SecureRF-Net: RF Device Classification and Unauthorized Transmitter Detection for Model-Identical Radios

Nour Badini[1,*], Carlos Fernando Quiroga Ruiz[1], Fabio Patrone[1] and Mario Marchese[1]

[1]*Department of Electrical, Electronics and Telecommunications Engineering, and Naval Architecture (DITEN), University of Genoa, Genoa, Italy*

Abstract

Device and user authentication are fundamental building blocks of secure wireless systems, but attacks, such as spoofing or device impersonation, can bypass higher-layer protections. Radio Frequency (RF) device fingerprinting mitigates this problem by exploiting hardware-generated signal artifacts that are difficult to forge and persist over time. This paper introduces SecureRF-Net, a deep learning framework for robust RF device identification that operates effectively with model-identical radios and under both closed-set and open-set conditions. Unlike existing approaches that maintain separability across varying channels by injecting transmitter-side impairments or relying on receiver equalization feedback, SecureRF-Net achieves robust classification directly from minimally preprocessed raw I/Q data by constructing a channel-independent spectrogram to suppress channel-induced distortions and employing a dual-loss optimization strategy to learn embeddings that are both discriminative and resilient to propagation variability. Experimental validation demonstrates that SecureRF-Net achieves high accuracy in both static and varying channel conditions when classifying six nominally identical transmitters. Furthermore, in open-set scenarios, the framework accurately detects unauthorized transmitters that use the same MAC addresses of authorized devices, with an accuracy exceeding 97%. These results confirm that SecureRF-Net advances the state of the art in physical-layer authentication by providing a practical, hardware-agnostic, and scalable solution for secure RF device identification in real-world wireless environments.

Keywords

RF device fingerprinting, device authentication, deep learning, channel-independent spectrogram, open-set recognition, unknown-device detection

1. Introduction

Securing modern wireless infrastructures requires ensuring not only the integrity of transmitted data but also the physical authenticity of the transmitting device. Traditional authentication mechanisms, even when enhanced with multiple security factors, remain vulnerable to identity-based attacks such as spoofing, replay, and credential theft. These mechanisms typically operate at higher protocol layers, where adversaries can exploit protocol weaknesses or stolen credentials to impersonate legitimate users. However, signals emitted by Radio Frequency (RF) hardware inherently carry subtle but unique characteristics resulting from unavoidable manufacturing imperfections, which can be leveraged to identify devices at the physical layer. RF device fingerprinting exploits these hardware-induced variations such as oscillator drift, power amplifier nonlinearities, and In-phase/Quadrature (I/Q) imbalance to uniquely characterize each transmitter. Since these impairments are difficult to replicate and remain relatively stable over time, RF device fingerprinting provides a hardware-rooted authentication mechanism that complements traditional security layers. It enables passive and continuous verification of the transmitter identity, even when upper-layer credentials are compromised or cloned, thereby offering an additional line of defense against spoofing and device impersonation.

Despite its promise, RF device fingerprinting faces significant challenges in practical deployment. The wireless channel introduces unpredictable time-varying distortions that can obscure device-specific patterns, reducing classification reliability. Variations in Signal-to-Noise Ratio (SNR), interference, and multipath propagation further complicate the extraction of consistent fingerprints. The problem becomes

Joint National Conference on Cybersecurity (ITASEC & SERICS 2026), February 09–13, 2026, Cagliari, IT

*Corresponding author: Nour Badini, nour.badini@edu.unige.it

nour.badini@edu.unige.it (N. Badini); carlos.quiroga@edu.unige.it (C. F. Q. Ruiz); fabio.patrone01@unige.it (F. Patrone); mario.marchese@unige.it (M. Marchese)

even more challenging when multiple transmitters share identical hardware and signal configurations, as is the case for model-identical radios. In such scenarios, even small channel variations can dominate the intrinsic differences among devices, making classification unstable. Previous research has sought to mitigate channel effects and maintain separability across devices by injecting transmitter-side impairments or employing receiver-side equalization feedback. Although these techniques can improve identification accuracy, they present several drawbacks as they require hardware modifications, rely on calibration feedback loops, and are impractical for large-scale or passive monitoring applications.

