Open-world Radio Frequency Fingerprint Identification via Augmented Semi-supervised Learning

Zehua Han[1], Jing Xiao[1*], Qirui Zhao[2†], Zhexuan Cui[2†], Yufeng Wang[1], Duona Zhang[3], Wenrui Ding[1]

[1] Beihang University
[2] Northeastern University
[3] North China University of Technology
{ShuasHan2020, changshengEVA, vin019843}@gmail.com, {xiaoj, wyfeng, zhangduona, ding}@buaa.edu.cn

Abstract

In complex electromagnetic environments, the identification and differentiation of diverse radio frequency (RF) emitters become particularly crucial. Existing RF fingerprinting methods demonstrate limitations when dealing with numerous unknown emitters, making it challenging for accurate classification and recognition. These limitations hinder the effective handling of specific unknown emitters. To address this issue, we introduce a novel RF fingerprinting method suitable for open-world conditions for the first time. We develop a novel RF fingerprinting model, Roinformer, to extract signal features with positional attention. We then leverage data augmentation strategies such as noise jitter and signal frame rearrangement to construct an effective pre-training model. Moreover, by incorporating instance-level similarity loss and a novel local entropy regularization approach, we significantly enhance the accuracy of known class identification and mitigate the catastrophic forgetting of known signal samples. Experimental results on three temporal signal datasets demonstrate that our method effectively recognizes both the known and unknown classes, outperforming several state-of-the-art methods by a large margin.

Code — https://github.com/ShuaS2020/OpenRFI

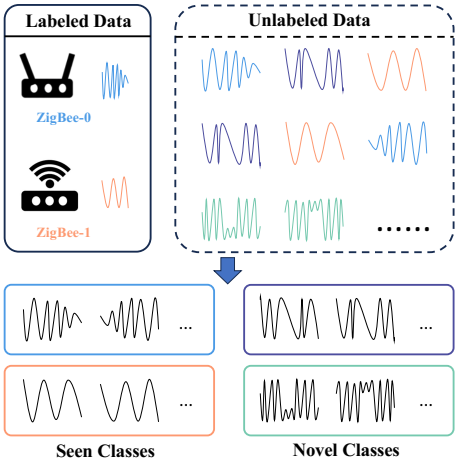

Figure 1: Illustration of open-world RFFI task. In this scenario, unlabeled samples contain RF fingerprint signals emitted by both known and unknown emitters. The objective is to identify RF fingerprint signals from known emitters while simultaneously discovering and recognizing signals from unknown emitters within the unlabeled samples.

Introduction

With the rapid advancement of wireless communication technologies, the Internet of Things (IoT) and associated sensor technologies also see significant progress. In recent years, the extensive deployment of IoT devices provides technological support for applications in smart homes, intelligent transportation systems, and smart cities. However, due to the openness of wireless communication channels, these devices are often exposed to threats of unauthorized access and communication security. To address these concerns, radio frequency fingerprint identification (RFFI) technology emerges, offering a novel solution.

RFFI technology enhances the security (Zhang et al. 2021) of wireless communication systems by analyzing unique characteristics carried by wireless signals during transmission for device identification and authentication. These fingerprints, unintentional by-products of hardware imperfections, are closely associated with the devices and exhibit a high degree of uniqueness, making them difficult to replicate or forge (Baldini and Steri 2017). Thus, radio frequency fingerprints serve as effective and unique identifiers for wireless device authentication.

Alongside continuous advancements in computer technology, deep learning-based RFFI methods (Hua et al. 2018; Liu et al. 2019; Polak and Goeckel 2015; Deng et al. 2023; Shen et al. 2023) significantly improve identification performance. Nevertheless, in complex and dynamic electromagnetic environments, the emergence of unknown emitters challenges the applicability of models trained on known device signal distributions. As shown in Fig. 1, to address this, some studies introduce RFFI techniques in open-set environments to effectively distinguish between known and un-

*Corresponding authors.
†These authors contributed equally.

known emitters. However, these methods typically categorize all unknown emitters into a single class, failing to further differentiate them. In response to these challenges, this study introduces a novel open-world Radio Frequency Fingerprint Identification method, OpenRFI. Initially, we develop a novel RFFI model, Roinformer, by integrating the intrinsic properties of radio frequency fingerprints with advanced deep learning models and encoding techniques. This model surpasses the benchmark Transformer (Vaswani et al. 2017) model in computational complexity and accuracy, validated through supervised learning. Subsequently, we incorporate two novel data augmentation methods and pre-train the Roinformer model using the self-supervised and semi-supervised learning framework, SimCLR (Chen et al. 2020), enabling it to learn latent features between different categories from unlabeled data. Ultimately, addressing the shortcomings of the OpenNCD (Liu et al. 2023) method, we introduce instance-level similarity loss and prototype-group-level local entropy regularization loss, constructing a more robust open-world new class discovery framework, OpenRFI, which achieves significant performance improvements on the RFFI dataset.

