# Federated Learning for Spectrum Sensing in THz Bands

Omar Serghini, Salvatore Serrano
*Department of Engineering*
*University of Messina*
Messina, Italy
omar.serghini@studenti.unime.it
salvatore.serrano@unime.it

Asmaa Maali
*Royal Naval School of Casablanca*
Casablanca, Morocco
maali.asmaa@gmail.com

Hayat Semlali
*Laboratory of Electrical Systems,*
*Energy Efficiency and Telecommunications*
*Cadi Ayyad University*
Marrakesh, Morocco
h.semlali@uca.ma

***Abstract*—Spectrum sensing at terahertz (THz) frequencies presents significant challenges due to extreme signal attenuation, device heterogeneity, and real-time processing constraints. In this work, we propose a lightweight federated learning (FL) framework for cooperative spectrum sensing, tailored to decentralized secondary users operating under non-identically distributed signal conditions. We adopt a lightweight convolutional neural network (CNN) with a multi-head attention refinement and soft attention-based pooling, enabling efficient processing of magnitude-domain MIMO-OFDM signals while preserving key spectral features. We evaluate our attention-enhanced CNN model within a federated learning setup, comparing it against a standard FL aggregation baseline, a centralized model trained on aggregated data, and a local-only model trained on a single device. Results show that our model with FedProx consistently outperforms all alternatives, highlighting the benefits of proximal regularization under non-IID conditions, and demonstrating that federated learning can match centralized performance while preserving privacy and scalability in heterogeneous THz environments.**

***Index Terms*—Attention Mechanism, Cognitive Radio, Cooperative Spectrum Sensing, Federated Learning, THz Communications**

## I. Introduction

The growing density of wireless devices, driven by IoT, autonomous systems, and emerging 6G infrastructure, has pushed the radio spectrum toward unprecedented levels of congestion [1]. Traditional static spectrum allocation, where frequency bands are pre-assigned to licensed users, has proven inefficient under such dynamic and heterogeneous conditions [2]. This has intensified research into more flexible spectrum access paradigms, particularly Cognitive Radio (CR) systems, which allow Secondary Users (SUs) to opportunistically detect and utilize idle spectrum without disrupting Primary Users (PUs) [10]. A key enabler of CR is spectrum sensing (SS), the task of detecting PU activity in real time [3]. The challenge becomes more acute at terahertz frequencies, which are central to next-generation wireless systems due to their ultra-wide bandwidth and low-latency capabilities [4]. However, THz bands suffer from extreme path loss, molecular absorption, phase noise, and severe signal attenuation, making traditional SS techniques like energy detection or cyclostationary analysis unreliable, especially under low signal-to-noise ratios (SNR) and in the presence of hardware impairments [5]. To overcome the limitations of traditional techniques, machine learning, particularly deep learning, has been increasingly applied to spectrum sensing, showing strong performance in extracting patterns from raw RF data [6] [7]. However, most approaches rely on centralized SS approaches, which are often infeasible in distributed sensing scenarios due to privacy, bandwidth, and scalability concerns. Cooperative Spectrum Sensing (CSS) mitigates this by combining inputs from multiple devices, but centralizing such data introduces latency and privacy risks. To address this, Federated Learning offers a decentralized alternative in which only model updates are shared [8]. Yet, applying FL to THz sensing remains challenging due to the heterogeneity of device environments, non-IID data, and the need for low latency, resilient training under resource constraints [9]. In this work, we propose a cooperative spectrum sensing framework for THz bands that combines deep learning with FL to enable decentralized, privacy-preserving inference. Unlike prior studies that rely on centralized training or focus on lower-frequency bands, our approach specifically addresses THz conditions by using local training on raw waveform features and a communication-efficient FL strategy to handle non-IID data and device heterogeneity. The rest of the paper is organized as follows: Section II formulates the problem, Section III describes the data generation pipeline, Section IV details the model architecture and FL setup, followed by results and discussion in Section V, and concluding remarks in Section VI.