This paper introduces SecureRF-Net, a deep learning framework that enhances the robustness of RF device fingerprinting without requiring any transmitter-side modification or channel feedback. SecureRF-Net achieves robust classification directly from minimally preprocessed raw I/Q data by constructing a channel-independent spectrogram to suppress channel-induced distortions and employing a dual-loss optimization strategy comprising focal and batch-hard triplet losses to learn embeddings that are discriminative and resilient to propagation variability. In addition, SecureRF-Net incorporates a confidence-based open-set detection mechanism, allowing the model to identify known transmitters and detect the presence of unauthorized devices.

The structure of the paper is as follows. Section 2 reviews RF device fingerprinting techniques in the literature. Section 3 details the SecureRF-Net processing chain, from signal model to network and training losses. Section 4 describes the experimental setup and reports the incurred results. Section 5 concludes the paper.

2. Related works

Contemporary solutions for emitter identification show a clear transition from closed supervised classifiers to open-set–oriented techniques capable of detecting unknown devices. A relevant early study is that of Sankhe et al. [1], who with Over-the-air Radio Communications Learning and Evaluation (ORACLE) demonstrated that deep Convolutional Neural Networks (CNNs) can identify Wi-Fi transmitters with $\approx 99\%$ accuracy in indoor Line of Sight (LoS) environments, but with severe degradation under dynamic multipath, which motivated the adoption of models more robust to channel effects. Along this line, Hanna et al. [2] evaluated open-set methods based on CNNs and autoencoders on a Wi-Fi Signature dataset (WiSig). In the Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT), they achieved $\approx 98\%$ detection of unknown devices in same-day captures and $\approx 80\%$ with temporal variation over days, highlighting the challenge posed by a nonstationary channel. To improve the separation between known and unknown classes, Ma et al. [3] proposed Multi-Task Prototypical Learning (MTPL), which integrates encoder–decoder, prototypical learning, and Extreme Value Theory (EVT). With 16 indoor multipath Wi-Fi transmitters, MTPL obtained an Area Under the Receiver Operating Characteristic curve (AUROC) ≈ 0.99 and open-set accuracy above 95%, gradually degrading as the number of unknown classes increases.

Puppo et al. [4, 5] presented a more sophisticated approach with Hidden-state Neural Radio Fingerprinting via Variational Analysis (HiNoVa), a Convolutional Neural Network–Long Short-Term Memory (CNN–LSTM) architecture that uses hidden states as a fingerprint. With 15 real multi-protocol devices, HiNoVa achieved an Area Under the Precision–Recall Curve (AUPRC) close to 1 and improvements of 8–15% over traditional CNN/LSTM baselines, showing notable resilience under low SNR. In Automatic Dependent Surveillance–Broadcast (ADS-B) signals, Huang et al. [6] introduced Metric Denoising Autoencoder (MeDAE), a metric denoising autoencoder that, using real data from 9 transmitters and injected Additive White Gaussian Noise (AWGN), improved Area Under the Curve (AUC) by 0.05 to 0.12 compared to conventional autoencoders, especially in low-SNR open-set scenarios.

From a boundary-modeling perspective, Wu et al. [7] proposed Open-Set Support Vector Data Description (OpenSVDD), a solution based on a One-Dimensional Residual Network-50 (1D ResNet-50) extractor and Support Vector Data Description (SVDD) descriptors. Evaluated on 5 Universal Software Radio Peripheral 2901 (USRP-2901) devices in dynamic indoor conditions and 49 real ADS-B aircraft, it exceeded 95% accuracy in closed-set settings and achieved AUROC ≈ 0.97 in open-set.

For embedded hardware and low latency, Zhang et al. [8] developed a lightweight model based on a Residual Network (ResNet)-type CNN, evaluated with 35 LoRa transmitters in an indoor environment. It reported $\approx 92\%$ open-set accuracy and ≈ 0.92 AUC, preserving strong performance on devices with processing constraints.

