Pragmatic Cybersecurity
https://www.sciltp.com/journals/pc

Article

Enhancing Channel Robustness in LoRa Radio Frequency Fingerprinting through a Physics-Informed Residual Spectrogram †

Elisée Toe [1,*], Laurent Ferrier [1,2] and Fehmi Jaafar [1]

[1] Department of Computer Science and Mathematics, University of Quebec at Chicoutimi , Chicoutimi, G7H 2B1 QC , Canada
[2] Cégep de Sept-Îles, Sept-Îles, QC G4R 5G7, Canada
* Correspondence: elisee.toe1@uqac.ca

† Extended from a non-archival presentation titled 'Physics-Informed Residual Learning for Channel-Robust LoRa Radio Frequency Fingerprinting', presented at the 31st Australasian Conference on Information Security and Privacy (ACISP 2026), Perth, WA, Australia, 6–9 July 2026.

How To Cite: Toe, E.; Ferrier, L.; Jaafar, F. Enhancing Channel Robustness in LoRa Radio Frequency Fingerprinting through a Physics-Informed Residual Spectrogram. *Pragmatic Cybersecurity* **2026**, *1*(2), 9. https://doi.org/10.53941/pc.2026.100009

Received: 21 July 2026
Revised: 24 July 2026
Accepted: 31 July 2026
Published: 10 August 2026

Abstract: Radio Frequency Fingerprinting Identification (RFFI) is a lightweight authentication method at the physical layer for LoRa IoT devices that avoids cryptographic overhead. However, current deep learning pipelines for RFFI face a robustness gap in that their identification accuracy can decline in time-varying multipath environments caused by real-world mobility, especially in non-line-of-sight (NLOS) settings. Earlier mitigations rely on data augmentation or implicit channel normalization, yet the resulting features often remain dominated by the deterministic LoRa chirp, which can overshadow device-specific hardware signatures. We proposed a physics-informed residual (PIR) framework that incorporates domain knowledge directly into the extraction of features, thereby enhancing the detection of these hardware imperfections. PIR aligns the received preamble and, in the Channel Independent Spectrogram (CIS) domain, subtracts an equally processed ideal LoRa reference. This operation emphasizes residual patterns linked to transmitter hardware impairments and improves numerical stability. We assess the method on a public dataset of commercial LoRa devices, using a homogeneous set of 45 SX1276-based transceivers split across training, open-set enrollment, and rogue device evaluation. The results indicate more robust performance in most scenarios while maintaining strong open-set rogue detection with Area Under the Curve (AUC) greater than 0.99. These findings suggest that physics-informed representation design can enhance the reliability of LoRa RFFI under dynamic deployment conditions.

Keywords: LoRa; radio frequency fingerprinting; physical-layer authentication; physics-informed learning; channel robustness

1. Introduction

The implementation of low-power wide-area networks (LPWANs), especially LoRa (Long Range), has facilitated the development of long-range Internet of Things (IoT) applications characterized by minimal energy use. However, the constraints inherent in these systems present significant security issues. Conventional cryptographic authentication methods can create considerable computational or communication burdens for limited edge devices and are susceptible to risks such as physical cloning of devices [1,2]. Radio Frequency Fingerprinting Identification (RFFI) offers a solution to these challenges by capitalizing on unique hardware imperfections introduced during manufacturing, such as oscillator frequency variances and in-phase and quadrature (I/Q) imbalances, thereby enabling the authentication of transmitters without necessitating alterations to hardware or communication protocols. Recent progress in deep learning has markedly enhanced the efficacy of RFFI by facilitating the automatic extraction of features from unprocessed radio signals and simulated data augmentation [3–8].

Nonetheless, a persistent challenge is the potential decline in RFFI effectiveness under conditions of dynamic multipath fading and Doppler shifts caused by the mobility of devices. This observation indicates a gap in robustness that stems from the signal representations utilized by neural networks [9–15]. Current leading preprocessing methods, particularly the Channel Independent Spectrogram (CIS) [14], seek to alleviate channel-related impacts by temporally dividing adjacent spectrogram windows. Although these techniques perform well in relatively static environments, it has been reported that CIS representations may exhibit decreased stability when device mobility is introduced. Additionally, the resultant representations are often dominated by the deterministic structure of LoRa chirps, which can mask subtler hardware-dependent residual patterns. As a result, neural networks may depend more on channel-specific macroscopic characteristics than on the intrinsic features of the devices, which reduce their ability to generalize across varying channel conditions.

This study introduces the Physics-Informed Residual (PIR), a hardware-oriented residual representation designed to emphasize device-dependent characteristics and improve the robustness of LoRa RFFI under dynamic channel conditions. PIR uses the CIS domain as an intermediate representation space in which the received preamble is compared with an identically processed ideal LoRa reference. By incorporating known LoRa signal parameters into the preprocessing stage, the proposed representation attenuates the common deterministic chirp structure and increases the relative visibility of hardware-dependent residual patterns. Unlike approaches that require target-domain data, domain adaptation, or modifications to the learning architecture, PIR operates entirely at the representation level and can be incorporated into existing open-set RFFI pipelines. The main contributions of this work are summarized as follows:

(1) Physics-Informed Residual Representation: We introduce a preprocessing pipeline that constructs a residual representation by subtracting an identically processed ideal LoRa reference. Signal-level analyses show increased intra-device stability and inter-device separability.
(2) Improved Robustness under Dynamic Channels: Across most evaluated zero-shot environments, PIR provides its largest gain under the most challenging NLOS mobility condition, indicating reduced sensitivity to channel-dependent artifacts.
(3) Preserved Open-Set Discrimination: Despite attenuating the dominant modulation structure, PIR retains strong device discrimination in open-set conditions and supports reliable rogue-device detection.

The remainder of this paper is structured as follows. Section 2 situates our contribution within the existing literature. Section 3 reviews LoRa modulation, introduces the RFFI baseline, and formalizes the problem statement. The proposed framework is detailed in Section 4. Section 5 presents the learning procedure and the open-set authentication pipeline, while Section 6 describes the empirical setup and evaluation scenarios. The principal experimental results are reported in Section 7, followed by an analysis of robustness mechanisms, implications, and limitations in Section 8. Threats to validity are examined in Section 9. Finally, Section 10 summarizes the key findings. All experimental material and source code are publicly available online: https://github.com/YetEliot/PIR-LoRa-RFFI (accessed on 24 July 2026).

2. Related Work

The identification of devices through radio frequency fingerprinting (RFFI) has emerged as a significant physical-layer enhancement to cryptographic methods for authenticating resource-limited Internet of Things (IoT) devices [8, 16]. Initial studies demonstrated that hardware imperfections, inherent to manufacturing processes, create stable signatures unique to each device, which can be extracted from raw baseband signals. Advances in deep learning have markedly increased the accuracy of closed-set identification by facilitating the automatic extraction of features directly from I/Q streams or from time-frequency representations [3,5,17,18].

Despite these advancements, ensuring robustness against channel variations poses a challenge for real-world RFFI applications. In environments characterized by dynamic multipath propagation and Doppler effects, the learned representations may inadvertently reflect correlations with the varying channel response, rather than solely focusing on the intrinsic characteristics of the devices. This can lead to a decline in the model's ability to generalize across different deployment conditions [11,17,19]. Earlier approaches to mitigate this issue have included strategies such as data augmentation [15], feature equalization [18], and methods for tracking carrier frequency offsets [20,21]. While these techniques have shown effectiveness in controlled settings, they may not translate well to environments with extreme dynamics, where channel variations surpass the assumptions inherent in the augmentation or equalization methodologies.

The Channel Independent Spectrogram (CIS) has become a prevalent preprocessing technique designed to diminish localized channel effects by temporally segmenting adjacent spectrogram windows, achieving commendable results in quasi-static environments [14]. Subsequent research has built upon CIS-based representations to enhance robustness against noise, adapt to varying lengths, and improve open-set discrimination [8, 22, 23].

Nevertheless, the reliance on a quasi-static channel model in CIS may falter under conditions involving mobility or non-line-of-sight propagation. Additionally, CIS representations are often predominantly influenced by the deterministic characteristics of the LoRa chirp signal, which may mask subtler device-specific hardware features and restrict the ability to generalize across different domains.

A complementary research axis explores the incorporation of physical or structural signal knowledge into RFFI representations. For instance, Baldini and Bonavitacola have suggested utilizing residuals obtained from Variational Mode Decomposition to modify the spectrogram from the primary modulation envelope [10]. Similarly, Pan et al. have illustrated that leveraging residual channel information can enhance contrastive representation learning in RFFI [23]. Zhang et al. have offered an extensive examination of physical-layer device fingerprinting, covering both theoretical and practical dimensions, and underscoring the importance of signal modeling in enhancing representation efficacy [8]. Collectively, these studies indicate that integrating domain-specific knowledge during the feature extraction process can lead to representations that are both more discriminative and more resilient than those generated through purely data-driven methods.

Domain adaptation strategies constitute another promising approach for mitigating the effects of channel-induced domain shifts. Al-Shawabka et al. [12] have introduced a scalable authentication system that is independent of channel characteristics, utilizing domain-adversarial training techniques. Gaskin et al. [19] have proposed a domain-agnostic authentication model for LoRa that employs transfer learning, while Xie et al. have created a hybrid model-and-data-driven framework for open-set authentication that addresses variability in deployment contexts [13]. Although these strategies are effective, they often necessitate additional training data from the target domains, architectural adjustments, or fine-tuning processes that may be challenging for resource-limited IoT gateways.

Prior work has enhanced the robustness of LoRa radio-frequency fingerprint identification (RFFI) through a range of channel-aware strategies, including preprocessing, data augmentation, feature equalization, and domain adaptation. Among existing preprocessing techniques, the Channel Independent Spectrogram (CIS) is widely regarded as a strong channel-oriented baseline representation for LoRa RFFI because it explicitly targets localized channel effects through temporal division of adjacent spectrogram windows. However, the problem of extracting stable, hardware-dependent representations under rapidly time-varying channel conditions remains only partially resolved, especially in zero-shot deployment scenarios.