The main contributions of this paper are summarized as follows:

- This is the first instance of extending radio frequency fingerprint identification from open-set to open-world recognition, facilitating further classification of anomalous emitters.
- We propose an efficient radio frequency fingerprint feature extraction model, Roinformer, incorporating RoPE method into the Transformer-based architecture.
- We introduce two time-series augmentation methods based on radio frequency fingerprints, which are effectively utilized within self-supervised and semi-supervised learning frameworks.
- To mitigate the catastrophic decline in identification performance for known emitters, we introduce instance-level similarity loss and prototype-group-level local entropy regularization methods, effectively preventing performance degradation.

Related Work

Semi-supervised Learning. In many real-world scenarios, acquiring a large amount of unlabeled data is relatively easy. However, manually labelling this unlabeled data requires specialized equipment and significant time investment, making it costly. Semi-supervised learning can leverage a small amount of labelled data along with a large amount of unlabeled data, addressing the high cost of data labelling in supervised learning and the inaccuracy of models in unsupervised learning. Traditional semi-supervised learning methods (Li, Xiong, and Hoi 2021; Xie et al. 2020; Sohn et al. 2020; Berthelot et al. 2019; Guan et al. 2025) assume that the labelled samples encompass all sample classes, which corresponds to a closed-set scenario. If unknown class samples are mixed with the unlabeled data, these unexpected samples will be misidentified as known classes, significantly affecting the performance of semi-supervised learning.

Open-set Recognition. Open set recognition (Geng, Huang, and Chen 2020; Shen et al. 2022; Xie et al. 2021; Huang et al. 2024; Yu et al. 2020; Guo et al. 2020; Huang et al. 2021) treats all unknown classes in unlabeled samples as anomalous, considering the new classes in the open set as low-accuracy outliers to minimize their impact on training. However, it cannot identify different categories within the new samples. Specifically, research in the field of RF fingerprinting open set recognition (Shen et al. 2022; Xie et al. 2021; Huang et al. 2024; Robinson and Kuzdeba 2021) has made notable progress and achieved significant results. In (Xie et al. 2021), the approach extends from a closed-set scenario, where the device set remains constant, to the discrimination of RF fingerprints from unknown devices. In (Huang et al. 2024), Fourier-based synchrosqueezing transform (FSST) and supervised contrastive learning (SCL) are employed to generate more concentrated and distinctive RF fingerprints, reducing open space risk and mitigating the impact of channel noise. In (Robinson and Kuzdeba 2021), the authors propose extensions to the Resampling in Frequency and Time Network (RiftNet) model, enabling simultaneous novel device detection and RF fingerprinting. Although these methods can detect and reject the presence of illegal devices, they cannot classify the illegal devices.

Open-world Recognition. Open-world recognition aims to address the limitations of traditional closed-world models by utilizing an additional, albeit different, labeled set to cluster new, unlabeled classes. In image recognition, ORCA (Cao, Brbic, and Leskovec 2022) uses an uncertainty-aware adaptive margin mechanism to avoid bias towards known classes; OpenNCD (Liu et al. 2023) employs a dual-level contrastive learning method that constructs prototypes and prototype groups to calculate contrastive loss from both group and prototype levels. Notably, OpenNCD's setup only considers 10% of the samples from known classes as labeled data, aligning more closely with real-world scenarios. To tackle the catastrophic forgetting problem of known classes in open-world recognition, LegoGCD (Cao et al. 2024) effectively mitigates this issue using a simple entropy regularization method called LER. However, RF fingerprint signals differ from images, thus requiring distinct open-world recognition methods tailored for RF fingerprinting.

Method

Problem Definition

We divide the entire dataset into two parts: $\mathcal{D}_l$ and $\mathcal{D}_u$. $\mathcal{D}_l = \{(x_i, y_i)\}_{i=1}^{M} \in \mathcal{X} \times \mathcal{Y}_l$ represents the labeled portion, where $\mathcal{Y}_l$ is the set of classes for the labeled samples. $\mathcal{D}_u = \{(x_i, y_i)\}_{i=1}^{N} \in \mathcal{X} \times \mathcal{Y}_u$ represents the unlabeled portion, where $\mathcal{Y}_u$ is the set of classes for the unlabeled samples. In the open-world setting, the unlabeled samples may contain novel classes that are not present in the labeled samples, meaning $\mathcal{Y}_l \subseteq \mathcal{Y}_u$. The set of novel classes is denoted as $\mathcal{Y}_{\text{novel}} = \mathcal{Y}_u / \mathcal{Y}_l$. In our method, the total number of classes $|\mathcal{Y}_u|$ is treated as prior knowledge, consistent with previous works (Zhong et al. 2021; Fini et al. 2021; Han et al. 2021; Zhao and Han 2021; Cao et al. 2024).

