E_{MAC}$ for M being number of multiply-accumulate operations [9], but the radio dominates: under the First-Order Radio Model [10], sending k bits over distance d costs

$$E_{TX}(k, d) = E_{elec} \cdot k + \epsilon_{amp} \cdot k \cdot d^2, \tag{3}$$

while reception costs $E_{RX}(k) = E_{elec} \cdot k$. A broadcast is charged once per transmission event at the worst-case range $d = R_{TX}$, and E_{RX} once per receiving neighbor, so at the mean degrees of Table I a node spends more on reception than on its own transmissions and dense nodes are reception-limited. The challenge is to sparsify $\mathcal{G}$ enough for nodes to survive without fracturing the network.

B. The EA-DFL Algorithm

Because each patient wears exactly one sink device, every node represents both a wearable and its human subject. Over a shared fixed topology, EA-DFL replaces continuous broadcasting with a probabilistic gossip mechanism driven by a localized urgency score. The protocol, formalized in Algorithm 1, runs in three phases.

Phase 1 (Initialization). Nodes discover neighbors within R_{TX}; reserves $E_i = E_{max}$ and probabilities $p_i = 1.0$ are set.

Phase 2 (Local training and urgency evaluation). Nodes run local mini-batch gradient descent, then node i evaluates

$$u_i = \alpha \hat{E}_i + \beta \Delta L_i + \gamma \rho_i. \tag{4}$$

Here $\hat{E}_i = E_i / E_{max} \in [0, 1]$ is normalized residual energy and $\rho_i = |\mathcal{N}_i|/(N-1) \in [0, 1]$ is node degree over the maximum possible degree. The learning term $\Delta L_i = \max(0,\, L_i^{(t-1)} - L_i^{(t)})$ is the drop between the losses of the two most recent local mini-batch updates, read at the communication step of the round. It is rectified, so a rising loss counts as stagnation. If u_i falls to or below a localized threshold τ_i, the node decays p_i.

Phase 3 (Probabilistic gossip and aggregation). Nodes activate the radio with probability p_i. When a broadcast begins, E_{TX} is deducted and neighbors aggregate the incoming parameters. Following the shared-representation principle [11], aggregation touches only the feature extraction layers, so personalized classification heads stay strictly local, and E_{RX} is deducted correspondingly. A silent node conserves radio energy and still progresses through its local gradients and earlier aggregations.

Following the Multi-Factor Weighting Function proposed for IoT [12], urgency is a linear sum, which keeps $\mathcal{O}(1)$ complexity so constrained microcontrollers spend no energy on non-linear functions. At initialization the node is at full

Algorithm 1 Energy-Aware DFL (EA-DFL)

Parameters: Nodes N, range R_{TX}, epochs E, rate η, energy E_{max}, costs E_{MAC}, E_{elec}, ϵ_{amp}, weights α, β, γ, decay rule, floor p_{min}, buffer b.

Initialize: node i: $\boldsymbol{\theta}_i = [\boldsymbol{\omega}_i; \boldsymbol{\phi}_i]$, $E_i = E_{max}$, $p_i = 1.0$, $\tau_i = \alpha + \gamma\rho_i - b$.
while global convergence not met **do**
  for each node i **do**
    Compute $\nabla_{\boldsymbol{\omega}} f_i$, $\nabla_{\boldsymbol{\phi}} f_i$; deduct E_{comp} from E_i
  end for
  if epoch $\pmod E == 0$ **then**
    for each node i **do**
      $u_i \leftarrow \alpha \hat{E}_i + \beta \Delta L_i + \gamma \rho_i$
      if $u_i \le \tau_i$ **then**
        decay p_i, floored at p_{min}; $\tau_i \leftarrow u_i - b$
      end if
      Broadcast $\boldsymbol{\omega}_i$ to $\mathcal{N}_i$ w.p. p_i; deduct E_{TX} if sent
    end for
    for each node i **do**
      $\boldsymbol{\omega}_i \leftarrow \sum_{j \in \mathcal{N}_i} \mathbf{W}_{ij} \boldsymbol{\omega}_j - \eta \nabla_{\boldsymbol{\omega}} f_i$; $\boldsymbol{\phi}_i \leftarrow \boldsymbol{\phi}_i - \eta \nabla_{\boldsymbol{\phi}} f_i$
      Deduct E_{RX} from E_i *// head $\boldsymbol{\phi}_i$ stays local*
    end for
  end if
end while