The use of synthetic outliers was explored by Wang et al. [9] with Synthetic Feature-Constrained Regularization (SFCR), which reinforces the open-set boundary by generating artificial examples in the embedding space. With 12 LoRa transmitters in indoor/outdoor environments, it increased AUROC from 0.92 to ≈ 0.97 even under adverse SNR. At the level of hybrid architectures, Cai et al. [10] developed Joint Radio-Frequency Fingerprinting with Siamese Comparison (JRFFP-SC), a CNN + siamese combination for pairwise comparison. In dynamic indoor channels, it achieved 4–7% improvements in open-set accuracy and > 0.02 in AUROC over baselines.

In large-population scenarios, Gritsenko et al. [11] addressed the detection of novel radios in IEEE 802.11b/g/n Wi-Fi using CNNs with I/Q slicing, evaluating up to 500 known emitters and 5 unknown ones in a real outdoor environment. They achieved $\approx 76\%$ detection with $< 10\%$ degradation in known classes, underscoring the difficulty of large-scale open-set recognition. Another supervised alternative based on real outliers was explored by Karunaratne et al. [12], who trained a One-vs-All (OvA) CNN with 10 authorized transmitters and 50 unauthorized ones. They achieved $\approx 98\%$ detection on the same day and $\approx 80\%$ across different days, again demonstrating the strong influence of the channel on open-set performance.

Finally, Gong et al. [13] extended the multi-task approach with CNN + Counting Generative Adversarial Network (CountGAN) + EVT on Industrial, Scientific, and Medical (ISM) signals (≈ 10 transmitters), achieving F1-score (F1) ≈ 0.91 and False Positive Rate (FPR) $\approx 6.3\%$, validating the usefulness of adversarial outlier generation in heterogeneous unlicensed-band environments. Along with these, Wang et al. [14] combined CNN with Squeeze-and-Excitation (CNN-SE), grouped convolutions, and prototypical networks with EVT in a mixed ZigBee/Wi-Fi scenario (15 transmitters), achieving $\approx 90\%$ accuracy and outperforming OpenMax under variable SNR. In LTE/5G settings, Yin et al. [15] proposed Multi-Channel Convolutional Neural Network (MCCNN) with multi-segment Deep Convolutional Time–Frequency (DCTF) features, achieving 98.7% in LoS, 89.4% in Non-Line of Sight (NLoS), and AUC ≈ 0.98 open-set with 10 devices. Taken together, these works consolidate the trend toward models that integrate open-set regularization, prototypical learning, EVT, outlier generation, temporal fingerprinting, and lightweight architectures, with the goal of maintaining robust performance under temporal variability, multipath, unknown devices, and real channel conditions.

Despite these advances, many existing methods still depend on transmitter-side impairments, receiver equalization feedback, or handcrafted features to counter channel variability, and often degrade when handling model-identical radios or low SNR. SecureRF-Net moves beyond these limitations by introducing a channel-independent spectrogram that suppresses propagation effects without hardware modification, together with a dual-loss embedding strategy that enhances both separability and channel robustness. Combined with a simple confidence-based open-set detector, the proposed SecureRF-Net provides reliable identification and unknown-device detection for model-identical radios under varying wireless conditions, advancing the state of the art in practical RF fingerprinting.

3. Methodology

This section describes the proposed SecureRF-Net framework, which introduces a channel-independent representation and a dual-loss optimization strategy for robust RF device fingerprinting. The approach is designed to extract discriminative and channel-robust features directly from minimally preprocessed I/Q data, without any transmitter-side modification or equalization feedback.

3.1. Signal Model and Preprocessing

The received complex baseband signal can be modeled as:

$$y(t) = h(t) * f(x(t)) + n(t), \tag{1}$$

where $x(t)$ denotes the ideal transmitted waveform, $f(\cdot)$ represents hardware-dependent nonlinearities and impairments introduced by the transmitter (e.g., oscillator drift, power-amplifier distortion, and I/Q imbalance), $h(t)$ the time-varying channel impulse response, and $n(t)$ additive white Gaussian noise. The objective is to extract features representative of $f(\cdot)$ while minimizing the impact of $h(t)$ and $n(t)$.

The analog signal $y(t)$ is sampled at rate F_s, producing a discrete-time complex sequence:

$$y[n] = h[n] * f(x[n]) + n[n], \tag{2}$$

where $y[n] = I[n] + jQ[n]$ are the complex baseband samples with in-phase ($I[n]$) and quadrature ($Q[n]$) components.