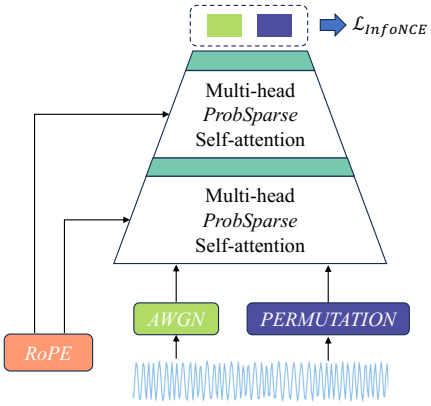

Figure 2: Training of the Roinformer model using the self-supervised learning framework SimCLR. In Roinformer, each Attention operation is preceded by RoPE encoding. The original RF signals undergo two types of data augmentation: AWGN and PERMUTATION, resulting in two positive samples with the same label. For simplicity, some components of SimCLR are omitted.

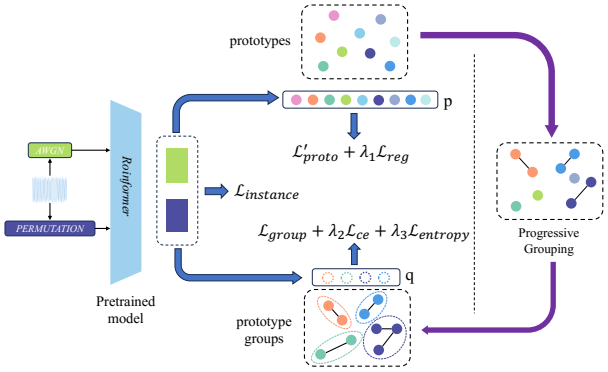

Figure 3: Framework of the proposed OpenRFI method. We incorporate instance-level similarity loss $\mathcal{L}_{instance}$ into features extracted by the encoder to improve recognition accuracy for known class samples. Next, we modify and obtain a new prototype-level similarity loss $\mathcal{L}'_{proto}$. Finally, we introduce prototype group-level similarity loss $\mathcal{L}_{entropy}$ to alleviate catastrophic forgetting of known class samples. The training process continuously optimizes prototype representations, resulting in a more reliable prototype distribution.

Backbone

In this section, we propose a novel RF fingerprint classification model named Roinformer. The Roinformer model adopts the encoder architecture of Informer (Zhou et al. 2021), an advanced time-series forecasting model based on the Transformer (Vaswani et al. 2017) architecture. Informer is designed to address the inefficiencies and accuracy issues of traditional time-series forecasting methods when handling long sequence data. By optimizing and improving upon these challenges, the Informer encoder serves as a powerful feature extractor, providing robust support for open-world recognition. Additionally, recognizing that signal components near the signal segments are more critical than those in other positions, we replace the traditional sinusoidal position encoding with Rotary Position Embedding (RoPE) (Su et al. 2024). RoPE introduces relative positional information of signal segments into the self-attention matrix, offering an advantage over conventional sinusoidal encoding in capturing the importance of signal segments. In summary, Roinformer combines the Informer encoder and RoPE to enhance feature extraction and positional awareness, achieving superior RF fingerprint classification performance.

Pre-trained Model

The pre-trained Roinformer model is trained using the unsupervised SimCLR framework (Chen et al. 2020). SimCLR is a simple contrastive learning framework for learning representation, which simplifies existing contrastive self-supervised learning algorithms by eliminating the need for specialized architectures or memory banks. Since SimCLR is originally designed for images, specific data augmentation methods for RF fingerprints (RFF) are needed. Here, we introduce two effective data augmentation methods for RFFI:

- AWGN (Additive White Gaussian Noise): AWGN is an unavoidable factor in measurement environments, and it inevitably interferes with the signals received by the receiver. By adding slight noise to the signal during augmentation, we simulate real-world conditions, enabling the model to learn the impact of AWGN on RFFI.
- PERMUTATION: This method randomly segments each signal into no more than 10 fragments on the time scale, rearranges them, and recombines them into the original shape (Qian, Tian, and Miao 2022). This enables the model to learn different views of signal representations. After permutation, the signal is further augmented using AWGN. Despite reordering, unique losses from emitter defects remain, as these losses are device-specific. Reordering also increases signal diversity. Thus, permutation is an effective data augmentation method for RFFI.

We utilize the self-supervised learning framework SimCLR to obtain a pre-trained version of Roinformer. The training process is illustrated in Fig. 2. This pre-trained model is used for open-world recognition of emitters in the OpenRFI framework. Experimental results show that allowing all layers of the pre-trained model to be fine-tuned yields better open-world recognition performance compared to fine-tuning only the final few layers. Therefore, in our approach, every layer of the pre-trained model is fine-tunable.