energy ($\hat{E}_i = 1.0$) with no loss history, so $\Delta L_i = 0$ and the baseline urgency and initial threshold are

$$U_{base} = \alpha + \gamma\rho_i, \quad \tau_i = U_{base} - b. \tag{5}$$

We use $\alpha = 0.3$, $\beta = 0.3$, $\gamma = 0.4$ and $b = 0.03$. Setting $\alpha = \beta$ gives an even trade-off between survival and learning utility, and against $\alpha = 0.3$ the buffer $b = 0.03$ corresponds to a 10% margin of battery capacity a stagnated node may spend before throttling begins. On each decay the threshold is re-anchored to $\tau_i \leftarrow u_i - b$, enforcing a monotonic step-down. These values were fixed and not tuned; Section IV returns to this.

Expanding the trigger $u_i \le \tau_i$ gives $\alpha \hat{E}_i + \beta \Delta L_i + \gamma \rho_i \le \alpha + \gamma \rho_i - b$, so that

$$0.3\hat{E}_i + 0.3\Delta L_i \le 0.27 \iff \hat{E}_i \le 0.90 - \Delta L_i. \tag{6}$$

A stagnated node ($\Delta L_i \approx 0$) throttles at 90% battery, whereas one still improving ($\Delta L_i = 0.05$) may drain to 85% first, so the threshold preserves the broadcasts of nodes that are still learning.

Remark 1 (the density term is inert on a static graph). Equation (6) follows by cancelling $\gamma\rho_i$ from both sides, and the cancellation recurs after every decay: with $\tau_i = u_i^{(r)} - b$. When ρ_i is time invariant, γ has no influence on when a node throttles and only shifts the absolute level of u_i and τ_i. We therefore make no claim that γ protects bridge nodes or prevents node death below. We retain it because it becomes active once ρ_i changes between the round anchoring τ_i and the round evaluating u_i, as under a time-varying topology, and

TABLE I
SIMULATION AND ENERGY MODEL PARAMETERS

Parameter	PathMNIST	MIT-BIH	UCI-HAR
Learning rate	0.001	0.001	0.05
Accuracy target	80%	85%	90%
Decay rule	Linear 0.2	Halving $\times 0.5$	Linear 0.2
Floor p_{min}	0.25	0.125	0.1
Eval. interval	1 epoch	1 epoch	2 epochs
Classes	9	5	6
Train / test	89,996 / 7,180	80/20 (cap 12k)	80/20
Model	Light CNN	1D CNN	MLP
Parameters	106,064	23,653	80,582
Mean degree	9.4	9.2	9.6
Payload k (bits)	3.39 M	640 k	2.58 M
E_{elec} (nJ/bit)	50	500	50
ϵ_{amp} (pJ/bit/m^2)	100	5000	100
E_{MAC} (nJ/MAC)	0.15	1.5	0.15
MACs per sample	318,192	60,000	241,746
Random seed	43	45	42

[a] MIT-BIH payload set from a nominal 20,000-parameter model.

because a degree-scaled buffer $b(\rho_i)$ would activate it on static graphs. In other words, it would prove useful in future work on dynamic topology.

III. Performance Evaluation

A. Experimental Setup

We evaluate EA-DFL on three datasets spanning different modalities and non-IID skews: PathMNIST [13], colon pathology images in 9 classes; MIT-BIH Arrhythmia [14], electrocardiogram signals from wearable sensors; and UCI-HAR [15], accelerometer and gyroscope signals over 6 daily activities.

Every environment models 30 identical nodes deployed uniformly in a 50×50 m area, linked within $R_{TX} = 20$ m and initialized to $E_{max} = 300$ J. Traditional DSGD, Probabilistic DSGD ($p = 0.5$), Periodic DSGD ($K \in \{2, 3\}$) and EA-DFL ran on the same random geometric graph per dataset, with the same partition, the same optimizer (Adam, weight decay 10^{-4}) and the same shared-extractor and local-head split, so accuracy differences are not attributable to personalization. PathMNIST and MIT-BIH used a Dirichlet partition ($\alpha = 5.0$); UCI-HAR gave each node one human subject. Table I lists the remaining parameters and the constants used in (3). Each configuration was run once, from a single seed and one graph realization, so differences of a few rounds should not be read as significant. A node reaching zero energy is marked dead and permanently excluded from transmitting, receiving and aggregating; the graph is not re-planned, so the surviving topology can fracture.