## II. Problem Formulation

We consider a distributed CSS scenario in the THz band, where multiple spatially distributed SUs aim to detect the presence of a PU transmission. The THz environment is characterized by severe path loss, atmospheric absorption, Doppler spread, and hardware impairments such as phase noise and beam misalignment, which collectively degrade sensing reliability. Each SU is equipped with a multi-antenna receiver and operates in a wideband Orthogonal Frequency-Division Multiplexing (OFDM) regime. The baseband received signal at the $m-th$ antenna of the $s-th$ SU is modeled as:

$$x_{s,m}(n) = \begin{cases} w_{s,m}(n), & \mathcal{H}_0 \text{ (PU absent)} \\ h_{s,m} \cdot s(n) + w_{s,m}(n), & \mathcal{H}_1 \text{ (PU present)} \end{cases} \quad (1)$$

where $h_{s,m}$ denotes the channel gain and $w_{s,m}(n)$ is the additive noise. The task is to distinguish between the two hypotheses: $\mathcal{H}_0$, where no PU signal is present, and $\mathcal{H}_1$, where the PU is actively transmitting. This forms a binary classification problem at the SU level, where each model must infer the presence or absence of a PU signal based on noisy, device-specific observations.

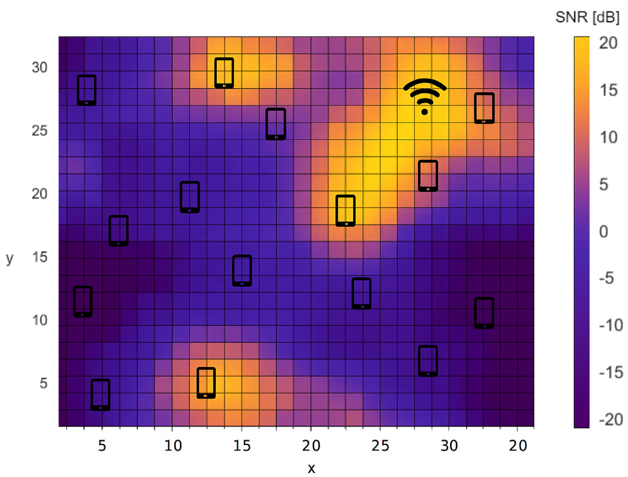


Fig. 1. Spatial distribution of SNR levels and SU positions in the simulated THz environment.

To reflect real-world variability, our simulation introduces heterogeneous SNR conditions across devices, mimicking spatial diversity without explicitly modeling coordinates. As illustrated in Fig. 1, some SUs are assumed to operate under favorable SNR conditions (shown in yellow), while others experience adverse scenarios (blue/purple regions) due to factors such as distance, blockage, or multipath fading. This variation results in non-IID sensing data across devices, complicating centralized learning and motivating the use of decentralized, privacy-preserving strategies.

## III. Data Generation and Preprocessing

To evaluate the proposed federated spectrum sensing framework under realistic 6G conditions, we developed a comprehensive simulation pipeline that generates physically realistic MIMO-OFDM signals in the THz band. Instead of relying on synthetic or random bit sequences, the transmitted content is drawn from a corpus of public domain novels encoded in ASCII, introducing semantically structured and temporally realistic traffic patterns that mimic real 6G applications. Although all devices process the same input corpus, each of the ten simulated sensing devices operates under its own fixed set of physical and environmental conditions such as distance from the transmitter, orientation, molecular absorption, and relative velocity. As a result, the system reflects realistic heterogeneity across nodes, leading to diverse signal characteristics that challenge and inform the design of decentralized learning methods. The system employs 16-QAM modulation and a 2×4 MIMO-OFDM setup, featuring 128 subcarriers, 14 symbols per frame, and a 32-sample cyclic prefix—parameters aligned with typical THz waveform configurations. Transmissions are simulated at a carrier frequency of 300 GHz with a 10 GHz bandwidth. Starting from Friis' equation, signal degradation is modeled with additional randomness to account for molecular absorption and occasional blockage events, in order to better reflect real-world conditions. Doppler shift and phase noise are introduced to reflect mobility and oscillator instability, respectively. Multipath fading is modeled using a four-tap delay profile, where each tap has a random delay and complex gain to capture realistic temporal dispersion. For each device and SNR level (ranging from $-20$ dB to 0 dB in 2 dB increments), 4,000 signal-present (H1) frames are generated. These frames are scaled to the target SNR before complex Additive White Gaussian Noise (AWGN) is added. Signal-absent (H0) frames are generated separately using noise only. Finally, each multi-antenna complex waveform is preprocessed by extracting its magnitude, applying z-score normalization, flattening the antenna channels, and interpolating to a fixed input length of 512 to ensure consistent input dimensionality. Unlike prior methods that rely on FFT-based spectral features or convolution over time-frequency maps, this lightweight, magnitude-only approach reduces computational cost and latency, key considerations for on-device learning in distributed THz sensing. Despite discarding phase and some timing cues, the simplified input remains effective for sensing under resource constraints.