OpenRFI

OpenRFI is an open-world recognition framework for RFFI, based on an advanced modification of OpenNCD (Liu et al. 2023). The open-world recognition process is as follows: To perform an initial coarse partitioning of the sample features, we first predefine a set of trainable prototypes $\mathcal{C} = \{\mathbf{c}_1, \ldots, \mathbf{c}_K\}$, where each prototype is represented as a vector of the same dimension, and the number of prototypes K is predefined, with $K \gg |\mathcal{Y}_u|$. Next, based on the prototype

similarity matrix, the prototypes are further clustered into coarse-grained prototype groups, with each group being at the same hierarchical level as the final class. The number of clusters corresponds to the size of $|\mathcal{Y}_u|$. Finally, the category of each sample is predicted based on the prototype group to which it is most likely to belong. During the training process, the representation of the prototypes is continuously optimized to achieve better prediction results.

The entire open-world recognition process is divided into three stages. The first stage is the Warming stage, where the initial prototype representations are constructed, and some objective functions are excluded from the training. The second stage is the Grouping stage. In this stage, OpenRFI generates a prototype similarity matrix, clusters the prototypes into prototype groups, and obtains $|\mathcal{Y}_u|$ different groups, and this stage also excludes some objective functions from the calculations. The final stage is the Fixing stage, where all objective functions are included in the training.

OpenNCD's loss function consists of four components:

$$\mathcal{L}_{OpenNCD} = \mathcal{L}_{proto} + \mathcal{L}_{group} + \lambda_1 \mathcal{L}_{reg} + \lambda_2 \mathcal{L}_{ce}, \quad (1)$$

where the first two terms represent the prototype-level similarity loss $\mathcal{L}_{proto}$ and the group-level similarity loss $\mathcal{L}_{group}$. $\mathcal{L}_{reg}$ is the regularization loss at the prototype level, and $\mathcal{L}_{ce}$ is the cross-entropy loss for labeled samples.

Samples $\mathbf{x}$ and $\mathbf{x}'$ are positive pairs generated through data augmentation.$\mathbf{x}$ and $\mathbf{x}'$ are processed by the feature extractor f_θ to generate instance features $\mathbf{z}$ and $\mathbf{z}'$.The distribution probability from instance $\mathbf{z}$ to the prototype set $\mathcal{C}$ is represented by $\mathbf{p} \in \mathbb{R}^{(m+n)\times K}$. The distribution probability of instance $\mathbf{z}$ to the k-th prototype $\mathbf{c}_k$ is given by:

$$\mathbf{p}^{(k)} = \frac{\exp\left(\frac{1}{\tau}\mathbf{z}\cdot\mathbf{c}_k^\top\right)}{\sum_{\mathbf{c}_{k'}\in\mathcal{C}}\exp\left(\frac{1}{\tau}\mathbf{z}\cdot\mathbf{c}_{k'}^\top\right)}, \quad (2)$$

where τ is the temperature parameter, and both $\mathbf{z}$ and $\mathbf{c}$ have undergone l_2-normalization. Similarly, the distribution probability of the positive instance $\mathbf{z}'$ to the k-th prototype $\mathbf{c}_k$ is denoted as $\mathbf{p}'^{(k)}$.

The prototype group $\mathcal{C}_g$ is hierarchically aligned with the classes, where $\mathcal{C}_g \subseteq \mathcal{C}$. Prototype groups have no shared prototypes. The distribution probability of instance feature $\mathbf{z}$ to prototype group $\mathcal{C}_g$ is denoted as $\mathbf{q}^{(g)}$:

$$\mathbf{q}^{(g)} = \frac{\sum_{\mathbf{c}_k\in\mathcal{C}_g}\exp\left(\frac{1}{\tau}\mathbf{z}\cdot\mathbf{c}_k^\top\right)}{\sum_{\mathbf{c}_{k'}\in\mathcal{C}}\exp\left(\frac{1}{\tau}\mathbf{z}\cdot\mathbf{c}_{k'}^\top\right)}, \quad (3)$$

where the distribution probability of the positive instance feature $\mathbf{z}'$ to the prototype group $\mathcal{C}_g$ is denoted as $\mathbf{q}'^{(g)}$.

The prototype-level similarity loss $\mathcal{L}_{proto}$ is specifically formulated as follows:

$$\mathcal{L}_{proto} = -\frac{1}{m+n}\sum_{i}^{m+n}\log\langle\mathbf{p}_i, \mathbf{p}'_i\rangle, \quad (4)$$

where $\langle\cdot\rangle$ denotes the Euclidean distance, m and n are the numbers of labeled and unlabeled samples in a training batch, respectively. $\mathcal{L}_{proto}$ is calculated for each batch to update the encoder and prototypes. Treating $\mathbf{p}$ and $\mathbf{p}'$ as pseudo-labels for cross-entropy loss better measures the similarity between the distributions, whereas Euclidean distance performs poorly. Thus, we modify $\mathcal{L}_{proto}$ as follows:

$$\mathcal{L}'_{proto} = -\frac{1}{m+n}\sum_{i}^{m+n}(\mathbf{p}'_i\log\mathbf{p}_i + \mathbf{p}_i\log\mathbf{p}'_i), \quad (5)$$

where $\mathcal{L}'_{proto}$ is also calculated for each training batch to update the encoder and prototypes.