This study addresses this gap by introducing a residual representation tailored for channel-robust LoRa RFFI. The proposed method operates exclusively at the representation level and is assessed within an otherwise unmodified open-set RFFI pipeline. This controlled experimental design ensures that the impact of the proposed representation is disentangled from potential confounding factors, such as modifications to the feature extractor, enrollment mechanism, classifier architecture, or rogue-device rejection criterion.

3. Background and Problem Formulation

This section provides an overview of the proposed Physics-Informed Residual (PIR) framework. We will then examine the LoRa signal model and analyze situations in which RFFI systems may underperform in changing wireless environments.

3.1. LoRa Modulation and Signal Model

LoRa utilizes Chirp Spread Spectrum (CSS) modulation, wherein information is encoded through frequency-modulated pulses, known as chirps [1,24]. The theoretical representation of an ideal transmitted baseband signal for a fundamental LoRa up-chirp can be defined as follows:

$$s_{\text{ideal}}(t) = A \cdot \exp\left(j2\pi\left(-\frac{B}{2} + \frac{B}{2T}t\right)t\right), \quad 0 \le t \le T \tag{1}$$

where A signifies the amplitude, B represents the bandwidth, and $T = 2^{SF}/B$ indicates the symbol duration associated with a specific spreading factor SF. A conventional LoRa preamble generally comprises multiple consecutive up-chirps, typically amounting to eight.

In real-world applications, the transmitted signal is subject to various hardware imperfections inherent to the devices, such as frequency offsets in oscillators, I/Q imbalance, and phase noise. Consequently, the received baseband signal can be characterized as:

$$y(t) = h(\tau, t) \otimes f(s_{\text{ideal}}(t)) + n(t) \tag{2}$$

In this equation, $h(\tau, t)$ represents the time-varying impulse response of the wireless channel, $f(\cdot)$ encapsulates the cumulative hardware imperfections of the transmitter that create the radio frequency fingerprint, and $n(t)$ stands for the additive noise. Figure 1 illustrates the material impairment present in the radio-frequency (RF) chain.

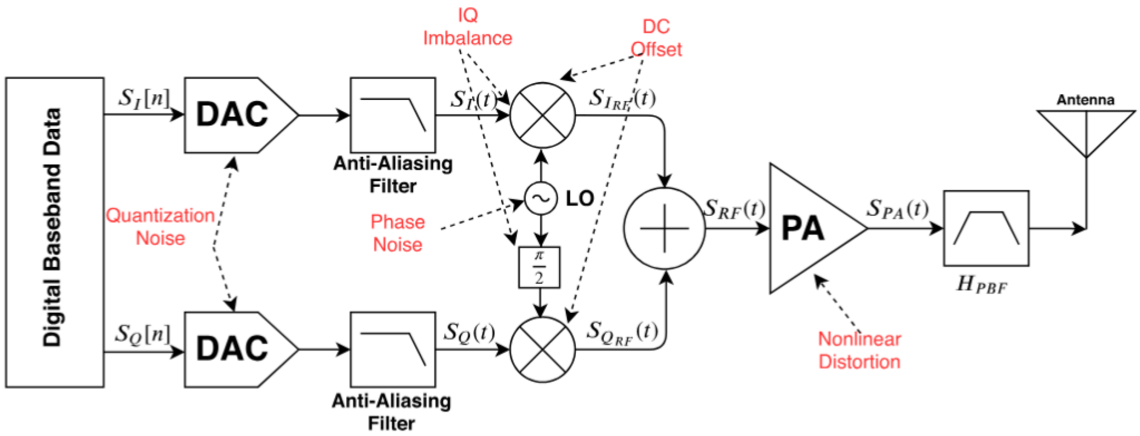

Figure 1. Typical transceiver with various radio-frequency (RF) impairments [17]. DAC—digital-to-analog converter; DC—direct current; LO—local oscillator; PA—power amplifier; PBF—pass-band filter.

Radio Frequency Fingerprinting Identification (RFFI) capitalizes on hardware imperfections induced during manufacturing, denoted as $f(\cdot)$, to authenticate wireless transmitters without the need for additional cryptographic measures. Although deep learning methodologies have exhibited impressive results in controlled settings, their efficacy tends to diminish under significantly varying channel conditions [9,16–18,25].

A significant obstacle is the nature of the representations acquired by neural networks. In practice, the features learned may retain some correlation with the macroscopic channel response $h(\tau, t)$ rather than solely encapsulating the intrinsic characteristics unique to the device, $f(\cdot)$. In scenarios characterized by Non-Line-of-Sight (NLOS) conditions or device mobility, multipath fading and Doppler effects can drastically modify the profile of the received signal, leading to domain shifts that hinder generalizability across diverse deployment contexts [10,12,15,19,22,26,27].

3.2. Channel Independent Spectrogram (CIS)

To address the sensitivity to channel variations, the Channel Independent Spectrogram (CIS) [14,23] has been introduced as a preprocessing representation derived from Short-Time Fourier Transform (STFT) spectrograms and other joint time-frequency representations [8,10,11,23]. CIS operates under the premise that the channel remains quasi-static within adjacent time intervals. Let $S[k, m]$ represent the STFT of the received signal, where k denotes the frequency bin and m indicates the time index. The CIS assumes a factorized form:

$$S[k, m] \approx F_k(X[k, m]) \cdot H[k, m] \tag{3}$$

In this formulation, $F_k(\cdot)$ captures the effects dependent on the transmitter, while $H[k, m]$ portrays the channel response. By segmenting adjacent time windows, CIS diminishes the localized effects of the channel:

$$CIS[k, m] = \frac{S[k, m+1]}{S[k, m]} \approx \frac{F_k(X[k, m+1])}{F_k(X[k, m])} \tag{4}$$

This representation is typically conveyed on a logarithmic scale. While CIS proves effective in quasi-static conditions, its stability may be compromised under rapidly fluctuating channel environments due to two principal limitations:

(1) Residual channel variation: The quasi-static assumption ($H[k, m] \approx H[k, m+1]$) may not be applicable during instances of swift fading or Doppler shifts, resulting in leftover channel components within the representation:

$$CIS \approx \frac{F_k(X[k, m+1])}{F_k(X[k, m])} \cdot \left(1 + \frac{\Delta H}{H}\right) \tag{5}$$

where ΔH represents the temporal variation of the channel.

(2) Numerical sensitivity during division: In regions of spectral nulls or low signal-to-noise ratios (SNR), the denominator $S[k, m]$ may approach zero. Such division in these conditions can exacerbate noise components and bring about unstable spectral features:

$$CIS_{\text{noise}} \approx \frac{s_1 + n_1}{s_2 + n_2} \tag{6}$$

4. Physics-Informed Residual Framework

The proposed Physics-Informed Residual (PIR) framework is structured into five principal phases: (1) acquisition of raw I/Q data along with synchronization; (2) compensation for Carrier Frequency Offset (CFO) and normalization of power levels; (3) generation of an ideal LoRa reference preamble; (4) transformation of the signal into a parallel Channel Independent Spectrogram (CIS); and (5) extraction of differential residuals within the CIS domain. In order to define the methodology, we dissect the received LoRa signal into components that can be physically understood. In its complex baseband representation, the transmitted signal can be formulated as follows:

$$s_{\text{tx}}(t) = A \exp\left(j \left[\phi_{\text{ideal}}(t) + \Delta\phi_{\text{hw}}(t)\right]\right) \tag{7}$$

where $\phi_{\text{ideal}}(t)$ represents the ideal deterministic phase trajectory of the LoRa chirp, and $\Delta\phi_{\text{hw}}(t)$ encompasses phase distortions attributable to hardware. Let $x_{\text{ideal}}(t)$ denote the deterministic modulation component, $x_{\text{imp}}(t)$ denote hardware-induced variations, $h(t)$ denote the channel impulse response, and $n(t)$ denote additive noise. Upon propagation through the wireless channel, the received baseband signal can be modeled as:

$$y(t) = h(t) \otimes (x_{\text{ideal}}(t) + x_{\text{imp}}(t)) + n(t) \tag{8}$$

Within a given transmission burst, $x_{\text{imp}}(t)$ tends to remain device-dependent and relatively stable, whereas $h(t)$ varies across environments and time. This distinction motivates separating hardware-related residual structure from channel-dependent effects through structured preprocessing. Rather than performing explicit channel estimation using methods such as Least Squares (LS) or Minimum Mean Square Error (MMSE), we use the CIS representation as an intermediate domain in which the residual operation is performed.

4.1. Motivating Hypothesis: Device Signatures in Residual Structure

Many representations have been studied, including raw IQ, Fast Fourier Transform (FFT), STFT, CFO-based features, and statistical features. Beyond sensitivity analyses of better spectrogram-based representations, including raw spectrogram and CIS representations, the deterministic structure of the LoRa chirp exerts a predominantly significant influence. Consequently, variations that depend on the device may be obscured by this shared modulation envelope. So how can we extract RF representations that preserve hardware imperfections while reducing channel influence without modifying the learning architecture?

We propose the hypothesis that device identity can be more accurately delineated by examining the residual signal obtained after accounting for the deterministic modulation component of LoRa. Let $x_{\text{ideal}}(t)$ denote a synthesized ideal LoRa preamble created from known transmission parameters. Post time-domain synchronization, the residual representation can be articulated as:

$$r(t) = y_{\text{aligned}}(t) - x_{\text{ideal}}(t) \tag{9}$$

Incorporating this physics-informed structure into the preprocessing phase may direct the feature extractor towards characteristics induced by hardware while minimizing reliance on macroscopic artifacts influenced by the channel. Our intuition starts from a simple observation: the LoRa preamble is deterministic. It is fully defined by known parameters, such as the spreading factor, the bandwidth, and the sampling rate. So, two devices transmitting the same preamble share the same ideal baseband signal. If we can remove this common ideal structure, then the remaining residual may better expose device-dependent distortions. A naive idea would be to subtract the ideal signal directly in the time domain. But this does not work reliably, because the ideal reference has not passed through the same wireless channel as the received signal. The residual would still be dominated by channel phase rotation and multipath effects. Following a thorough evaluation of the residual metric using explicit channel estimation techniques, it was determined that performing the subtraction subsequent to CIS processing would be more advantageous. The CIS domain functions as an intermediate layer, where the channel has already undergone partial normalisation, thereby rendering residual channel effects more conducive to learning processes.