## IV. Model Architecture and Federated Learning Setup

To handle the challenging propagation conditions of the THz band, we design a compact neural architecture tailored for decentralized edge deployment. Building on the use of magnitude-domain inputs to reduce preprocessing overhead, we adopt a 1D convolutional backbone that balances efficiency with representational power. This architecture captures key spectral-temporal patterns in multi-antenna OFDM waveforms while avoiding the complexity of 2D convolutions or frequency-domain transforms, making it well suited for fast, low-latency inference on resource-constrained devices. We use LeakyReLU activations to improve gradient flow in low-SNR conditions and avoid dead neurons. Batch normalization stabilizes training, and regularization helps prevent overfitting under non-IID, resource-constrained local data. To extend the receptive field of the model without relying on recurrent layers, we added a multihead self-attention module after the convolutional blocks. A custom attention pooling layer follows, applying soft weights across time to retain the most informative activations, allowing the model to capture long-range dependencies without adding architectural complexity or latency. The final dense layers use Gaussian Error Linear

Unit (GELU) activations for smoother nonlinearities, followed by a softmax for binary output. Inputs are resampled to 512 to standardize duration and enable fast inference.

TABLE I
FEDERATED LEARNING AND MODEL HYPERPARAMETERS

| Parameter | Value | Description |
|---|---|---|
| Input length | 512 | Signal frame size |
| Classes | 2 | Binary output (PU/no PU) |
| Activations | LeakyReLU / GELU | Conv / dense layers |
| Conv filters | 96 / 64 / 32 | Per conv block |
| Kernel sizes | 9 / 9 / 5 | Per conv block |
| Pooling | MaxPooling1D | Temporal downsampling |
| Attention | 4 × 32 | Heads × key dim |
| Pooling after attn. | Attention-based | Learns time weights |
| Dropout rate | 0.2 | Applied in dense layers |
| Regularization | L2 ($\lambda$ = 0.0005) | Conv layer weights |
| FedProx $\mu$ | 0.01 | Local objective term |

We adopt the FedProx algorithm for federated training, which extends FedAvg by adding a proximal term that penalizes divergence from the global model. This helps stabilize learning in our non-IID setting, where device heterogeneity stems from differences in distance, mobility, and channel conditions. For aggregation, we use weighted averaging based on local dataset sizes to ensure fair contributions across devices. This setup improves training consistency and reduces client drift, particularly under low-SNR conditions. Each device trains locally for a few epochs using the Adam optimizer before sending updates to the server. Throughout the paper, we refer to this federated configuration as FedProx. In each round, 80% of devices are selected, and training runs for 4 rounds to keep communication overhead low, reflecting the constraints of distributed THz sensing. Key settings are listed in Table I.

## V. RESULTS AND DISCUSSION

Experiments were run locally on a multi-core CPU using TensorFlow, without GPU acceleration, reflecting an edge-like setup. The simulation pipeline spans both MATLAB and Python environments: realistic MIMO-OFDM signals were generated in MATLAB under varying THz channel conditions, while all preprocessing steps, model training, and evaluation were conducted in Python using TensorFlow 2.x and Keras, which also served to implement the neural network architecture and the federated learning procedure.

To evaluate the effectiveness of our federated learning framework for cooperative spectrum sensing in the THz band, we compare four models: FedProx, FedAvg (the same architecture trained without proximal regularization), a centralized CNN baseline trained on aggregated data from all devices, and a local-only model trained using data from a single sensing device. These represent a progression from realistic decentralized training (FedProx and FedAvg), to an upper-bound centralized setting with full data access, and finally to a minimal case with isolated, limited training data. This setup enables a clear assessment of key trade-offs, particularly the impact of data heterogeneity, isolated training, and distribution mismatch. Our evaluation focuses on two core performance metrics: the probability of detection as a function of SNR under a fixed false alarm rate ($P_{fa} = 0.1$), and ROC curves generated at a fixed SNR of $-16dB$, where model performance gaps are most visible. For Pd–SNR analysis, we evaluated all models across SNR levels from $-20dB$ to $-8dB$ in $2dB$ steps, using standardized, device-specific signal frames that were not seen during training. This provides a detailed view of model robustness under degraded THz conditions. All measurements were performed on preprocessed magnitude-domain features extracted from complex MIMO-OFDM waveforms. These evaluation procedures are designed to reflect the practical challenges of distributed THz sensing, including non-IID conditions, privacy preservation, and latency constraints. By benchmarking across centralized, federated, and local baselines, we aim to quantify what is gained, and what is reasonably sacrificed, when moving toward more scalable and realistic spectrum sensing solutions.