Each received recording is divided into fixed-length windows of N samples. Since absolute signal power can vary due to propagation conditions or hardware gain, each segment is RMS-normalized as:

$$\tilde{s}[n] = \frac{y[n]}{\sqrt{\frac{1}{N}\sum_{n=0}^{N-1} |y[n]|^2}}, \tag{3}$$

ensuring that subsequent feature extraction focuses on waveform shape and spectral characteristics rather than amplitude.

The normalized discrete signal $\tilde{s}[n]$ is analyzed in the time–frequency domain using the Short-Time Fourier Transform (STFT):

$$S(f,t) = \sum_{n=0}^{N_{\text{win}}-1} \tilde{s}[n + tH]\, w[n]\, e^{-j2\pi \frac{fn}{M}}, \tag{4}$$

where $w[n]$ denotes a Hamming window of length N_{win}, $H = N_{\text{win}} - N_{\text{overlap}}$ the hop size, f indexes the frequency bins ($0 \le f < M$), and t indexes time frames.

The magnitude $|S(f,t)|$ describes the instantaneous spectral energy of the signal, influenced by both the transmitter hardware and the propagation channel.

Assuming the channel response changes slowly over consecutive STFT frames, it can be considered approximately constant within a short temporal interval Δt:

$$|H(f,t)| \approx |H(f,t+1)| \approx |H(f)|. \tag{5}$$

Under this assumption, the magnitude spectrogram of the received signal within Δt can be approximated as:

$$|S(f,t)| \approx |H(f)| \cdot |F(X(f,t))|, \tag{6}$$

where $|F(X(f,t))|$ encodes the spectral content shaped by transmitter hardware impairments.

To suppress the static channel term $|H(f)|$, we exploit the relative stability of consecutive frames and compute the ratio:

$$S'(f,t) = \frac{|S(f,t+1)|}{|S(f,t)|}. \tag{7}$$

This operation cancels the quasi-static channel gain while preserving the spectral transitions caused by the device's front-end characteristics:

$$S'(f,t) \approx \frac{|F(X(f,t+1))|}{|F(X(f,t))|}. \tag{8}$$

The resulting two-dimensional ratio $S'(f,t)$ forms a channel-independent spectrogram that is then used as input to the proposed SecureRF-Net network for feature extraction and classification.