4.2. Residual Extraction Method

The received signal is subjected to a series of deterministic preprocessing steps that adhere to established RFFI protocols. These procedures facilitate reproducible alignment and normalization prior to the extraction of residuals [14,16,17,26]. Although CIS partially mitigates channel effects through temporal division, its ratio-based formulation may remain sensitive under low-SNR conditions or rapidly varying channels. The proposed residual representation measures the deviation of the received representation from an ideal LoRa reference within the CIS domain. Following alignment, the normalized received signal $y_{\text{norm}}[n]$ inherently incorporates macroscopic multipath distortions $h(\tau, t)$. A straightforward time-domain subtraction approach (i.e., $y_{\text{norm}}[n] - x_{\text{norm}}[n]$) proves inadequate for reliably isolating hardware identity, as it fails to properly address swift channel phase rotations that can lead to unpredictable destructive interference. Instead of resorting to unstable explicit equalization in the presence of Doppler shifts, we conduct the differential measurement within the established CIS subspace.

(1) Ideal Reference Synthesis: An ideal discrete-time complex baseband preamble $x_{\text{ideal}}[n]$ is synthesized for each received packet using the spreading factor, bandwidth, and sampling rate associated with that reception. Although the spreading factor and bandwidth may vary dynamically through Adaptive Data Rate (ADR) or other link-configuration mechanisms, their effective values for a given packet can be obtained from the gateway configuration, reception metadata, or appropriate protocol and signal-processing parsing. The sampling rate is determined by the software-defined radio (SDR) acquisition configuration.
(2) Time-Domain Synchronization: The incoming signal $y_{\text{raw}}[n]$ is cross-correlated with the ideal reference through an FFT-based implementation to pinpoint the preamble. Taking Δ_n as the estimated delay to ensure precise sample-level alignment, the aligned sequence is then extracted as:

$$y_{\text{sync}}[n] = y_{\text{raw}}[n + \Delta_n] \tag{10}$$

(3) Carrier Frequency Offset (CFO) Compensation : Residual CFOs, which arise from oscillator mismatches and environmental drifts, are estimated using a delay-and-multiply technique and subsequently compensated to avert coarse frequency shifts from distorting the learned representations.
(4) Power Normalization: Both the synchronized received signals and the ideal references are normalized using Root Mean Square (RMS) scaling to prevent variations in path loss from being misinterpreted as differences in device identity.
(5) Differential Residual Computation: Both normalized received and ideal signals are transformed into the CIS domain using consistent parameters and the proposed residual matrix is derived through element-wise subtraction.

$$CIS_{\text{rx}} = \mathcal{C}y_{\text{norm}}[n], \quad CIS_{\text{ideal}} = \mathcal{C}x_{\text{norm}}[n] \tag{11}$$

$$PIR = CIS_{\text{rx}} - CIS_{\text{ideal}} \tag{12}$$

When the transformation parameters are held fixed, shared structural components are attenuated, enabling the residual representation to emphasize device-dependent deviations from the idealized physical model. Relative to ratio-based CIS features, this subtraction-based formulation can improve numerical stability and reduce the influence of channel-correlated artifacts, while still preserving hardware-specific variability.

It is essential to emphasize that the PIR subtraction step is not intended to eliminate or fully compensate for the wireless channel. Instead, the CIS transformation serves as an intermediate equalization domain, consistent with prior work in which the multiplicative channel response is partially suppressed through temporal division before applying a differential operation [14]. The subsequent subtraction of CIS_{ideal} from CIS_{rx} operates on representations that have already undergone this implicit channel normalization, thereby recasting residual channel effects as additive perturbations rather than dominant multiplicative components. Accordingly, PIR should be interpreted as a representation-level transformation that restructures the feature space, not as a technique for explicit channel estimation or equalization. Figure 2 depicts this transformation for the same received LoRa packet: whereas the CIS domain preserves clearly identifiable macro-modulation patterns, the PIR representation accentuates residual structures associated with device-specific characteristics.

On the left, the unprocessed Short-Time Fourier Transform (STFT) spectrogram is dominated by the high-power LoRa chirp component, which largely masks the underlying hardware-specific signatures. The middle panel, corresponding to the Channel-Independent Spectrogram (CIS), partially compensates for channel-induced effects; however, residual chirp-related structure remains visible. In contrast, the right panel shows that the proposed Physics-Informed Residual (PIR) representation substantially attenuates the global chirp pattern and reveals localized residual features, most prominently around chirp transition regions. These regions are of particular interest because they can exhibit manifestations of oscillator instability, phase noise, and in-phase/quadrature (I/Q) imbalance. As a result, PIR effectively alters the input presented to the neural network, redirecting its emphasis from dominant modulation characteristics toward hardware-specific deviations. The degree to which these residual features are discriminative across devices is quantitatively evaluated in Section 7.1.

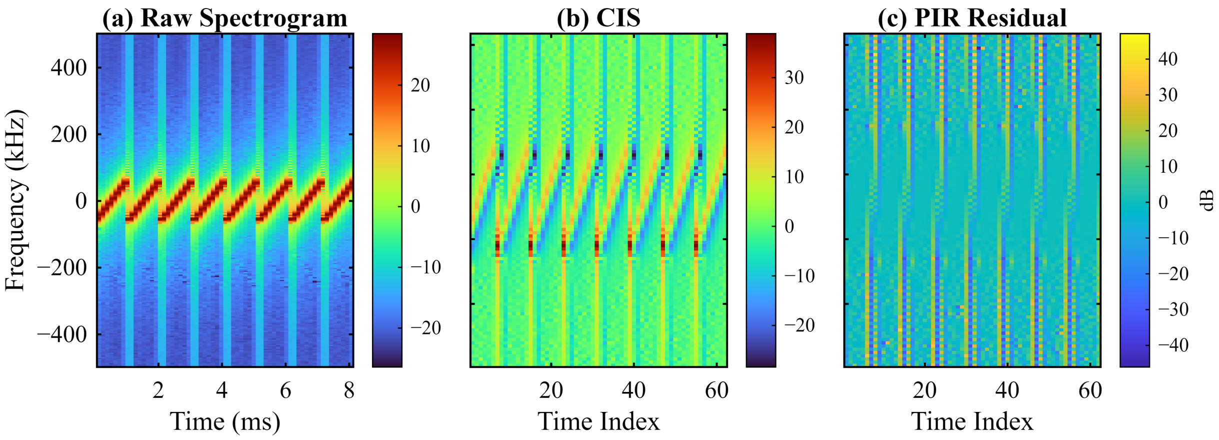

Figure 2. Illustrative comparison of (**a**) STFT spectrogram; (**b**) CIS representation; and (**c**) proposed PIR residual representation.

5. Learning and Classification Pipeline Integration

The primary objective of this study is to modify the representation in order to enhance the attenuation of the common chirp structure and to accentuate the visibility of hardware-dependent residuals. Our proposed method can be incorporated into existing Radio Frequency Fingerprinting Identification (RFFI) frameworks. In order to evaluate the proposed representation's impact, PIR is seamlessly incorporated into the pre-existing RFFI pipeline established by Shen et al. [14] which supports enrolling new devices and rejecting unauthorized transmitters. This study does not introduce any new architecture for feature extraction, nor does it alter the enrollment process, the nearest-neighbor classification approach, or the rules for rejecting rogue devices. Instead, the focus is placed on varying the input representation while all other components of the learning and authentication pipeline remain constant. This deliberate setup ensures that any observed performance differences can be primarily linked to the representation itself, eliminating confounding factors related to the classifier or the open-set decision-making process. The comprehensive integration process is illustrated in Figure 3.

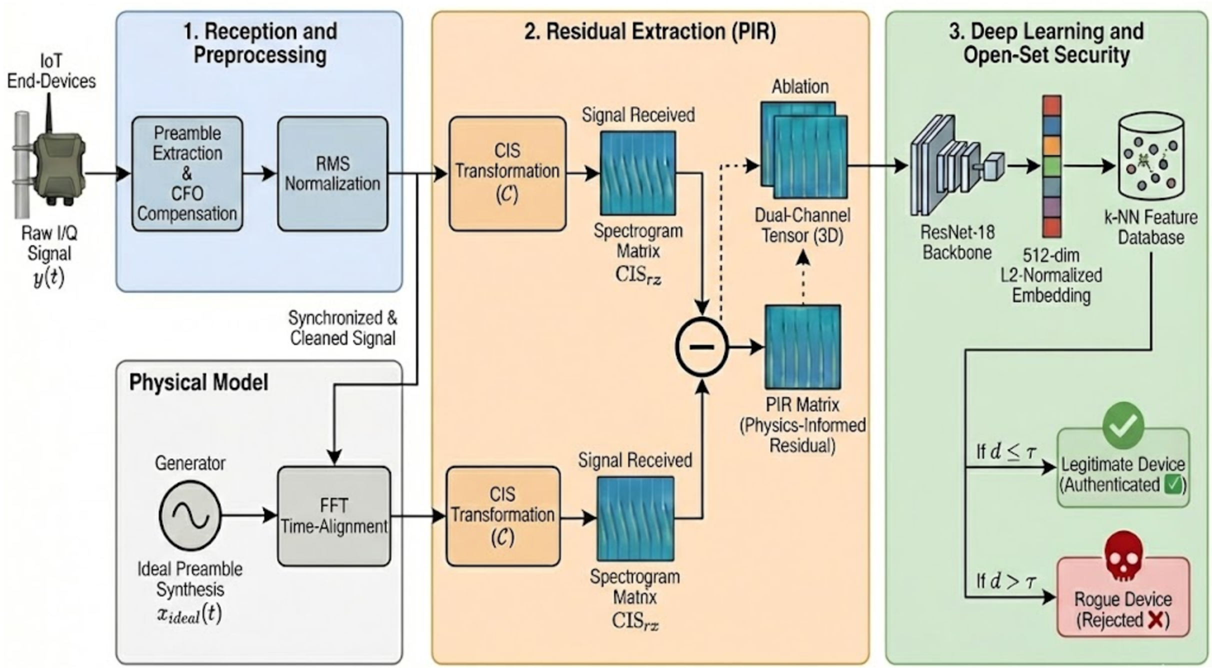

Figure 3. Overview of the proposed RFFI pipeline, including feature extraction, optional dual-channel fusion, embedding projection, and open-set authentication. k-NN—k-nearest neighbors.