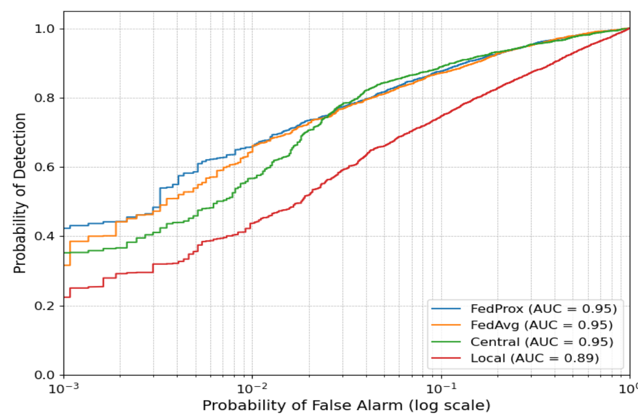


Fig. 2. ROC curves under $-16$ dB SNR condition

Figure 2 shows the ROC curves for the four evaluated models, FedProx, FedAvg, a centralized model, and a local-only baseline, at an SNR of $-16dB$, a mid-to-low range where THz spectrum sensing becomes especially difficult. The x-axis uses a logarithmic scale to emphasize differences in the low false alarm region, relevant in sensitive spectrum-sharing. The local-only model, trained on a single device and tested on a distinct one, exhibits the weakest performance ($AUC = 0.89$), reflecting its limited ability to generalize under unseen channel and device conditions. In contrast, the centralized model, trained on aggregated data from all devices, performs significantly better ($AUC = 0.95$), leveraging full data diversity in an idealized training scenario. This model sets an upper-bound reference for comparison. Both federated models (FedAvg and FedProx) also reach AUCs of 0.95, matching the centralized model despite operating under stricter constraints. This highlights the strength of collaborative training in capturing cross-device variations without requiring centralized data aggregation. Compared to FedAvg, FedProx maintains a consistent advantage in the low false-alarm region, where false detections are most costly. This suggests that the added proximal term helps stabilize training across heterogeneous clients and reduces divergence, which is especially beneficial in non-IID environments. Overall, the results confirm that federated learning, when paired with a lightweight CNN-attention architecture, can approach centralized-level performance in challenging THz conditions—while preserving data privacy and enabling scalable, device-aware deployment.