B. Accuracy and Convergence

Mean accuracy is read together with milestone rounds, the round at which a given percentage of nodes reaches target, since a method can look stable on average while leaving part of the network in local minima.

On PathMNIST, Traditional DSGD reaches the 95% milestone in 63 rounds and EA-DFL in 75, while Probabilistic DSGD ($p = 0.5$) stalls at the 90% milestone. Traditional DSGD and EA-DFL hold accuracy standard deviation in a 1% to 2% band, whereas the periodic methods oscillate between 4% and 10% late into training.

On MIT-BIH the 90% condition is met in 43 rounds by EA-DFL, 46 by Probabilistic DSGD and 47 by Traditional DSGD. These fall within four rounds of each other in a single run, so we read them as comparable rather than as evidence that EA-DFL converges faster; We can claim that throttling did not slow convergence.

UCI-HAR is where the communication policy matters most. Traditional DSGD plateaus near 85% and flattens as nodes exhaust their batteries, freezing local accuracies, so it never clears the 80% condition. Probabilistic DSGD reaches 90% after 313 rounds and neither the 95% nor the 100% condition. EA-DFL keeps the topology intact, reaches 80% in 104 rounds, and is the only configuration bringing 100% of the network to target, in 191 rounds.

C. Energy and Network Lifespan

Traditional DSGD bounds expenditure from above (Fig. 1), since every active node broadcasts its full payload at every opportunity. On PathMNIST it consumes close to 130 J, over 40% of capacity, in 60 rounds, much of it after local models have stabilized; EA-DFL reaches the same milestone on roughly 110 J by round 75.

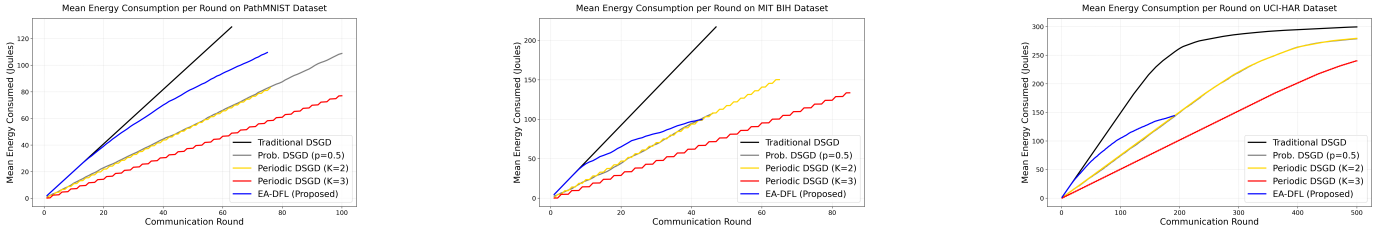
On MIT-BIH, Traditional DSGD consumes over 200 J. Rather than halving communication uniformly, EA-DFL transmits freely during early feature exploration, so its curve initially tracks above the static methods; as the local loss decrease shrinks it steps the probabilities down and tapers off, finishing on close to 100 J.

On UCI-HAR, baselines were run to 500 rounds to observe longevity. Under Traditional DSGD the curve approaches the 300 J ceiling then flattens near round 200 because nodes are dying: by round 500, 28 of 30 were depleted. The static methods delay this, with 19 deaths for Probabilistic DSGD and $K = 2$ and 5 for $K = 3$, and this fracturing coincides with the accuracy plateau of Section III-B. Because $\hat{E}$ enters the throttling condition, EA-DFL grows more conservative as reserves fall and reaches 100% convergence on 145 J with no node death in this run.

D. Communication Volume

Energy use is driven by the number of messages the radio handles, and dense nodes are the most exposed. Converging on PathMNIST required Traditional DSGD nodes to process a mean of 592 incoming messages, which Periodic DSGD ($K = 2$) reduces to 357 at the cost of the instability noted above. On MIT-BIH, EA-DFL reduces the mean reception burden to 193 messages against 432 for Traditional DSGD.

Both comparisons are taken at each method's own milestone round, which favours methods that stop earlier and should be read with that caveat. For UCI-HAR, where several baselines never converge, volumes are instead snapshotted at the common round 191: Traditional DSGD had forced 1,624 received messages per node, static $p = 0.5$ and $K = 2$ roughly 900, and EA-DFL 907. EA-DFL matched the static methods in



(a) PathMNIST (b) MIT-BIH (c) UCI-HAR

Fig. 1. Mean energy consumed per communication round. Traditional DSGD drains the 300 J budget close to linearly and flattens on UCI-HAR as nodes die; EA-DFL tapers once local losses plateau.

volume while being the configuration in which the network reached target, which is consistent with, but not proof of, better message selection; confirming it would require the ablation we did not run.