In alignment with their methodology, a streamlined ResNet-18 architecture serves as the feature extractor. This network is adapted to accommodate inputs from either a single channel (utilizing PIR or CIS independently) or from a dual-channel configuration, achieved by stacking PIR and CIS along the channel dimension. Each input sample is structured as a time-frequency matrix of dimensions 102×62. When operating in dual-channel mode, the input tensor is represented as $\mathbf{T}_{\text{fused}} \in \mathbb{R}^{102\times62\times2}$. The dual-channel PIR_CIS setup is included as an ablation study to assess whether the two representations provide complementary information. The early fusion approach enables the same feature extractor to concurrently process both streams of information. However, this fused representation is not presented as a novel architectural advancement or as the optimal deployment option. The performance of standalone CIS, standalone PIR, and the dual-channel PIR_CIS fusion is evaluated under consistent training and authentication conditions in Section 7.3.

The subsequent processes for enrollment, identification, and rejection adhere to the open-set protocol established by Shen et al. [14]. During the enrollment phase, normalized embeddings derived from verified authorized devices are recorded in a reference database denoted as $\mathcal{D}$. This method allows for the incorporation of additional authorized devices post-training without necessitating alterations to the feature extractor. During inference, a query embedding $\hat{\mathbf{z}}$ is matched against the embeddings stored in $\mathcal{D}$. The $k = 15$ nearest neighbors are identified based on Euclidean distance. Since the embeddings are normalized using the ℓ_2 norm, the ordering of nearest neighbors remains consistent whether using Euclidean distance or cosine similarity. If the query is recognized as authorized, its device identity is assigned based on the majority vote from the selected neighbors. Their rejection of rogue devices employs the same distance-based decision-making criterion and the authentication score is calculated as the mean distance between the query embedding and its k nearest neighbors (with τ representing the threshold for rejection). Conversely, if this condition is not satisfied, the query is accepted and assigned the majority label based on its nearest neighbors. The experimental framework modifies solely the input representation provided to the established pipeline. The ResNet-18 backbone, the dimensionality of embeddings, the enrollment procedure following training, the k-nearest neighbor approach, and the rejection criteria are consistent across CIS, PIR, and PIR_CIS. This methodological design facilitates a controlled evaluation of the role of physics-informed residual extraction in closed-set identification, cross-environment generalization, and the detection of rogue devices within an open-set context.

6. Experimental Setup

The evaluation of the proposed framework is conducted utilizing the publicly available LoRa RFFI dataset introduced by Shen et al. [14,28]. The dataset comprises baseband I/Q streams captured at a sampling frequency of 1 MHz through a USRP N210 software-defined radio gateway. Every transmission utilizes a spreading factor of SF = 7, a bandwidth of 125 kHz, and a center frequency of 868.1 MHz. To focus on device-specific hardware traits instead of architectural distinctions at the chipset level, our analysis is confined to a homogeneous group of devices featuring identical Semtech SX1276 transceivers.

The evaluation methodology adheres to the standard partitioning outlined in prior work. The training dataset comprises 1000 packets per device, enhanced with multipath and Doppler variations, covering devices numbered 1 to 30. Closed-set identification is tested on these same 30 devices, utilizing 400 packets collected under typical static line-of-sight conditions. For the open-set evaluation, an additional 10 legitimate devices are introduced after training, with 400 packets collected in a residential setting, while five previously unseen devices function as unauthorized transmitters for the purpose of rogue rejection assessment. Cross-environment robustness is evaluated across 10 zero-shot channel scenarios, as summarized in Table 1, using devices that were not included in the training set. The protocol therefore assesses generalization along two dimensions: unseen device identities and previously unseen channel conditions. This joint evaluation provides a stringent test of the model's ability to operate under realistic deployment variations.

All representations are processed using the same pipeline before feature extraction. This pipeline includes time alignment, fractional CFO compensation, and RMS normalization. The learning architecture, training procedure, post-training enrollment protocol, and open-set decision rule are kept unchanged across configurations, allowing the effect of the input representation to be examined independently. The network is trained with RMSprop using an initial learning rate of 10^{-3}, which is progressively reduced to 10^{-6}. Categorical cross-entropy is minimized using a stratified 90/10 split of the training set, with 90% of the samples used for parameter optimization and 10% reserved for validation. Early stopping with a patience of 15 epochs is applied to limit overfitting and maintain a consistent convergence criterion across representations. All experiments are implemented in TensorFlow (2.19.1) and executed on an NVIDIA H100 GPU (Nvidia Corporation, Santa Clara, California, USA). Classification accuracy, macro F1-score, and rogue-device detection AUC are reported as mean $\pm$ standard deviation over three independent random seeds.

Table 1. Evaluation scenarios for channel robustness assessment.

Scenario	Dataset File	Condition
A (LOS)	`channel_problem/A.h5`	Static LOS
B (LOS)	`channel_problem/B.h5`	Static LOS
C (LOS)	`channel_problem/C.h5`	Static LOS
D (NLOS)	`channel_problem/D.h5`	Static NLOS
E (NLOS)	`channel_problem/E.h5`	Static NLOS
F (NLOS)	`channel_problem/F.h5`	Static NLOS
B (Moving Object)	`channel_problem/B_walk.h5`	LOS with moving obstacle
F (Moving Object)	`channel_problem/F_walk.h5`	NLOS with moving obstacle
Mobile Office	`moving_office.h5`	LOS device mobility
Mobile Meeting	`moving_meeting_room.h5`	NLOS device mobility

7. Results

This section encompasses the empirical assessment of the proposed PIR framework. We will first analyze the signal-level discriminative attributes using both quantitative metrics and qualitative spectrogram evaluations. Subsequently, we will discuss closed-set identification efficacy, cross-environment resilience across ten zero-shot channel situations, and open-set detection of rogue devices. All metrics presented are based on the mean $\pm$ standard deviation calculated from three separate training seeds.

7.1. Signal-Level Physical Analysis and Validation

We start by investigating whether the elimination of the deterministic LoRa modulation element accentuates device-specific traits prior to the feature learning phase. This statistical analysis is conducted directly in the representation space, independently of the learning model, to determine whether the observed differences arise from the PIR transformation itself. Figures 4 and 5 summarize the results. PIR exhibits two complementary and measurable properties:

- Intra-device stability: The power spectral densities (PSD) of PIR residuals remain consistent across packets from the same transmitter, resulting in a high intra-device Pearson correlation ($\rho_{\text{intra}} \approx 0.97$). This suggests that the residual structure effectively captures stable hardware-dependent features rather than mere noise or channel artifacts.
- Inter-device separability: Compared to CIS representations, PIR features demonstrate increased normalized pairwise Euclidean distances across all assessed device pairs. The observed average improvement in inter-device separability is approximately 81.6% relative to CIS, indicating a substantial enhancement in discriminative structure prior to any neural feature extraction.

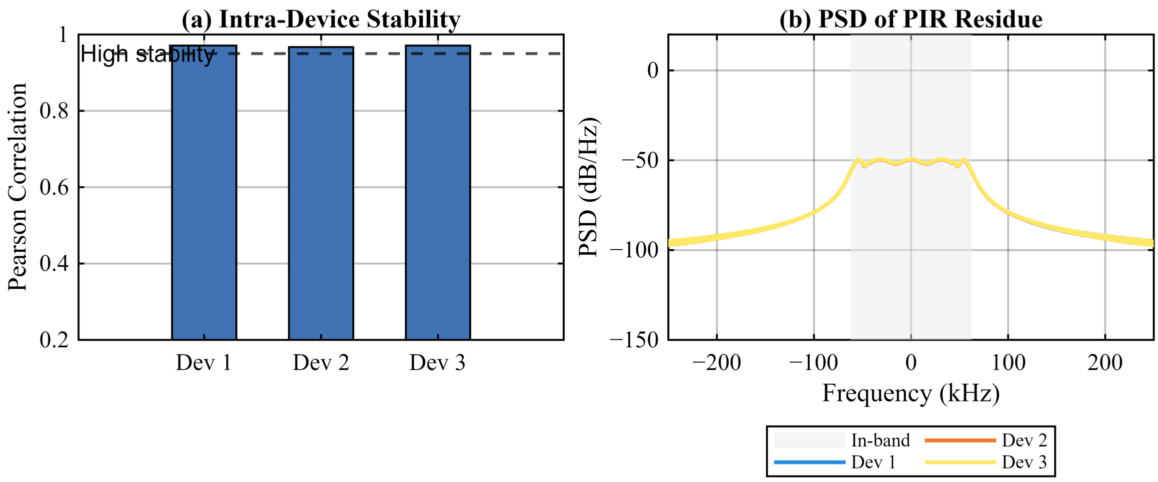

Figure 4. Correlation of residual power spectral densities across packets from the same transmitter.

