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/.

On the Impact of Differential Privacy on Federated Neuromorphic Learning Accuracy

Luiz Pereira[1], Dalton Valadares[2], Mirko Perkusich[1], and Kyller Gorgônio[1] *

1- VIRTUS Research, Development and Innovation Center (UFCG)
Campina Grande, PB - Brazil

2- Federal University of Paraíba - Agricultural School Vidal de Negreiros
Bananeiras, PB - Brazil

Abstract. Federated Neuromorphic Learning (FNL) applies Spiking Neural Networks (SNNs) to enable energy-efficient collaborative learning on devices without centralizing data. However, integrating Differential Privacy (DP) introduces critical changes to the SNN firing dynamics, which propagate to server coordination strategies. This paper investigates DP-induced firing-rate distortions and their influence on global model convergence and generalization. Experimental ablation studies across privacy budgets and clipping bounds highlight firing distortions directly related to global accuracy degradation. Additionally, client selection instabilities related to DP noise degrade the model aggregation performance. The results reinforce that firing-rate-based FNL strategies are fragile under DP and require precise calibration to maintain the effectiveness of federated coordination.

1 Introduction

Federated Neuromorphic Learning (FNL) explores collaborative Federated Learning (FL) of on-device Spiking Neural Networks (SNNs). SNNs have emerged as a promising alternative to conventional Artificial Neural Networks (ANNs) because event-driven computation can substantially reduce power, latency, and memory traffic [1]. These features enable collaborative on-device training across resource-constrained clients [2].

In practical deployments, however, FL on user-generated data must satisfy additional privacy guarantees [3]. An established privacy mechanism in the literature is the Differentially Private Stochastic Gradient Descent (DP-SGD), where clipped per-example gradients and Gaussian noise are applied [4]. When DP-SGD is used for SNN training, such perturbations do not inject noise into spikes. However, they can still indirectly alter the parameters controlling synaptic spikes and thresholds [5], potentially distorting SNN firing rates. Recent studies have introduced firing-rate–aware coordination strategies for FNL [6, 7].

*This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. It has also been partially funded by the project "*NeuroMemIC Base: Desenvolvimento de Arquiteturas Neuromórficas e Near-Memory para Circuitos Integrados de Alta Eficiência*" supported by CENTRO DE COMPETÊNCIA EMBRAPII VIRTUS EM HARDWARE INTELIGENTE PARA INDÚSTRIA - VIRTUS-CC, with financial resources from the PPI HardwareBR of the MCTI grant number 055/2023, signed with EMBRAPII.

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/.

However, the impact of differential privacy mechanisms on these rate-based coordination dynamics remains insufficiently understood and lacks a more in-depth characterization.

In contrast to prior studies that primarily evaluated the sensitivity of federated SNNs to privacy noise from a coordination-dynamics and resource-efficiency perspective [8], this work investigates the DP-noise influence on training convergence and the stability of global model aggregation. We detail the LIF dynamics under stochastic settings and demonstrate that the bias introduced by DP-SGD perturbs firing rates, altering their statistical distribution, which in turn reshapes the trajectory of global accuracy during training. This establishes a direct relationship between DP-induced firing-rate distortions and global accuracy degradation. Furthermore, through ablations on an event-driven dataset, we quantify both how privacy budgets (ε, δ) and clipping C reshape firing-rate statistics and how they influence the global model accuracy and convergence.

2 LIF Neuron Dynamics under Stochastic Settings

Primary computational units of Spiking Neural Networks (SNNs) are commonly modeled as a Leaky Integrate-and-Fire (LIF) neuron model [9]. The LIF neuronal subthreshold dynamics in continuous time are given by:

$$\tau_m \frac{dV_t}{dt} = -\big(V_t - V_{\text{rest}}\big) + R\, I_t, \tag{1}$$

where τ_m is the membrane time constant, V_t is the membrane potential, V_{rest} is the resting potential, and $R\, I_t$ is the input current scaled by membrane resistance R. Let V_{th} denote the firing threshold and V_r the reset potential. A spike s is emitted whenever $V_{t^-} \geq V_{\text{th}}$. Upon a spike, the state is reset $V_{t^+} \leftarrow V_r$ and held during an absolute refractory period τ_{ref}.