To further enhance fine-grained contrastive learning at the instance level and reduce the adverse effects of pre-trained models on open-world recognition, we introduce the instance-level similarity loss $\mathcal{L}_{instance}$. Here, the InfoNCE loss is applied. The $\mathcal{L}_{instance}$ is defined as follows:

$$\mathcal{L}_{instance} = -\frac{1}{2B}\sum_{i=1}^{B}\left[\log\frac{\exp(\frac{\text{sim}(\mathbf{z}_i,\mathbf{z}'_i)}{\tau_s})}{\sum_{j=1}^{2B}\mathbb{1}_{[j\neq i]}\exp(\frac{\text{sim}(\mathbf{z}_i,\mathbf{z}_j)}{\tau_s})} + \log\frac{\exp(\frac{\text{sim}(\mathbf{z}'_i,\mathbf{z}_i)}{\tau_s})}{\sum_{j=1}^{2B}\mathbb{1}_{[j\neq i]}\exp(\frac{\text{sim}(\mathbf{z}'_i,\mathbf{z}_j)}{\tau_s})}\right], \quad (6)$$

where $\mathbb{1}_{[j\neq i]} \in \{0,1\}$ is an indicator function, $\text{sim}(\cdot,\cdot)$ denotes the cosine similarity function, and τ_s is the temperature parameter. $B = m+n$, where m and n are the number of labeled and unlabeled samples in a training batch, respectively. Unlike unsupervised learning, for a labeled sample in a training batch, the instance feature $\mathbf{z}'_i$ of its positive sample is replaced by a different augmented view $\mathbf{z}'_i$ of the same label sample from the same batch. By this approach, we can inject information from labeled samples, thereby improving the accuracy of known class samples.

With the above modifications, OpenRFI has shown a significant improvement in open-world recognition accuracy. However, an issue of catastrophic forgetting of known classes has emerged. In (Cao et al. 2024), the authors use local entropy regularization on unlabeled samples with high confidence in being identified as known classes to optimize the probability distribution of these samples. However, in the context of RFF data, this approach can also lead to a decrease in the accuracy of unknown class samples.To address this issue, we propose a novel local entropy regularization loss function $L_{entropy}$. The function's impact will be analyzed in the experimental section.

The specific calculation process is as follows: for a training batch B containing m labeled samples and n unlabeled samples, define a binary mask $\mathcal{U}_{mask} = [\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{m+n}] \in [0,1]^{m+n}$ to indicate whether a sample is unlabeled. If sample $\mathbf{x}_i$ is unlabeled, then the corresponding $\mathbf{u}_i = 1$; otherwise, $\mathbf{u}_i = 0$.

Next, we define a binary mask $\mathcal{W}_{mask} = [\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_{m+n}] \in [0,1]^{m+n}$ to indicate whether a sample belongs to a specific prototype group with high confidence. Specifically, for sample $\mathbf{x}_i$, the prototype group distribution probability is $\mathbf{q}_i = [q_1, q_2, \ldots, q_{|\mathcal{Y}_u|}]$.For the elements in $\mathcal{W}_{mask}$, we have:

$$\mathbf{w}_i = \mathbb{1}(\max(\mathbf{q}_i) \geq \gamma), \tag{7}$$

where γ is the threshold for determining whether $\mathbf{x}_i$ is a high-confidence sample.

Since $\mathcal{L}_{ce}$ aligns the predicted labels of known samples with true labels, we can directly use the OpenRFI model's output to determine whether a sample belongs to a known or unknown class. Therefore, a binary mask $\mathcal{Y}_{mask} = [\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_{m+n}] \in [0,1]^{m+n}$ is defined, where:

$$\mathbf{y}_i = \mathbb{1}(\arg\max(\mathbf{q}_i) \notin \mathcal{Y}_l). \tag{8}$$

Finally, a binary mask $S_{\text{mask}} = [\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_{m+n}] \in [0,1]^{m+n}$ is used to indicate whether a sample is considered a high-confidence unknown sample in the unlabeled set. The calculation process is as follows:

$$\mathcal{S}_{mask} = \mathcal{U}_{mask} \odot \mathcal{W}_{mask} \odot \mathcal{Y}_{mask}. \tag{9}$$

where $\odot$ denotes the element-wise multiplication operation.