Effective fingerprinting requires intra-device stability across repeated transmissions and clear inter-device separability between transmitters. Their joint improvement with PIR provides signal-level evidence that subtracting the ideal LoRa reference enhances stable device-dependent residual structure rather than merely redistributing noise. More broadly, the shift from a signal-dominated to a residual-dominated representation suggests that the subtraction step directs the feature extractor toward deviations associated with transmitter hardware.

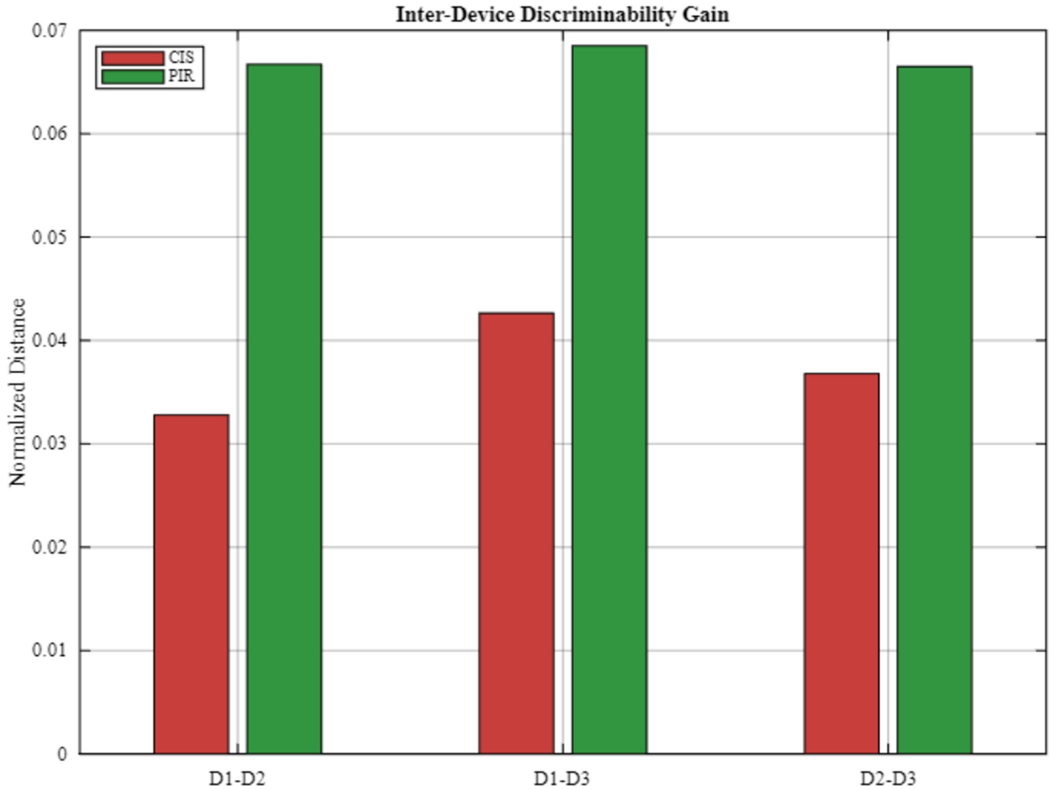

Figure 5. Normalized Euclidean distance between device pairs in CIS and PIR representation.

7.2. Closed-Set Identification Performance

The evaluation of closed-set performance is conducted under static Line-of-Sight conditions, which simulates a typical deployment setting. Table 2 summarizes the identification accuracy and macro F1-score for all configurations evaluated. The CIS baseline achieves an identification accuracy of 99.01%, aligning with previously documented outcomes. The PIR representation on its own reaches an accuracy of 98.63%, signifying that the residual representation retains adequate device-discriminative information for precise closed-set identification, despite the removal of the dominant modulation energy. The dual-channel PIR_CIS configuration yields a comparable performance of 98.12%, indicating that the incorporation of residual features does not compromise identification accuracy under static conditions. The confusion matrix for PIR (Figure 6) indicates that misclassifications are infrequent and randomly distributed, with no discernible confusion patterns, suggesting that the residual structure captures device-specific traits rather than broader group-level characteristics.

Table 2. Closed-set identification performance.

Representation	**Accuracy (%)**	**Macro F1 (%)**
CIS (baseline)	99.01 ± 0.21	99.00 ± 0.22
PIR (standalone)	98.63 ± 0.31	98.61 ± 0.33
PIR_CIS (fusion)	98.12 ± 0.44	98.09 ± 0.46

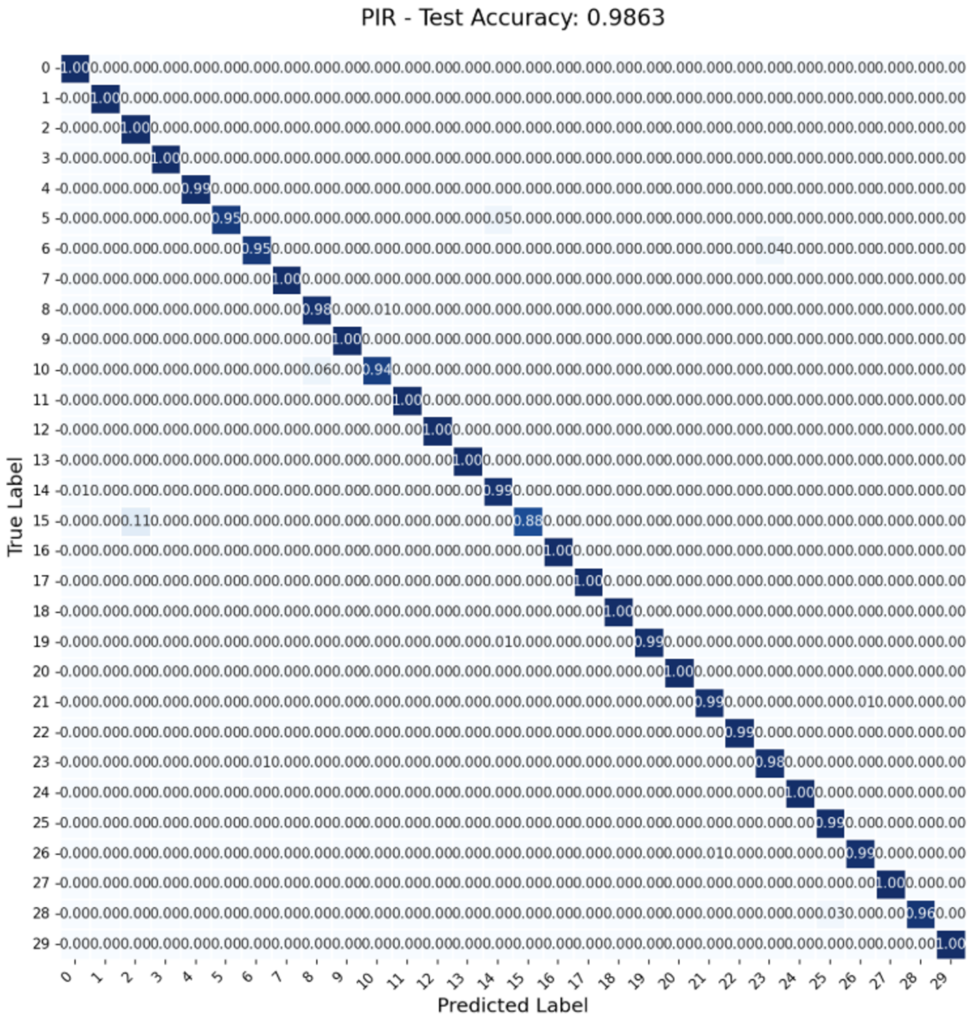

Figure 6. Closed-set confusion matrix for the PIR representation.

7.3. Channel Robustness Under Zero-Shot Scenarios

In this section, we evaluate cross-environment robustness across ten zero-shot channel scenarios encompassing static LOS, static NLOS, moving-obstacle, and device-mobility conditions. Because the evaluated devices and channel conditions were not observed during training, this dual zero-shot protocol assesses generalization to both unseen device identities and unseen propagation environments. Figure 7 presents four representative scenarios, namely LOS A, NLOS F, NLOS Moving F, and Mobile Meeting, while Table 3 reports the complete results across all ten environments. Under severe mobility, CIS accuracy decreases to 61.90% in the Mobile Office scenario and 63.53% in the Mobile Meeting scenario, indicating substantial sensitivity to channel-induced domain shift. In comparison, PIR-based representations improve accuracy in most evaluated environments, with standalone PIR providing its largest gain under NLOS device mobility. Figure 8 summarizes the per-scenario accuracy improvement of PIR over the CIS baseline. Four main observations emerge from these results:

- Broad robustness improvements: Standalone PIR outperforms CIS in eight of the ten evaluated environments, whereas the dual-channel PIR_CIS configuration outperforms CIS in all ten. At the category level, PIR improves mean accuracy by 4.67 percentage points under Static NLOS conditions (84.84% versus 80.17%), by 3.35 points under Dynamic Fading, and by 7.17 points under Severe Mobility. Under Static LOS conditions, the mean difference between PIR and CIS remains marginal (82.76% versus 82.69%). These results indicate that the benefits of the residual representation become more pronounced as the propagation conditions depart from the static LOS environment.
- Standalone PIR under severe mobility: In the most challenging Mobile Meeting scenario, standalone PIR achieves $75.55\% \pm 4.41\%$, compared with $63.53\% \pm 3.71\%$ for CIS, corresponding to a gain of 12.02 percentage points. PIR also exceeds the dual-channel PIR_CIS configuration in this scenario, which reaches $72.10\% \pm 6.29\%$. Mobile Meeting is the only evaluated environment in which standalone PIR outperforms the fused representation. Across the two severe-mobility scenarios, PIR obtains a mean accuracy of 69.89%, slightly above the 69.62% obtained by PIR_CIS, while exhibiting a lower standard deviation ($\pm 6.08\%$ versus $\pm 8.43\%$).