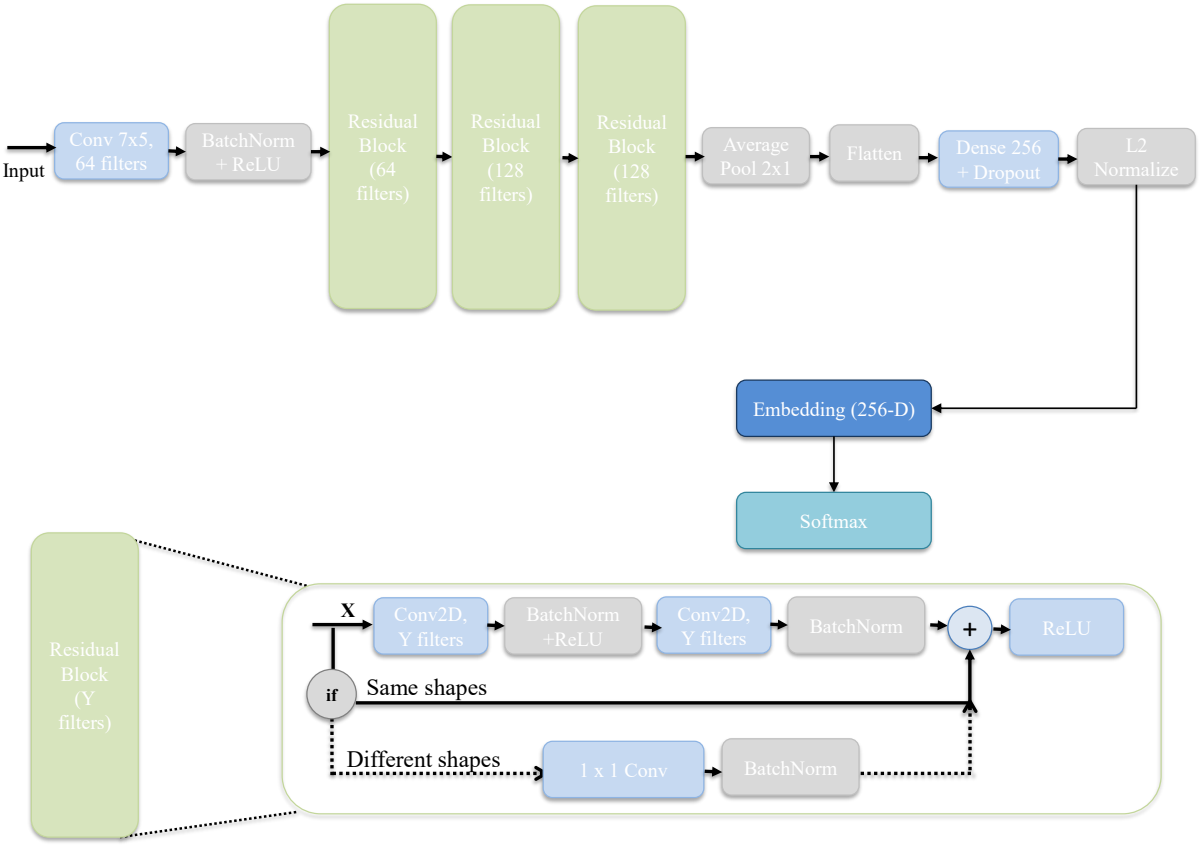

Figure 1: Block diagram of the proposed SecureRF-Net framework. Raw I/Q samples are converted into a channel-independent spectrogram (CISE) and processed by a residual CNN trained under dual-loss supervision.

3.2. Network Architecture

The architecture of SecureRF-Net, shown in Figure 1, follows a residual CNN design inspired by the ResNet family. Its goal is to learn channel-invariant, device-discriminative embeddings from the channel-independent spectrogram.

The network begins with a 7×5 convolutional layer with 64 filters, followed by batch normalization and ReLU activation. This is followed by several residual blocks, each comprising two 3×3 convolutional layers with batch normalization and ReLU activations, connected by either identity or 1×1 projection shortcuts to preserve feature flow. The feature maps are progressively downsampled via average pooling and flattened. A fully connected layer with dropout produces a 256-dimensional feature vector, which is ℓ_2-normalized to form the embedding $\mathbf{z} \in \mathbb{R}^{256}$. In parallel, a softmax head provides posterior probabilities over the K known transmitters.

Training is formulated as a dual-loss optimization problem that balances classification confidence and geometric separability of embeddings. The total loss is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{focal}} + \lambda \mathcal{L}_{\text{triplet}}^{(\text{batch-hard})}, \tag{9}$$

where λ controls the relative contribution of each term.

The focal loss enhances the learning of difficult or low-confidence samples by down-weighting easy examples, encouraging sharper decision boundaries and mitigating class imbalance. The batch-hard triplet loss operates on embeddings: for each anchor, it identifies the hardest positive (same-class) and hardest negative (different-class) examples in the batch and enforces a margin between them. This ensures compact intra-class clusters and large inter-class separation, producing a structured embedding

space suitable for both closed-set classification and open-set detection.

Training is performed using the Adam optimizer with empirically tuned λ and margin values. This dual-objective formulation allows SecureRF-Net to learn robust, channel-invariant, and device-specific features that generalize to unseen environments, providing a foundation for both reliable classification and detection of unauthorized transmitters.

4. Performance evaluation

For the evaluation, the publicly available ORACLE dataset [1] was used. This dataset contains IEEE 802.11a-compliant baseband transmissions generated by several USRP X310 SDRs that are nominally identical. Signals are transmitted at 2.45 GHz and captured at 5 MHz sampling frequency. We selected six transmitters to test the ability to distinguish devices of the same model. The dataset is organized into different transmitter–receiver distance folders, each corresponding to distinct propagation conditions and SNRs.