Let $\theta_{i,j}$ be the synaptic weight connecting a pre-synaptic neuron i in layer $\ell - 1$ to neuron j in layer ℓ. The input current I_j to a neuron j can be expressed as follows:

$$I_j = \sum_{i}^{n_{\ell-1}} \theta_{i,j} \cdot s_i \tag{2}$$

Therefore, based on LIF dynamics (Equation 1), we can consider that the firing frequency of neuron j is a function of I_j.

In a deterministic case (without noise in I_j) and $V_{\text{rest}} = V_r$, if I_j is large enough to reach the threshold before complete decay, the firing frequency can be approximated by $f_j \approx [\tau_{\text{ref}} + \tau_m \ln(I_j/[I_j - I_{\text{th}}])]^{-1}$, where I_{th} is the minimum current required to reach the firing threshold [5]. Note that the firing frequency is inversely proportional to the total time required for the membrane potential to reach threshold, which in turn depends on how strong I_j is. Furthermore, f_j (and therefore ν_j) is a monotonically increasing function of the input current I_j. For $I_j \ll I_{th}$, ν grows slowly, approaching zero near threshold. For $I_j \gg I_{th}$, ν increases and converges asymptotically to $1/\tau_{\text{ref}}$.

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/.

The dynamics discussed are expected to be in stochastic settings when I_t contains an additive white-noise component ("noise-assisted" spiking is a classic result for IF models [5]). Based on Equation 2, note that a stochastic contribution to synaptic weight θ can cause the membrane potential to drift away from the reference trajectory. Figure 1 describes the behaviour of a LIF neuron in this stochastic case. Note that even if s is constant (or approximately constant), the neuron may fire stochastically due to variations in synaptic weights. Because the weights θ, in this case, follow a random trajectory due to the added noise, the membrane potentials fluctuate stochastically, and this results in ν exhibiting higher variance across training iterations compared to the deterministic case. Furthermore, if V_t is close to V_{th}, even the small weight perturbations can flip a neuron from silent ($s = 0$) to spiking ($s = 1$) or vice versa. This threshold sensitivity amplifies the variance in ν_j.

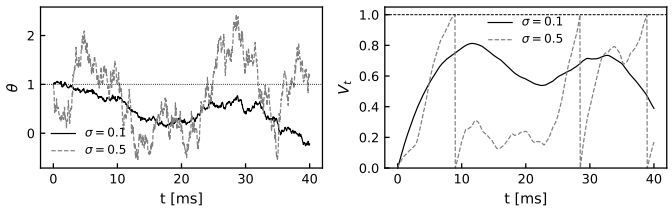

Fig. 1: LIF behavior under noised synaptic weight θ. The pre-synaptic spike remains constant $s = 1$ and $V_{\text{rest}} = V_r$. Right. θ under Gaussian noise with variance σ. Left. V_t under noise. We adopted $R = 1$, $\tau_m = 10$ ms and $V_r = 0$.

3 Privacy-Preserving FNL

Federated Neuromorphic Learning (FNL) explores collaborative FL for on-device neuromorphic computation models, such as Spiking Neural Networks (SNNs) [2]. A standard cross-device FNL setting comprises a central server and K resource-constrained edge clients. In each communication round r, the server selects a subset P of clients and broadcasts the current global parameters Θ_r. Each selected client k performs E local training steps on its private dataset D_k, producing a local update $\Theta_{k,r}$, which is then transmitted back to the server. The server aggregates the received client updates to compute the next global model parameters.

Additional privacy requirements are incorporated in FNL by the introduction of DP-based noisy gradients into training methods [3]. At example-level DP, each FNL client is noised during local training via Differentially Private Stochastic Gradient Descent (DP-SGD) [4]. Formally, in DP-SGD with per-example clipping at norm C and Gaussian noise multiplier σ, the noisy mini-batch gradient at local step t is defined by:

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/.

$$\tilde{\mathbf{g}}_t = \bar{\mathbf{g}}_t + \frac{\sigma C}{B}\,\boldsymbol{\xi}_t, \quad \boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{3}$$

where $\bar{\mathbf{g}}_t$ is the clipped average gradient measured by:$\bar{\mathbf{g}}_t = 1/B \sum_b \mathrm{clip}_C(\mathbf{g}_{t,b})$.

To improve FNL coordination, strategies were developed that are based on the firing rate. We can calculate the empirical firing rate of a neuron from its spike-emission behaviour [8]. Let $s_{b,j,t} \in \{0, 1\}$ be the spike of neuron j at time t to example b. For a time window of length T and batch of size B, we can calculate the per-neuron firing rate by $\nu_j = 1/B\,T \sum_b \sum_t s_{b,j,t}$. We can derive a neuron-weighted layer rate ν^ℓ for a layer ℓ, and a network-wide rate via across layers, i.e., $\nu_k = 1/n \sum_\ell \omega_\ell \nu^\ell$ with $\omega_\ell \propto n_\ell$.