- Trade-offs in fusion: The dual-channel PIR_CIS achieves the highest category means under Static LOS (87.50%), Static NLOS (88.25%), and Dynamic Fading (91.05%). This suggests that CIS retains complementary structural information when channel variations remain static or moderately dynamic. However, under severe device mobility, the channel-dependent variability retained in CIS may reduce the benefit of fusion, allowing standalone PIR to provide slightly greater robustness. The fused configuration is therefore considered an ablation for evaluating representational complementarity rather than a universally preferable deployment choice.
- Variability across training seeds: PIR-based configurations exhibit larger standard deviations in several static environments. For example, under LOS C, the standard deviation increases from $\pm 4.49\%$ for CIS to $\pm 8.87\%$ for PIR. This suggests greater sensitivity to stochastic training effects in some conditions. Nevertheless, in the Mobile Meeting scenario, PIR achieves a substantial accuracy gain while maintaining a standard deviation close to that of CIS. Because these statistics are computed from only three independent seeds, they should be interpreted as indicators of training variability rather than as a complete statistical characterization. Results with overlapping variability, particularly under static conditions, should therefore be regarded as empirical trends rather than statistically established improvements.

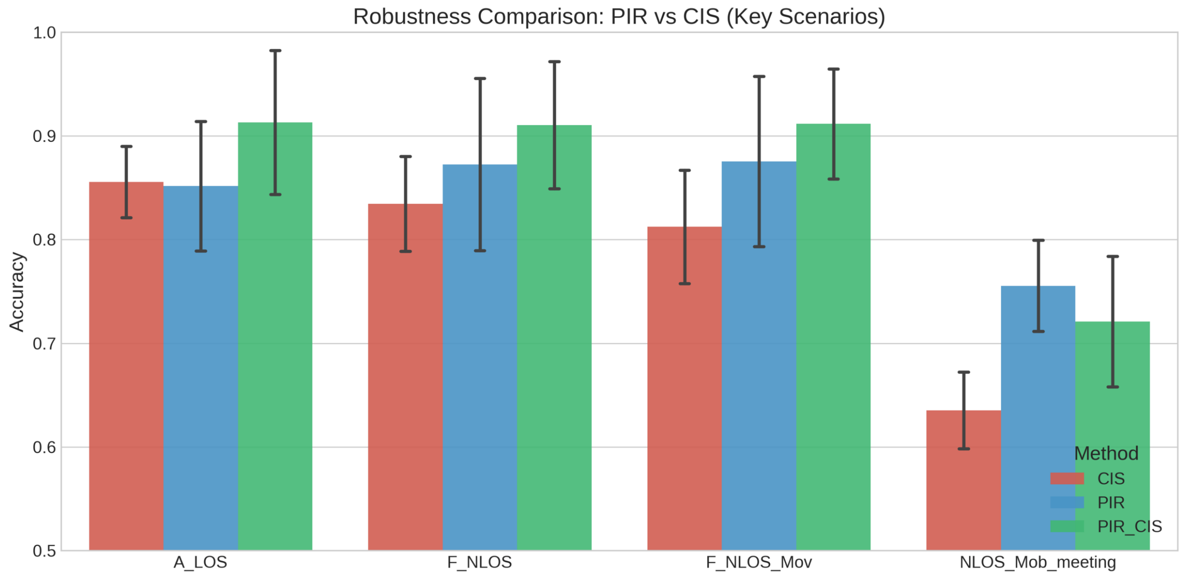

Figure 7. Per-scenario accuracy comparison.

Table 3. Robustness accuracy (mean $\pm$ std %) across zero-shot channel environments (three seeds).

Category	Environment	CIS	PIR	PIR_CIS
Static LOS	LOS A	85.57 ± 3.43	85.17 ± 6.25	91.30 ± 6.94
	LOS B	87.30 ± 3.21	86.12 ± 6.21	91.88 ± 6.78
	LOS C	75.20 ± 4.49	76.98 ± 8.87	79.33 ± 9.56
	LOS Mean	82.69 ± 3.71	82.76 ± 7.11	87.50 ± 7.76
Static NLOS	NLOS D	82.73 ± 3.96	83.62 ± 9.03	87.02 ± 7.73
	NLOS E	74.33 ± 4.43	83.65 ± 8.18	86.67 ± 7.41
	NLOS F	83.45 ± 4.57	87.25 ± 8.31	91.05 ± 6.14
	NLOS Mean	80.17 ± 4.32	84.84 ± 8.51	88.25 ± 7.09
Dyn. Fading	LOS Moving B	85.55 ± 4.13	85.92 ± 8.19	90.92 ± 6.65
	NLOS Moving F	81.23 ± 5.46	87.55 ± 8.21	91.17 ± 5.30
	Fading Mean	83.39 ± 4.80	86.74 ± 8.20	91.05 ± 5.98
Severe Mobility	Mobile Office	61.90 ± 3.22	64.22 ± 7.74	67.13 ± 10.56
	Mobile Meeting	63.53 ± 3.71	75.55 ± 4.41	72.10 ± 6.29
	Mobility Mean	62.72 ± 3.47	69.89 ± 6.08	69.62 ± 8.43

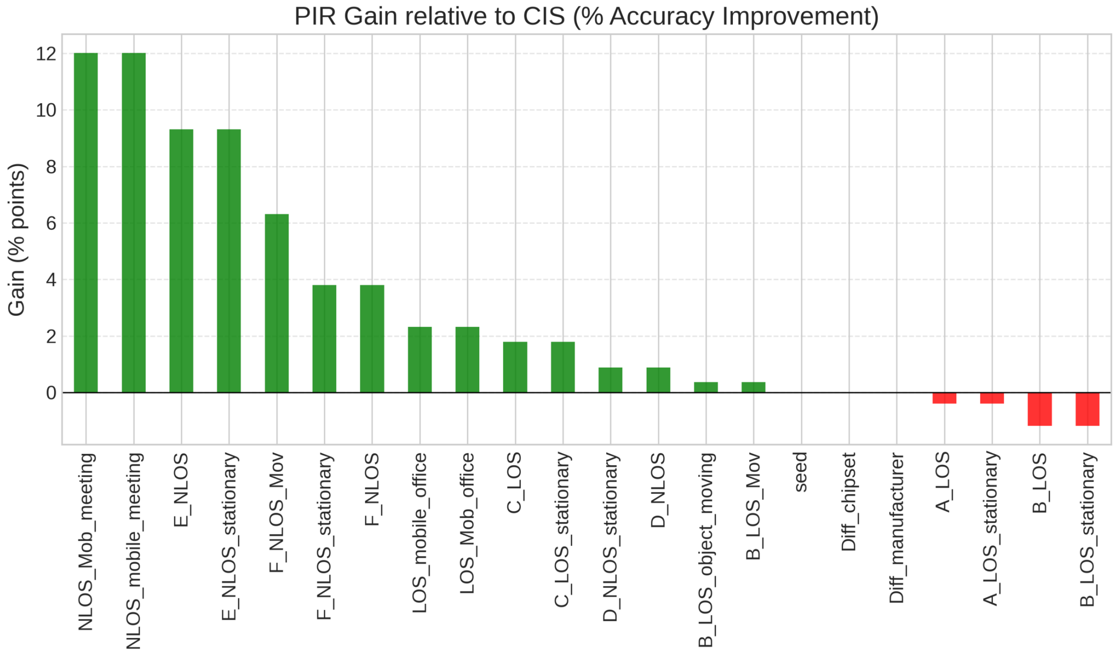

Figure 8. Per-scenario accuracy improvement of PIR.

7.4. Open-Set Rogue Device Detection

We next evaluate whether the proposed representations preserve open-set discrimination under the rogue-device detection protocol defined in Sections 5 and 6. As shown in Figure 9, all three configurations achieve strong discrimination between enrolled and previously unseen transmitters. The CIS baseline obtains an area under the Receiver Operating Characteristic curve (AUC) of 0.9984, while standalone PIR and PIR_CIS achieve AUC values of 0.9980 and 0.9968 respectively. The largest absolute difference relative to CIS is therefore 0.0016. Despite attenuating a substantial portion of the dominant LoRa modulation structure, PIR retains sufficient device-dependent information for reliable rogue-device rejection. In particular, the standalone PIR representation maintains a clear separation between enrolled and unauthorized transmitters in the embedding space. This result indicates that the residual structure preserves identity-related characteristics that are not solely determined by the dominant chirp component.

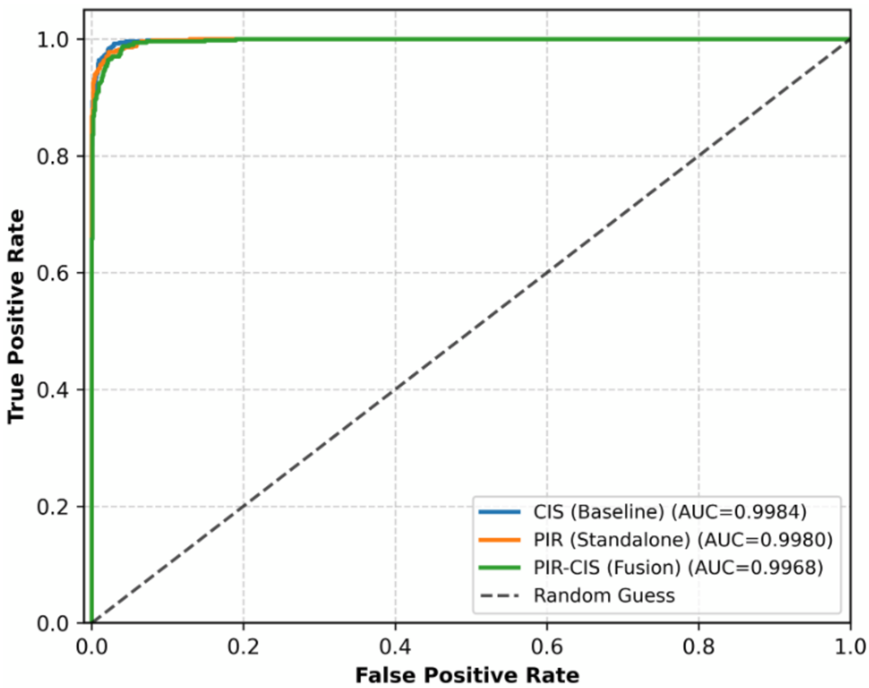

Figure 9. ROC curves for open-set rogue-device detection.