For a training batch, let the number of unlabeled samples that are identified as high-confidence unknown samples be b. The average prototype group distribution probability of these samples is denoted as:

$$\mathbf{q}_{group} = \frac{1}{b} \sum_i s_i \mathbf{q}_i. \tag{10}$$

We design a prior distribution $\mathbf{q}_{prior}$ to encourage the prototype group is uniformly distributed to achieve the purpose of maximizing entropy:

$$\mathbf{q}_{prior} = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_{|\mathcal{Y}_u|}], \quad \forall \mathbf{q}_i = \frac{1}{|\mathcal{Y}_u|}. \tag{11}$$

We minimize the Kullback-Leibler (KL) divergence to reduce the difference between the two distributions:

$$\mathcal{L}_{entropy} = KL(\mathbf{q}_{group} \parallel \mathbf{q}_{prior}). \tag{12}$$

Since the regularization loss $\mathcal{L}_{entropy}$ requires a certain level of prototype representation to be effective, it is only applied during the Fixing stage of model training.

In summary, the overall objective function of OpenRFI is as follows:

$$\begin{aligned} \mathcal{L}_{OpenRFI} &= \mathcal{L}'_{proto} + \mathcal{L}_{group} + \lambda_1 \mathcal{L}_{reg} + \lambda_2 \mathcal{L}_{ce} \\ &\quad + \mathcal{L}_{instance} + \lambda_3 \mathcal{L}_{entropy} \\ &= \mathcal{L}'_{OpenNCD} + \mathcal{L}_{instance} + \lambda_3 \mathcal{L}_{entropy}. \end{aligned} \tag{13}$$

Experiments

Dataset

To validate the effectiveness of Roinformer, we use a 32-class RF fingerprint dataset for supervised learning. The dataset contains approximately 155,000 samples, with about 124,000 samples in the training set, and roughly 15,000 samples each in the validation and test sets.

Dataset	RFF	UCIHAR	SHAR
Known classes	5	3	8
Total classes	10	6	17
Labeled samples	2.3k	0.5k	0.5k
Unlabeled samples	41k	9.6k	7.7k

Table 1: Datasets used in open-world recognition. Known classes correspond to $|\mathcal{Y}_l|$, total classes correspond to $|\mathcal{Y}_u|$, labeled samples correspond to $|\mathcal{D}_l|$, unlabeled samples correspond to $|\mathcal{D}_u|$.

Following the setup used for the CIFAR-10 in OpenNCD (Liu et al. 2023), we randomly select 10 classes from the 32-class RFF dataset. Half of these classes are designated as known classes, with 10% of the samples from the known classes labeled, while the remaining samples are unlabeled. Additionally, to verify the generalization of our open-world method, we also test OpenRFI's performance on the UCI-HAR dataset, which contains 6 classes of human activity recognition data, and the SHAR dataset, which contains 17 classes of human activity recognition data. The labeling rules for the labeled and unlabeled samples in these datasets are consistent with those used for the 10-class RFF dataset. The distribution of samples in the datasets used for open-world recognition is shown in Table 1.

Implementation Details

When using the self-supervised learning framework SimCLR (Chen et al. 2020) to pre-train a model on a dataset, we split all data into training, validation, and test sets in a 6:2:2 ratio. During training, we implement early stopping to prevent overfitting. The contrastive loss on the validation set is monitored as the key indicator. At the end of each epoch, we calculate the validation loss. If the loss does not decrease for 50 consecutive epochs, early stopping is triggered.

In the OpenRFI setting, the number of prototypes is generally much larger than the number of actual classes. Here, we set the number of prototypes to be ten times the number of known classes. The dimension of the instance features obtained from the pre-trained model is set to 32. The temperature parameter τ_s is set to 0.05. The batch size is set to 128, and the learning rate for the Adam optimizer is 0.002.

Performance of the Feature Extractor

We analyze the performance of Roinformer using a supervised learning approach. Theoretically, Roinformer has a time and space complexity of $O(N \log N)$, which is lower than that of Transformer ($O(N^2)$), precisely because it adopts the Informer model. As shown in Table 2, Roinformer achieves better results with fewer parameters and lower time complexity.

Additionally, the incorporation of RoPE (Su et al. 2024) significantly improves model performance. RoPE does not add model parameters or change the time and space complexity of the model, yet it markedly enhances the performance of both Transformer and Roinformer, speeding up Roinformer's convergence (as can be seen from the Max_epoch parameter). Furthermore, from an application

Model	Acc	Max_epoch	Parameters	FLOPs
Transformer	52.72	57	80.96	25.82
Transformer (w/ RoPE)	94.99	82	80.96	25.82
Informer	61.19	165	51.21	13.98
Roinformer	96.64	57	51.21	13.98

Table 2: Performance comparison and analysis of feature extractors using supervised learning on 32-class RFF dataset. Acc: Accuracy on the test set, Max_epoch: The number of epochs corresponding to model convergence, Parameters ($M\times10^{-2}$): The number of parameters, which reflects the model's spatial complexity, FLOPs ($G\times10^{-3}$): Floating point operations, which reflect the model's time complexity.

perspective, RoPE reduces the result of dot products between Q and distant K positions, effectively applying a windowed approach to the self-attention matrix. Because RoPE can enhance the efficiency of the model, it allows the model to quickly learn the "window" structure that the attention mechanism needs to focus on during training, saving training time. Conversely, this also demonstrates that for RF fingerprinting, the signal segments close to the current position are more important than those further away.