Three experiments are carried out:

1. **Static channel**: train, validation, and test samples all come from the same distance folder.
2. **Cross-location**: the three splits are taken from disjoint distance folders, so that the test channel is unseen during training.
3. **Unauthorized / Unknown Detection**: to further assess the system's ability to handle open-set scenarios, we evaluate its performance in the presence of previously unseen transmitters (unauthorized).

Figure 2 shows the 2D t-SNE projection of training embeddings for the static scenario. Clusters corresponding to different devices are clearly separated, which indicates that SecureRF-Net has learned device-specific signatures when the channel is fixed. The normalized confusion matrix in Figure 3 confirms this, as the diagonal dominates and the average accuracy is about 94%. This level of performance is consistent with the expectation for a channel-normalized representation in a non-varying environment.

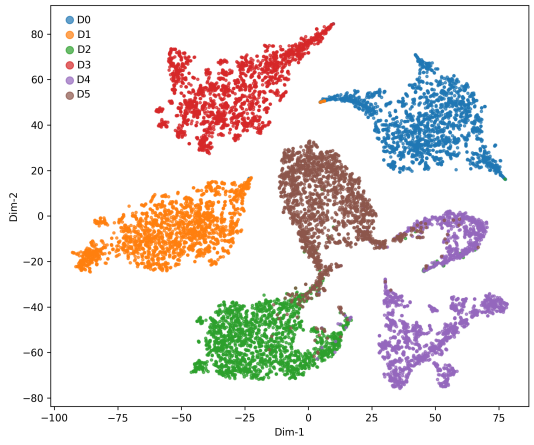

Figure 2: t-SNE visualization of training embeddings for six transmitters under static-channel conditions.

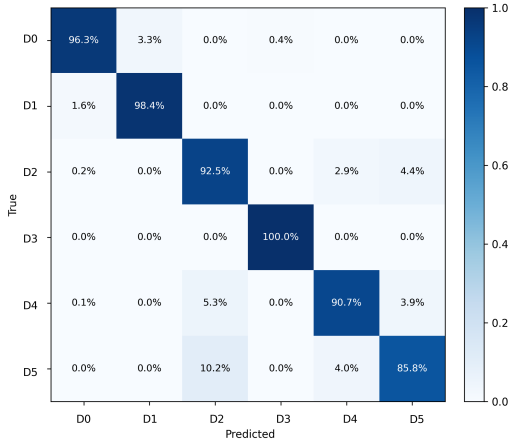

Figure 3: Normalized confusion matrix for six devices under static-channel conditions.

To test robustness to unseen channels, we adopt the distance-disjoint split. Figure 4 depicts the t-SNE projection under cross-location scenario. As shown in the figure, devices still tend to form clusters, but some overlap appears due to the change in propagation. Classification is performed using the softmax head. The row-normalized confusion matrix in Figure 5 shows strong diagonal elements and an average accuracy of around 83%. This indicates that SecureRF-Net can still discriminate among model-identical devices even when the test measurements originate from a receiver position not seen during training.

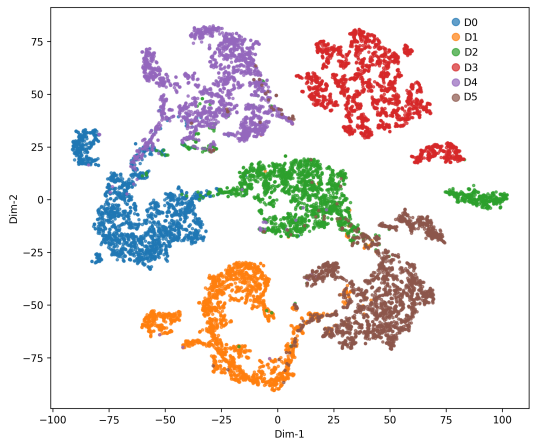

Figure 4: t-SNE visualization of embeddings for six transmitters under cross-location conditions.