This difference among the three AUC values suggests that under the evaluated conditions, open-set rogue detection is only weakly affected by the choice of input representation. More importantly, PIR preserves open-set performance while providing the robustness gains reported in Section 7.3. The residual representation therefore improves cross-environment identification without materially compromising the rejection of previously unseen transmitters. A compact cross-metric overview is provided (Figure 10) to summarize the relative performance of the three representations across selected channel scenarios and the open-set detection setting.

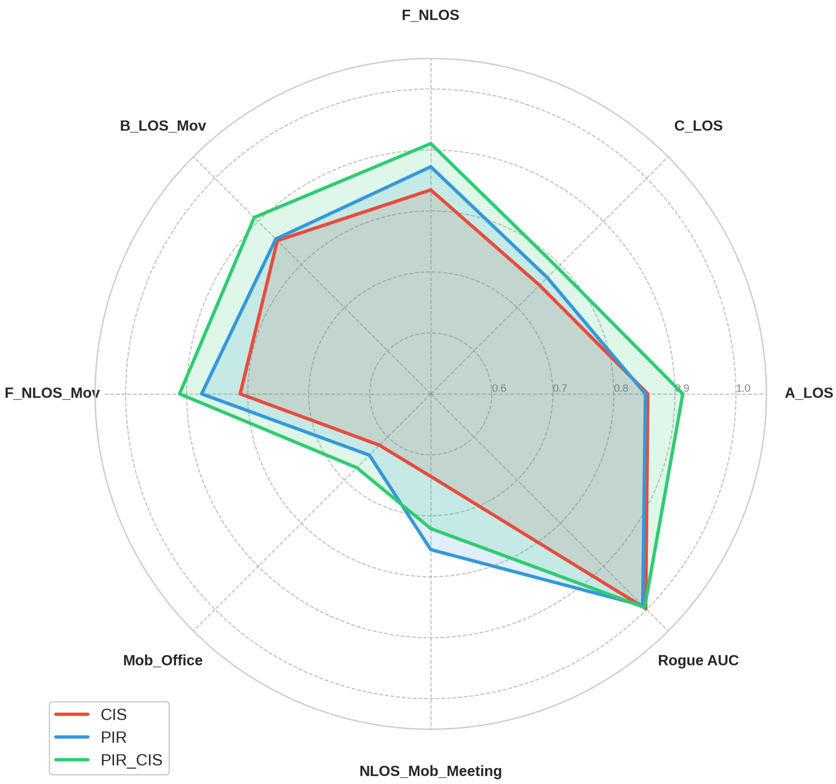

Figure 10. Performance profile across key evaluation dimensions.

8. Discussion

The empirical findings reveal that the limitations in robustness associated with traditional LoRa RFFI pipelines stem not just from the model's capacity but fundamentally from the nature of the input representations. In the preprocessing phase, although local channel effects are somewhat addressed through temporal segmentation, the resulting representations remain heavily influenced by the deterministic structure inherent to LoRa chirps. In dynamic channel scenarios, this prevailing component may still interact with residual multipath distortions, leading to a representation bias where the learned embeddings encapsulate a mixture of device-specific hardware characteristics and variations induced by the channel. This issue becomes particularly significant in NLOS propagation and mobility contexts, where the quasi-static assumption that underpins the CIS approach is violated. The proposed PIR framework addresses this challenge through a reference-guided residual representation. Unlike residual approaches based on signal decomposition or learned domain adaptation, PIR uses an analytically generated LoRa preamble as a physical reference and computes the residual after identical CIS-domain processing. Through this operation, PIR attenuates structural components that are common to devices transmitting the same LoRa preamble. As a result, the resulting feature space gives greater relative weight to deviations from the ideal physical reference. This redistribution of spectral energy reduces the impact of dominant low-frequency modulation components and exposes localized residual patterns that may reflect transmitter hardware impairments. Consequently, the framework enables the construction of more discriminative and more channel-robust embeddings.

Importantly, the formulation based on subtraction modifies the statistical characteristics of the input features in a manner that is distinct from CIS normalization. While CIS relies on division, which can amplify noise in low signal-to-noise ratio (SNR) regions or near-zero denominators, PIR aims to reduce the dominance of such unstable components but does not fully eliminate the numerical sensitivity inherited from the CIS transformation. However, the subsequent subtraction does not introduce an additional ratio operation and reorganizes the resulting features as deviations from a common reference. It is imperative to acknowledge that the PIR subtraction step does not serve to neutralize the wireless channel. The CIS transformation functions as an intermediary equalization domain, a property that has been previously evidenced in studies demonstrating the partial suppression of the multiplicative channel response through temporal division prior to the execution of the differential operation. A convenient interpretation is to regard the CIS representation as containing a device-dependent component together with residual channel effects, with $\varepsilon_{\text{chan}}$ capturing residual channel artifacts. Because the ideal reference undergoes the same transformation, its representation contains the deterministic transformation structure and calibration-related terms, denoted by $\varepsilon_{\text{calib}}$. This approximation is intended as an explanatory model rather than a formal proof of channel cancellation.

$$CIS_{\text{rx}}[k,m] \approx f(\Delta\phi_{\text{hw}}) + \varepsilon_{\text{chan}} \tag{13}$$

$$PIR[k,m] \approx f(\Delta\phi_{\text{hw}}) + (\varepsilon_{\text{chan}} - \varepsilon_{\text{calib}}) \tag{14}$$

In the context of mobility, this representation may make channel-related variation less dominant relative to device-dependent residual structure. This interpretation is consistent with the larger gains observed in NLOS mobility scenarios, particularly the 12% accuracy improvement reported in the Mobile Meeting setting, as opposed to the marginal gains obtained under static conditions. The comparative evaluation of standalone PIR and the dual-channel PIR_CIS fusion further clarifies the complementary roles of these two representation modalities. The CIS representation preserves coarse-grained temporal and spectral structures that can offer valuable contextual information under relatively stable channel conditions. However, under high-mobility regimes, CIS features tend to become less stable, thereby introducing additional variability into the embedding space when fused with PIR. This additional variability can partially offset the benefits afforded by residual-based feature extraction. This behavior accounts for the cases in which standalone PIR slightly outperforms the fused configuration under extreme channel dynamics, suggesting that PIR alone may constitute the more robust choice for deployment in highly dynamic mobility environments. These findings indicate that improvements in RFFI robustness do not inherently require more complex model architectures or larger training datasets. Instead, a representation design grounded in the physical structure of the signal can yield comparable, or even superior, benefits. Because the proposed transformation is applied entirely at the preprocessing stage, it can be incorporated into existing processing pipelines without necessitating architectural modifications or domain-specific retraining an attribute that is compatible with gateway-based RFFI implementations, although its computational cost should be characterized more systematically in future work.

Despite these advantages, certain limitations persist. The efficacy of the PIR approach is contingent upon precise time-domain synchronization and preamble alignment; persistent synchronization errors or extreme Doppler effects could impair the quality of residual extraction. Moreover, the evaluation is confined to a homogeneous subset of Semtech SX1276 devices, which, while allowing for controlled comparisons, does not account for variability across different chipsets. Future research will explore generalization across heterogeneous LoRa hardware platforms, including devices based on the SX1262. Additionally, the current assessment assumes a passive adversarial model; strategies to ensure robustness against active impersonation attempts utilizing software-defined radios or generative techniques merit further exploration. Finally, several promising avenues can be pursued to extend the current framework. Establishing formal criteria under which the residual term $(\varepsilon_{\text{chan}} - \varepsilon_{\text{calib}})$ remains bounded as a function of channel coherence time and Doppler spread would enhance the theoretical foundation of the PIR approach. Substituting the fixed k-NN authentication protocol with metric-learning strategies, such as margin-based losses or prototype-based classification, could bolster embedding separability amidst domain shifts while lowering retrieval costs during enrollment. Moreover, the incorporation of additional residual features associated with power-amplifier nonlinearities, oscillator instability, phase noise, and I/Q imbalance offers further potential for enhancing device-dependent residual structures without the necessity for alterations to the core processing pipeline.

9. Threats to Validity

In order to enhance transparency and provide context for the empirical results, we identify several potential threats to validity, categorized into three key areas:

- Internal Validity: Internal validity is contingent upon precise time synchronization, alignment of preambles, compensation for carrier frequency offsets, and normalization of power levels. Continued timing inaccuracies, significant Doppler shifts, or residual phase misalignment could modify the extracted residuals and influence the learned representations. While the conditions assessed did not reveal any significant degradation related to synchronization, it is essential to investigate the framework's resilience under more extreme mobility scenarios.
- External Validity (Generalizability): The external validity of the findings arises from evaluations performed on a uniform subset of devices that utilize Semtech SX1276 transceivers. This controlled environment effectively

mitigates variability associated with manufacturing processes; however, it fails to demonstrate generalizability across diverse chipsets, various receiver platforms, or wider deployment contexts. Therefore, it is critical to conduct further experiments that incorporate a range of other LoRa transceivers and acquisition systems.
- Threat-Model Validity: Regarding the validity of the threat model, the open-set protocol assesses the ability to reject previously unencountered transmitters within a passive adversarial context. However, it does not account for active impersonation efforts that could involve signal replay, software-defined radios, or generative models adept at mimicking device-specific characteristics. As a result, the findings related to rogue detection should not be construed as indicative of resilience against adaptive threats.