Comparison with the Baselines

In Table 3, we compare the performance of OpenRFI, OpenNCD (Liu et al. 2023), and ORCA (Cao, Brbic, and Leskovec 2022) methods on open-world recognition across three datasets. For all three methods, we use the total number of classes as the prior knowledge input to the model. For the k-means algorithm (MacQueen et al. 1967), we directly cluster the instance features output by the pre-trained model. Since k-means cannot distinguish between known and unknown classes, we treat the top $|\mathcal{Y}_l|$ predicted classes as known and the remaining classes as unknown.As the results show, our method consistently achieves the best performance across the three datasets.

Novel Local Entropy Regularization

LegoGCD (Cao et al. 2024) discovers and mitigates the issue of catastrophic forgetting in known class samples within open-world recognition. In the LegoGCD setup, its local regularization primarily focuses on samples that are highly confident to be known class samples. The difference is that we apply local entropy regularization to samples that are highly confident to be unknown classes. In Fig. 4, we analyze the effect of $\mathcal{L}_{entropy}$ in three scenarios: when applied to samples that are highly confident to be unknown, when applied to samples that are highly confident to be known, and when not applied at all. Fig. 4 shows the accuracy of the total unlabeled samples, known class samples, and unknown class samples over 160 epochs.

As shown in Fig. 4(b), $\mathcal{L}_{entropy}$ significantly alleviates the catastrophic forgetting issue in known class samples, whether applied to samples that are highly confident to be unknown or highly confident to be known. Furthermore, it can be analyzed that $\mathcal{L}_{entropy}$ optimizes the distribution of known class samples mistakenly identified as unknown and known class samples mistakenly identified as other known classes. From both perspectives, $\mathcal{L}_{entropy}$ helps mitigate the catastrophic forgetting issue in known class samples.

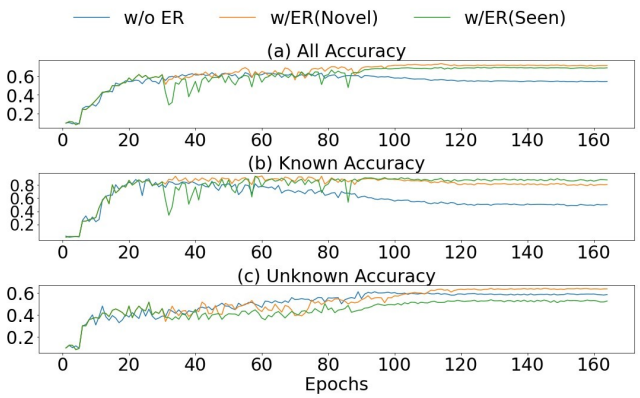

Figure 4: The effect of $\mathcal{L}_{entropy}$ when applied to samples that are highly confident to be unknown classes, when applied to samples that are highly confident to be known classes, and when not applied at all is shown over consecutive epochs. The results are represented by yellow (w/ ER (Novel)), green (w/ ER (Seen)), and cyan (w/o ER) lines, where ER represents $\mathcal{L}_{entropy}$. (a), (b) and (c) show the total accuracy, known class accuracy, and unknown class accuracy. The results were obtained from the RFF dataset.

However, as shown in Fig. 4(c), if local regularization is applied to samples that are highly confident to be known classes, it increases the number of samples mistakenly identified as unknown classes, thereby increasing the difficulty of recognizing unknown classes and leading to a decrease in the accuracy of unknown classes. This impact is more destructive than the effect of local regularization applied to samples that are highly confident to be unknown on the accuracy of known classes.

Additionally, because the identification accuracy of known class samples is relatively high, applying local regularization to samples that are highly confident to be known classes can lead to a "cliff-like" drop in the accuracy of known class samples. This phenomenon is clearly observed in Fig. 5, where, in our setup, $\mathcal{L}_{entropy}$ begins to take effect starting from the 30th epoch.

Ablation Experiment

Impact of Objective Loss Functions. Table 4 analyzes the contributions of our newly added loss functions, including instance-level similarity loss $\mathcal{L}_{instance}$, local entropy regularization at the prototype level $\mathcal{L}_{entropy}$, and the modified prototype-level similarity loss $\mathcal{L}'_{proto}$. To investigate the importance of these functions, we conduct ablation studies by individually removing each item, removing any two items, or removing all the items from the objective function. The results show that all the added objective functions contribute to improving the accuracy of the model. Specifically, $\mathcal{L}_{instance}$ has the most significant effect, $\mathcal{L}_{entropy}$ effectively prevents catastrophic forgetting for known classes, $\mathcal{L}'_{proto}$ markedly improves accuracy for unknown classes.