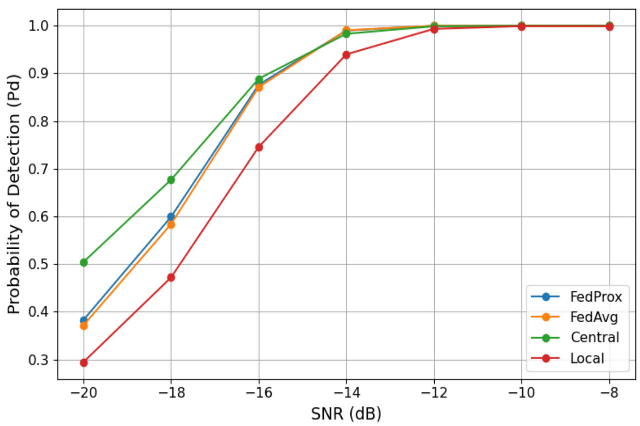


Fig. 3. Pd-SNR curves with a fixed Pf a of 0.1

Figure 3 presents the $P_d$ across a range of SNR levels, evaluated at a fixed false alarm rate of ($P_{fa} = 0.1$). As expected, the centralized model, trained on the full aggregated dataset, achieves the highest detection rates in the lowest SNR conditions ($-20dB$ and $-18dB$, benefiting from complete data visibility and globally optimized gradients. However, both FedProx and FedAvg deliver detection performance that closely tracks the centralized model from $-16dB$ and upward, despite being trained under decentralized, non-IID, and resource-constrained conditions. Notably, FedProx maintains a consistent advantage over FedAvg at lower SNRs, thanks to its proximal regularization, which mitigates client drift and improves training stability. The local-only model, by comparison, shows the weakest performance across all SNRs, underscoring the limitations of isolated training with limited data diversity. These results demonstrate a successful trade-off. While the centralized model offers optimal accuracy under idealized assumptions, the federated models achieve comparable sensing performance without compromising data privacy or requiring centralized coordination. This gain in scalability, privacy preservation, and adaptability to device heterogeneity is particularly important for real-world THz deployments. The fact that these benefits are realized with only minimal performance loss speaks to the robustness of the proposed federated architecture, anchored by a lightweight CNN-attention backbone and streamlined preprocessing, which enables reliable detection even under severe channel impairments. Taken together, the two analyses provide a consistent view of the proposed framework's effectiveness. The federated models, especially FedProx, strike a compelling balance between accuracy and deployment feasibility, approaching centralized performance while accommodating realistic constraints like non-IID data, device heterogeneity, and decentralized processing. These findings affirm the suitability of FL for robust THz spectrum sensing and highlight the value of thoughtful architectural and training choices in challenging wireless environments.

## VI. Conclusion

In this work, we have proposed a federated deep learning framework for cooperative spectrum sensing in the THz band, addressing the unique challenges of high-frequency propagation, non-ID signal distributions and decentralized data availability. Our architecture combines lightweight convolutional layers with multi-head attention and permutation-based pooling to efficiently capture the key spectral-temporal characteristics of MIMO-OFDM waveforms in the magnitude domain. Through extensive simulations under realistic THz conditions, we have demonstrated that federated training with FedProx consistently outperforms FedAvg and generalizes better than a centralized CNN baseline under non-IID and low-SNR conditions. Our results highlight the practical value of incorporating targeted architectural and optimization choices to improve detection robustness in highly heterogeneous, resource–constrained environments. Overall, the findings confirm the viability of applying federated learning to spectrum sensing in challenging THz scenarios and suggest that lightweight attention–based models, trained collaboratively, offer a promising direction for scalable and privacy-aware wireless intelligence in future 6G systems.

## References


[1] Cao, X., Yang, B., Wang, K., Li, X., Yu, Z., Yuen, C., ... & Han, Z. (2024). AI-empowered multiple access for 6G: A survey of spectrum sensing, protocol designs, and optimizations. Proceedings of the IEEE, 112(9), 1264-1302.

[2] Muzaffar, M. U., & Sharqi, R. (2024). A review of spectrum sensing in modern cognitive radio networks. Telecommunication Systems, 85(2), 347-363.

[3] Serghini, O., Semlali, H., Maali, A., Ghammaz, A., & Serrano, S. (2023). 1-D convolutional neural network-based models for cooperative spectrum sensing. Future Internet, 16(1), 14.

[4] Chaccour, C., Soorki, M. N., Saad, W., Bennis, M., Popovski, P., & Debbah, M. (2022). Seven defining features of terahertz (THz) wireless systems: A fellowship of communication and sensing. IEEE Communications Surveys & Tutorials, 24(2), 967-993.

[5] Elbir, A. M., Mishra, K. V., Chatzinotas, S., & Bennis, M. (2024). Terahertz-band integrated sensing and communications: Challenges and opportunities. IEEE Aerospace and Electronic Systems Magazine.

[6] Miuccio, L., Panno, D., and Riolo, S. (2022). A flexible encoding/decoding procedure for 6G SCMA wireless networks via adversarial machine learning techniques. IEEE transactions on vehicular technology, 72(3), 3288-3303.

[7] Miuccio, L., Panno, D., and Riolo, S. (2022). An energy-efficient DL-aided massive multiple access scheme for IoT scenarios in beyond 5G networks. IEEE Internet of Things journal, 10(9), 7936-7959.

[8] Wang, H., & Xu, J. (2024). Online vertical federated learning for cooperative spectrum sensing. IEEE transactions on cognitive communications and networking.

[9] Jiang, W., Zhou, Q., He, J., Habibi, M. A., Melnyk, S., El-Absi, M., ... and Leung, V. C. (2024). Terahertz communications and sensing for 6G and beyond: A comprehensive review. IEEE Communications Surveys & Tutorials, 26(4), 2326-2381.

[10] Serghini, O., Serrano, S., Semlali, H., and Maali, A. (2024, September). Robust DNN-Enabled Cooperative Spectrum Sensing. In 2024 International Conference on Software, Telecommunications and Computer Networks (SoftCOM) (pp. 1-6). IEEE.