10. Conclusions

This study introduced the Physics-Informed Residual (PIR) framework, designed to enhance the robustness and reliability of LoRa radio-frequency fingerprinting (RFFI) in time-varying wireless environments. The PIR methodology operates by subtracting an analytically derived, ideal LoRa reference signal—processed using the same signal chain—from the received signal in the CIS domain. This operation effectively suppresses deterministic modulation structures while amplifying residual components associated with hardware-induced nonidealities of the transmitter. Because this transformation is applied exclusively during the preprocessing stage, it can be seamlessly integrated into existing RFFI pipelines without requiring modifications to downstream feature extraction procedures or open-set authentication protocols. Across ten zero-shot channel scenarios, empirical evaluations demonstrated that PIR preserves device-discriminative characteristics while substantially improving robustness to channel variability. At the signal level, PIR achieved an intra-device correlation of approximately 0.97 and improved inter-device separability by an average of 81.6% relative to CIS. Under the most adverse high-mobility, non-line-of-sight conditions, the standalone PIR approach increased device identification accuracy from 63.53% to 75.55%, while maintaining strong rogue-device detection performance in open-set settings, with an AUC exceeding 0.99. The results further indicate that a fusion of PIR and CIS can provide complementary information under moderately dynamic channel conditions, whereas standalone PIR is more reliable in environments characterized by pronounced mobility. The present evaluation is limited to a homogeneous set of devices based on the SX1276 chipset and assumes accurate synchronization and preamble alignment. Future work will therefore extend this framework to heterogeneous transceiver platforms, multi-receiver deployments, severe Doppler conditions, and active impersonation attacks targeting RF fingerprints. Overall, the findings demonstrate that a physics-informed representation design can substantially enhance the robustness and generalizability of LoRa RFFI systems without the need for additional training data or more complex learning architectures.

Author Contributions

E.T.: conceptualization, methodology, software, writing; L.F.: supervision, validation, editing; F.J.: supervision, reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Source code for reproducibility is available at: https://github.com/YetEliot/PIR-LoRa-RFFI.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation and revision of this manuscript, we used ChatGPT-based writing tools available through the SciSpace platform to improve grammar, structure, and overall readability, including limited paraphrasing, summarization, and restructuring of selected passages originally written by us. We also used Gemini to assist with the visual design of the overall pipeline architecture figure, based on the technical prompt and specifications defined by us. We did not use these tools to formulate the research problem, develop the proposed methodology or mathematical formulation, conduct the experiments, generate or modify the data, numerical results, and data-derived figures. We manually reviewed, verified, and corrected all AI-assisted outputs and took full responsibility for the final content of the manuscript.

References

1. Milarokostas, C.; Tsolkas, D.; Passas, N.; et al. A Comprehensive Study on LPWANs with a Focus on the Potential of LoRa/LoRaWAN Systems. *IEEE Commun. Surv. Tutor.* **2023**, *25*, 825–867.
2. Toe, E.; Jaafar, F.; Ferrier, L. Securing LoRaWAN in the AIoT Era: A Systematic Mapping Study and an MITRE-Based Threat Matrix. *IEEE Internet Things J.* **2026**, *13*, 3827–3849.
3. Zhang, T.; Ren, P.; Xu, D.; et al. DFSNet: Deep Fractional Scattering Network for LoRa Fingerprinting. In Proceedings of the GLOBECOM 2022—2022 IEEE Global Communications Conference, Rio de Janeiro, Brazil, 4–8 December 2022; pp. 4897–4902.
4. Nair, M.; Dang, S.; Beach, M.A. IoT Device Authentication Using Self-Organizing Feature Map Data Sets. *IEEE Commun. Mag.* **2023**, *61*, 162–168.
5. Li, B.; Cetin, E. Design and Evaluation of a Graphical Deep Learning Approach for RF Fingerprinting. *IEEE Sens. J.* **2021**, *21*, 19462–19468.
6. Hamdaoui, B.; Elmaghbub, A. Deep-Learning-Based Device Fingerprinting for Increased LoRa-IoT Security: Sensitivity to Network Deployment Changes. *IEEE Netw.* **2022**, *36*, 204–210.
7. Huang, D.; Al-Hourani, A. Physical Layer Spoof Detection and Authentication for IoT Devices Using Deep Learning Methods. *IEEE Trans. Mach. Learn. Commun. Netw.* **2024**, *2*, 841–854.
8. Zhang, J.; Ardizzon, F.; Piana, M.; et al. Physical Layer-Based Device Fingerprinting for Wireless Security: From Theory to Practice. *IEEE Trans. Inf. Forensics Secur.* **2025**, *20*, 5296–5325.
9. Elmaghbub, A.; Hamdaoui, B. Distinguishable IQ Feature Representation for Domain-Adaptation Learning of WiFi Device Fingerprints. *IEEE Trans. Mach. Learn. Commun. Netw.* **2024**, *2*, 1404–1423.
10. Baldini, G.; Bonavitacola, F. LoRa Radio Frequency Fingerprinting with Residual of Variational Mode Decomposition and Hybrid Machine-Learning/Deep-Learning Optimization. *Electronics* **2024**, *13*, 1925.
11. Huan, X.; Wu, C.; Lei, Y.; et al. LLRF: Towards Long-Term LoRa Radio Frequency Fingerprint Identification Based on Transfer Learning. *IEEE Trans. Veh. Technol.* **2025**, *74*, 19921–19926.
12. Al-Shawabka, A.; Pietraski, P.; Pattar, S.B.; et al. SignCRF: Scalable Channel-Agnostic Data-Driven Radio Authentication System. *IEEE Trans. Mobile Comput.* **2025**, *24*, 9383–9394.
13. Xie, R.; Xu, W.; Chen, Y.; et al. A Generalizable Model-and-Data Driven Approach for Open-Set RFF Authentication. *IEEE Trans. Inf. Forensics Secur.* **2021**, *16*, 4435–4450.
14. Shen, G.; Zhang, J.; Marshall, A.; et al. Towards Scalable and Channel-Robust Radio Frequency Fingerprint Identification for LoRa. *IEEE Trans. Inf. Forensics Secur.* **2022**, *17*, 774–787.
15. Al-Shawabka, A.; Pietraski, P.; Pattar, S.B.; et al. DeepLoRa: Fingerprinting LoRa Devices at Scale Through Deep Learning and Data Augmentation. In Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing; ACM: Shanghai, China, 26–29 July 2021; pp. 251–260.
16. Ahmed, A.; Quoitin, B.; Gros, A.; et al. A Comprehensive Survey on Deep Learning-Based LoRa Radio Frequency Fingerprinting Identification. *Sensors* **2024**, *24*, 4411.
17. Elmaghbub, A.; Hamdaoui, B. LoRa Device Fingerprinting in the Wild: Disclosing RF Data-Driven Fingerprint Sensitivity to Deployment Variability. *IEEE Access* **2021**, *9*, 142893–142909.
18. Zhou, X.; Hu, A.; Li, G.; et al. A Robust Radio-Frequency Fingerprint Extraction Scheme for Practical Device Recognition. *IEEE Internet Things J.* **2021**, *8*, 11276–11289.
19. Gaskin, J.; Hamdaoui, B.; Wong, W.K. Tweak: Towards Portable Deep Learning Models for Domain-Agnostic LoRa Device Authentication. In Proceedings of the 2022 IEEE Conference on Communications and Network Security (CNS), Austin, TX, USA, 3–5 October 2022; pp. 1–9.
20. Jian, T.; Rendon, B.C.; Ojuba, E.; et al. Deep Learning for RF Fingerprinting: A Massive Experimental Study. *IEEE Internet Things Mag.* **2020**, *3*, 50–57.
21. Shen, G.; Zhang, J.; Marshall, A.; et al. Radio Frequency Fingerprint Identification for LoRa Using Spectrogram and CNN. In Proceedings of the IEEE INFOCOM 2021—IEEE Conference on Computer Communications, Vancouver, BC, Canada, 10–13 May 2021; pp. 1–10.
22. Shen, G.; Zhang, J.; Marshall, A.; et al. Toward Length-Versatile and Noise-Robust Radio Frequency Fingerprint Identification. *IEEE Trans. Inf. Forensics Secur.* **2023**, *18*, 2355–2367.
23. Pan, R.; Chen, H.; Shen, G.; et al. Residual Channel Boosts Contrastive Learning for Radio Frequency Fingerprint

Identification. *IEEE Wirel. Commun. Lett.* **2025**, *14*, 1728–1732.

24. Maleki, A.; Nguyen, H.H.; Bedeer, E.; et al. A Tutorial on Chirp Spread Spectrum Modulation for LoRaWAN: Basics and Key Advances. *IEEE Open J. Commun. Soc.* **2024**, *5*, 4578–4612.
25. Elmaghbub, A.; Hamdaoui, B. Comprehensive RF Dataset Collection and Release: A Deep Learning-Based Device Fingerprinting Use Case. In Proceedings of the 2021 IEEE Globecom Workshops (GC Wkshps), Madrid, Spain, 7–11 December 2021; pp. 1–7.
26. Restuccia, F.; D'Oro, S.; Al-Shawabka, A.; et al. DeepRadioID: Real-Time Channel-Resilient Optimization of Deep Learning-Based Radio Fingerprinting Algorithms. In Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing, Catania, Italy, 2–5 July 2019; pp. 51–60.
27. Chen, Z.; Peng, L.; Hu, A.; et al. Generative Adversarial Network-Based Rogue Device Identification Using Differential Constellation Trace Figure. *EURASIP J. Wireless Commun. Netw.* **2021**, *2021*, 72.
28. Shen, G.; Zhang, J.; Marshall, A. LoRa_RFFI_dataset. Available online: https://ieee-dataport.org/open-access/lorarffidataset (accessed on 6 September 2025).