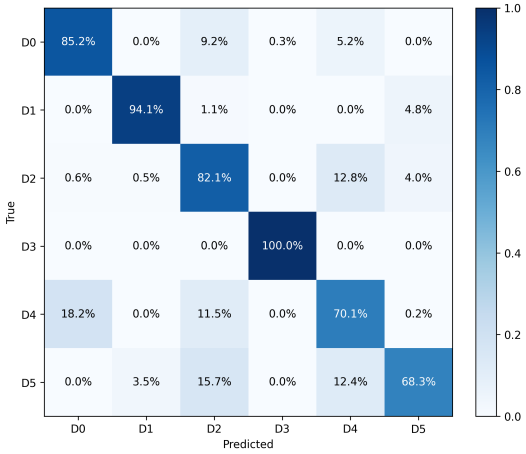

Figure 5: Row-normalized confusion matrix for six devices under cross-location evaluation.

To assess the behavior of the system in open-set conditions, we introduce additional transmitters from the same ORACLE pool that are not used during training. At test time, the classifier receives a mix of:

- signals from the six enrolled (authorized) devices;
- signals from other, previously unseen devices (unauthorized).

SecureRF-Net outputs a softmax probability vector over the known classes. Let $p_{\max}$ be the largest softmax score, we apply a fixed threshold τ as follows:

$$\text{if } p_{\max} \geq \tau \text{ then accept and assign class; else label as “unknown”.}$$

This approach exploits the tendency of out-of-distribution samples to produce lower softmax confidence.

In our experiments, the threshold is set to $\tau = 0.92$. The resulting row-normalized confusion matrix, now including the extra “unknown” class, is shown in Figure 6. Unauthorized signals are rejected with an accuracy above 97%, while the recognition rate of legitimate devices remains high. This demonstrates that the confidence-aware version of SecureRF-Net can generalize beyond the closed set of training classes and is suitable for real wireless security scenarios.

5. Conclusion

This paper presented SecureRF-Net, a deep-learning framework for RF device fingerprinting that is capable of discriminating among nominally identical Wi-Fi transmitters and detecting unauthorized devices. The solution builds a channel-independent spectrogram by dividing consecutive STFT frames, thus suppressing static channel effects while highlighting transmitter-induced spectral dynamics. A residual CNN trained with a dual-loss objective (focal plus batch-hard triplet) produces embeddings that are both discriminative and geometrically organized. On top of the classifier, a softmax confidence threshold enables open-set operation.

Evaluation on the ORACLE dataset shows that SecureRF-Net reaches 94% accuracy in a static channel scenario and about 83% accuracy under cross-location conditions for six devices of the same model. When unseen transmitters are introduced for the test phase, the system identifies them as unknown with more than 97% accuracy, while still maintaining high performance on authorized devices. These results confirm that channel-independent representations combined with confidence-aware classification are a viable path toward practical, physical-layer security mechanisms.

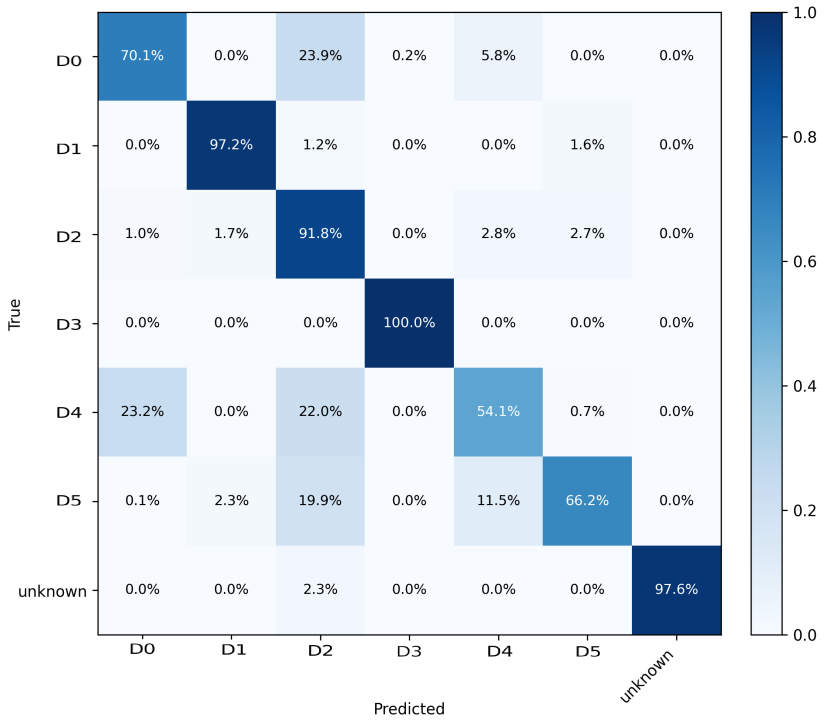

Figure 6: Row-normalized confusion matrix including the "unknown" class for open-set evaluation.