In [6], the server computes a rate-dependent coefficient for each client k in round r as follows:

$$\zeta_{k,r} = \frac{1}{\sqrt{2\pi}\,\sigma_r} \exp\Big(- \frac{(\nu_{k,r} - \mu_r)^2}{2\sigma_r^2} \Big), \tag{4}$$

and performs an global aggregation with $\lambda_{k,r} = \kappa\,\beta_{k,r}\,\psi_k\,\zeta_{k,r}$ and $\boldsymbol{\Theta}_r = (1 - \lambda_{k,r})\boldsymbol{\Theta}_{r-1} + \lambda_{k,r}\boldsymbol{\Theta}_{k,r}$, where μ_r and σ_r are round-wise statistics of firing rates, $\boldsymbol{\Theta}$ is a parameter set, $\beta_{k,r}$ accounts for information age, and ψ_k for the sample size of client k. In [7], the client selection is posed as a credit-assignment problem in which the credit is a function of firing rate. In each FL round, the server samples a candidate client set and selects the top-P clients by the squared rate change $\Delta\nu_k = \sum_l [\nu_k(\boldsymbol{\Theta}_k) - \nu_k(\boldsymbol{\Theta})]^2$, where $\nu_k(\boldsymbol{\Theta})$ denotes the network-wide firing rate using local ($\boldsymbol{\Theta}_k$) or global ($\boldsymbol{\Theta}$) parameters.

In these settings, however, DP-SGD can significantly perturb them, affecting the rate-based FNL global aggregation and client selection [8]. Large inter-client variability of firing rates (high σ_r) can cause poor discrimination across clients, and concentrate the global update on near-median clients. In addition, it can lead to global updates dominated by clients that do not necessarily have genuinely informative gradients. Note that these perturbations can induce a bias in global updates, directly pull down the global model accuracy, and slow the learning convergence time, especially under non-IID data. Additionally, instability and randomization in client selection ($\Delta\nu_k$ perturbation) lead to the server frequently picking different clients across rounds, even when underlying conditions are similar. This can result in large accuracy gaps because the server can promote clients that should have low influence and demote those that should be high.

4 Ablation Studies

We utilize the Google Speech Commands (GSC) dataset as a representative edge-relevant, event-driven benchmark. Audio waveforms are converted into spike trains via the Speech2Spikes (S2S) encoder with a temporal resolution of $T = 200$. Training is federated across K=10 clients. We adopt class-conditional Dirichlet splits with a concentration parameter $\alpha = 1.0$ to emulate realistic

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/.

data heterogeneity. We use a compact SNN provided by NeuroBench for the GSC task [10]. All LIF layers use the fast-sigmoid surrogate for gradients and $V_{\text{th}} = 1.0$. Each experiment runs for 10 global rounds and 1 local round.

We summarize the privacy budget and clipping effects on rate statistics and global model accuracy in Figure 2. The federated learning is performed using the standard cross-device FNL setting, and local models are aggregated based on the coefficient derived from the related client's average firing rate [6]. When DP is enabled, we apply DP-SGD with target privacy budgets $\epsilon \in \{0.5, 1, 2, 5\}$ at fixed δ and per-example clipping norm $C \in \{0.2, 1\}$. Non-DP runs (i.e., ¬DP) serve as references.

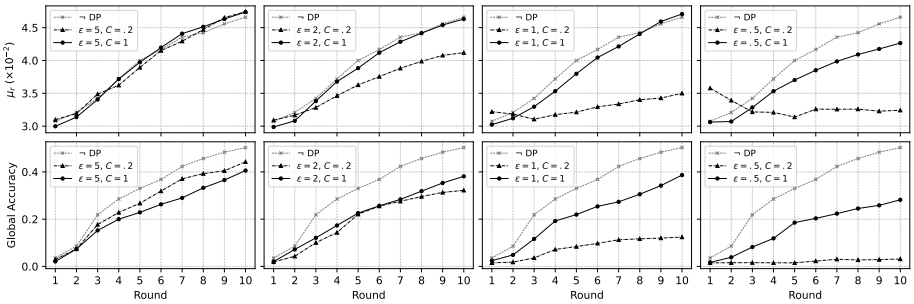

Fig. 2: Average network-wide firing rate μ_r and global model accuracy of ablations over 5 seeds.