IV. Conclusion, Limitations and Future Work

We presented EA-DFL, an energy-aware decentralized federated learning algorithm that throttles gossip broadcasts using a localized urgency score over residual battery, local learning progress and topological density. Each node evaluates the score from state it already holds, at $\mathcal{O}(1)$ cost and without any coordinator, so the mechanism stays within reach of a wearable microcontroller.

In simulation on three non-IID datasets, EA-DFL reached the accuracy levels of 80%, 85% and 90% that was reported for the activity recognition tasks [16] without slowing convergence. On MIT-BIH it met the final milestone in 43 rounds, comparable to the 46 and 47 of the static schemes, while cutting the mean reception burden from 432 to 193 messages per node. On UCI-HAR, where unthrottled broadcasting exhausted 28 of 30 batteries and left the network short of target, EA-DFL brought every node to target on 145 J without a single node death. Conditioning the broadcast decision on local hardware state therefore recovers much of the energy and traffic saving of a static communication cut while avoiding the loss of consensus it incurs.

These results carry clear limits. The evaluation is a simulation rather than a deployment, so the reported energy is modelled rather than measured and the link layer is idealized, leaving out the losses, contention and protocol overheads that weigh on real BLE and WBAN budgets. Each configuration was run once on a single static topology of homogeneous devices, so the differences of a few rounds reported above are indicative rather than significant. We report no ablation of the urgency weights, and we do not compare against compression-based decentralized methods [6], [7], which reduce how much is sent rather than whether to send and are complementary to EA-DFL.

Addressing these points is the subject of ongoing work, which moves beyond the fixed urgency rule used here and toward a less idealized network model.

References

[1] J. Muñoz et al., "Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges," Cyber-Defence Campus, 2024.
[2] R. A. Khan and A.-S. K. Pathan, "The state-of-the-art wireless body area sensor networks: A survey," *Int. J. Distrib. Sensor Netw.*, vol. 14, no. 4, 2018.
[3] S. Roy et al., "BrainTorrent: A peer-to-peer environment for decentralized federated learning," *arXiv:1905.05665*, 2019.
[4] L. Zhu, X. Jin, and Q. Li, "Decentralized federated learning via SGD over wireless D2D networks," *IEEE Trans. Wireless Commun.*, vol. 21, no. 2, pp. 1180–1193, 2021.
[5] I. Hegedűs, G. Doka, and M. Jelasity, "Gossip learning as a decentralized alternative to federated learning," *Distrib. Parallel Databases*, vol. 39, no. 1, pp. 73–114, 2021.
[6] C. Hu, J. Jiang, and Z. Wang, "Decentralized federated learning: A segmented gossip approach," in *Proc. IJCAI*, 2019, pp. 2616–2622.
[7] A. Koloskova, S. U. Stich, and M. Jaggi, "Decentralized deep learning with arbitrary communication compression," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2020.
[8] D. Han et al., "Decentralized federated learning with energy harvesting devices," *arXiv:2602.14051*, 2026.
[9] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proc. IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
[10] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proc. 33rd Annu. Hawaii Int. Conf. Syst. Sci.*, 2000.
[11] L. Collins, H. Hassani, A. Mokhtari, and S. Shakkottai, "Exploiting shared representations for personalized federated learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2021, pp. 2089–2099.
[12] L. Altoaimy et al., "Context-aware gossip-based protocol for internet of things applications," *Sensors*, vol. 18, no. 7, p. 2233, 2018.
[13] J. Yang et al., "MedMNIST v2: A large-scale lightweight benchmark for 2D and 3D biomedical image classification," *Sci. Data*, vol. 13, no. 1, p. 41, 2023.
[14] G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," *IEEE Eng. Med. Biol. Mag.*, vol. 20, no. 3, pp. 45–50, 2001.
[15] D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz, "A public domain dataset for human activity recognition using smartphones," in *Proc. ESANN*, 2013, pp. 437–442.
[16] Y. Ren, M. Liu, Y. Yang, L. Mao, and K. Chen, "Clinical human activity recognition based on a wearable patch of combined tri-axial ACC and ECG sensors," *Digital Health*, vol. 10, 2024.