Acknowledgments

This work was funded by the Resilient Trust project supported by the Chips JU and its members under Grant agreement No 101112282

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] K. Sankhe, M. Belgiovine, F. Zhou, et al., No radio left behind: Radio fingerprinting through deep learning of physical-layer hardware impairments, IEEE Transactions on Cognitive Communications and Networking 5 (2019) 481–497.

[2] S. Hanna, S. Karunaratne, D. Cabric, Open set wireless transmitter authorization: Deep learning approaches and dataset considerations, IEEE Transactions on Cognitive Communications and Networking 7 (2020) 59–72.

[3] Z. Ma, S. Fang, Y. Fan, Open-set radio frequency fingerprint identification method based on multi-task prototype learning, Sensors 25 (2025) 5415.

[4] L. Puppo, W.-K. Wong, B. Hamdaoui, A. Elmaghbub, Hinova: A novel open-set detection method for automating rf device authentication, in: 2023 IEEE Symposium on Computers and Communications (ISCC), IEEE, 2023, pp. 1122–1128.

[5] L. Puppo, W.-K. Wong, B. Hamdaoui, A. Elmaghbub, L. Lin, On the extraction of rf fingerprints from lstm hidden-state values for robust open-set detection, ITU Journal on Future and Evolving Technologies 5 (2024).

[6] S. Huang, L. Guo, X. Fu, Y. Peng, Y. Guo, Y. Wang, Q. Zhang, G. Gui, H. Sari, Open-set specific

emitter identification leveraging enhanced metric denoising autoencoders, IEEE Internet of Things Journal 12 (2024) 3453–3462.

[7] C. Wu, S. Chen, G. Sun, H. Fang, Open set rf fingerprint identification for wireless communication devices, IEEE Wireless Communications Letters (2024).

[8] B. Zhang, T. Zhang, Y. Ma, Z. Xi, C. He, Y. Wang, Z. Lv, A low-latency approach for rff identification in open-set scenarios, Electronics 13 (2024) 384.

[9] J. Yang, S. Feng, Y. Wang, X. Wu, M. Yan, Openrfi: Open-set radio frequency fingerprint identification via test-time fine-tuning, IEEE, 2025.

[10] D. Cai, J. Shan, N. Gao, B. He, Y. Chen, S. Jin, P. Fan, Open set rf fingerprinting identification: A joint prediction and siamese comparison framework, arXiv preprint arXiv:2501.15391 (2025).

[11] A. Gritsenko, Z. Wang, T. Jian, J. Dy, K. Chowdhury, S. Ioannidis, Finding a 'new'needle in the haystack: Unseen radio detection in large populations using deep learning, in: 2019 IEEE international symposium on dynamic spectrum access networks (DySPAN), IEEE, 2019, pp. 1–10.

[12] S. Karunaratne, S. Hanna, D. Cabric, Open set rf fingerprinting using generative outlier augmentation, in: 2021 IEEE Global Communications Conference (GLOBECOM), IEEE, 2021, pp. 01–07.

[13] J. Gong, X. Qin, X. Xu, Multi-task based deep learning approach for open-set wireless signal identification in ism band, IEEE Transactions on Cognitive Communications and Networking 8 (2022) 121–135. doi:10.1109/TCCN.2021.3118456.

[14] C. Wang, Y. Wang, Y. Zhang, H. Xu, Z. Zhang, Open-set specific emitter identification based on prototypical networks and extreme value theory, Applied Sciences 13 (2023) 3878.

[15] P. Yin, L. Peng, G. Shen, J. Zhang, M. Liu, H. Fu, A. Hu, X. Wang, Multi-channel cnn-based open-set rf fingerprint identification for lte devices, IEEE Transactions on Cognitive Communications and Networking 10 (2024) 1788–1800.