The ablations over privacy budgets and per-example clipping norms show that enabling DP-SGD systematically distorts firing-rate statistics. At client level, we obtained a mean Root Mean Square Error (RMSE) of 1.5×10^{-2} across ν_k of reference (¬DP) and differentially private models, while the mean RMSE across layer rate ν^ℓ was 5×10^{-2}, indicating that DP affects firing dynamics not only globally but also heterogeneously across layers. Beside this, still from Figure 2, we observe a clear relationship between the round-wise average firing rate across clients (μ_r) and global model accuracy. At server level, distortions track the degradation in global model accuracy. In particular, tighter privacy (smaller ϵ) and more aggressive clipping (smaller C) introduce stronger gradient bias and noise, which drive the network toward a lower-firing regime (reduced μ_r) and yield consistently lower global accuracy relative to the non-DP baseline. Moreover, decreasing C reduces μ_r across rounds, consistent with stronger clipping bias attenuating effective updates. Conversely, a larger C more faithfully preserve the true gradient, allowing a high-firing regime. Hence, the model can capture specific data features and deliver a superior global performance (i.e., better generalization).

When using the client selection strategy based on $\Delta\nu_k$ (with $P = 5$), we observed displacements of up to 9 positions in the client influence ranking and a considerable degradation in global accuracy between the reference and differentially private learning. Under considerable gradient bias ($C = .2$), when the

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/.

Kendall-τ rank correlation coefficient reaches -0.468, private learning reduces its global accuracy by 6%, delaying the round at which it achieves top accuracy. We observe that the primary reason for this degradation is that the server overweights suboptimal updates and selects inconsistent clients across rounds.

5 Conclusion

This study investigates how example-level DP-SGD disrupts rate-aware coordination in FNL with LIF-based SNNs, demonstrating that the bias and variance induced by DP in firing rates act as primary drivers for global accuracy degradation. We show that DP-induced shifts in firing-rate statistics closely track global accuracy degradation, revealing a structural relation between firing dynamics and performance. Furthermore, DP destabilizes client selection by altering influence rankings, leading to a suboptimal model aggregation performance. Our analysis is restricted to LIF-based SNNs and a single federated event-driven task, leaving broader tasks and per-layer clipping/accounting strategies for future investigation.

References

[1] Y. Venkatesha, Y. Kim, L. Tassiulas, and P. Panda. Federated learning with spiking neural networks. *IEEE Transactions on Signal Processing*, 69:6183–6194, 2021.

[2] N. Skatchkovsky, H. Jang, and O. Simeone. Federated neuromorphic learning of spiking neural networks for low-power edge intelligence. In *IEEE ICASSP 2020*, pages 8524–8528. IEEE, 2020.

[3] S. Shukla, S. Rajkumar, A. Sinha, M. Esha, K. Elango, and V. Sampath. Federated learning with differential privacy for breast cancer diagnosis enabling secure data sharing and model integrity. *Scientific Reports*, 15(1):13061, 2025.

[4] M. Abadi, A. Chu, I. Goodfellow, B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.

[5] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

[6] Y. Wang, S. Duan, and F. Chen. Efficient asynchronous federated neuromorphic learning of spiking neural networks. *Neurocomputing*, 557:126686, 2023.

[7] Q. Zhan, J. Cao, X. Xie, M. Zhang, H. Tang, and G. Liu. Sfedca: Credit assignment-based active client selection strategy for spiking federated learning. *arXiv preprint arXiv:2406.12200*, 2024.

[8] Luiz Pereira, Mirko Perkusich, Dalton Valadares, and Kyller Gorgônio. On the sensitivity of firing rate-based federated spiking neural networks to differential privacy. *arXiv preprint arXiv:2602.12009*, 2026.

[9] J. Eshraghian, M. Ward, E. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. Jeong, and W. Lu. Training spiking neural networks using lessons from deep learning. *Proceedings of the IEEE*, 111(9):1016–1054, 2023.

[10] J. Yik, K. Berghe, D. Blanken, Y. Bouhadjar, M. Fabre, P. Hueber, W. Ke, M. Khoei, D. Kleyko, N. Pacik-Nelson, et al. The neurobench framework for benchmarking neuromorphic computing algorithms and systems. *Nature communications*, 16(1):1545, 2025.