Methods	RFF			UCIHAR			SHAR		
	Seen	Novel	All	Seen	Novel	All	Seen	Novel	All
*k-means**	16.67	16.39	15.54	45.52	56.92	50.71	24.53	18.54	20.66
ORCA	9.47	25.30	16.11	97.93	86.87	91.94	53.03	17.88	31.06
OpenNCD	55.72	44.19	47.99	98.00	84.09	90.52	53.08	20.40	36.99
OpenRFI	**85.55**	**63.58**	**73.95**	**98.94**	**88.09**	**93.11**	**59.69**	**24.08**	**40.46**
Δ	0.61	0.36	0.48	0.05	0.08	0.06	2.87	1.27	0.88

Table 3: Open-world recognition results on the RFF, UCIHAR, and SHAR datasets. An asterisk (*) indicates that the accuracy of known classes, unknown classes, and overall accuracy is calculated using the Hungarian algorithm; otherwise, the accuracy of known classes is calculated using the conventional method. Bold values indicate the results of our method.

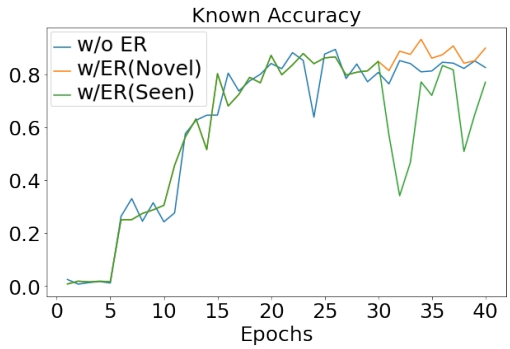

Figure 5: "Cliff-like" accuracy drop occurs when $\mathcal{L}_{entropy}$ is applied to samples highly confident as known classes. It shows accuracy over the first 40 epochs when $\mathcal{L}_{entropy}$ is applied to unknown classes, known classes, or not applied, represented by yellow, green, and cyan lines, respectively.

Impact of the Confidence Threshold. Table 5 shows the ablation study on the confidence threshold γ with the weight $\lambda_3 = 45$ in $\mathcal{L}_{entropy}$. We test the performance of OpenRFI with γ values ranging from 0.5 to 0.9. The results indicate that as γ increases, all three accuracy metrics exhibit an upward trend followed by a decline. The highest known class accuracy of 90.23% occurs at $\gamma = 0.6$, while the highest unknown class accuracy and overall accuracy of 63.58% and 73.95%, respectively, occur at $\gamma = 0.7$. Since open-world recognition places more emphasis on the accuracy of unknown classes, We choose $\gamma = 0.7$ as the optimal value and set it as the default.

Impact of Novel Local Entropy Regularization Weight. Table 6 shows the results of our ablation study on the weight λ_3. We tested λ_3 values from 30 to 50, assessing the performance of OpenRFI. The results show that as λ_3 increases, all three accuracy rates first increase and then decrease. The best overall accuracy and known class accuracy are achieved at $\lambda_3 = 40$, with maximum values of 89.51% and 74.52%, respectively. The best unknown class accuracy is achieved at $\lambda_3 = 45$, with a maximum value of 63.58%. Similarly, since our study focuses more on the accuracy of unknown classes, we chose $\lambda_3 = 45$ as the final weight for OpenRFI in our open-world recognition results.

Methods	Seen	Novel	All
w/o $\mathcal{L}_{instance}$	21.51	39.96	21.10
w/o $\mathcal{L}_{entropy}$	57.80	60.50	59.23
w/o $\mathcal{L}'_{proto}$	89.48	38.41	58.99
w/o $\mathcal{L}'_{proto} + \mathcal{L}_{instance}$	35.79	37.31	19.70
w/o $\mathcal{L}'_{proto} + \mathcal{L}_{entropy}$	73.88	43.94	58.08
w/o $\mathcal{L}_{instance} + \mathcal{L}_{entropy}$	22.02	40.23	21.24
OpenNCD	55.72	44.19	47.99
OpenRFI	**85.55**	**63.58**	**73.95**

Table 4: Analysis of the fusion of novel objective functions. We report results on the 10-class RFF dataset, where 50% of the classes are known, with 10% of the samples from the known classes labeled, and the remaining samples unlabeled. The results without the ablation experiment are indicated in bold.

γ	Seen	Novel	All
0.5	87.85	48.15	62.06
0.6	**90.23**	56.27	72.30
0.7	85.55	**63.58**	**73.95**
0.8	82.78	60.76	71.15
0.9	83.86	49.16	65.54

Table 5: Ablation study on the confidence threshold value conducted on 10-class RFF dataset. Bold indicates the best results. Underlined indicates the selected value.

λ_3	Seen	Novel	All
30	86.79	56.65	70.88
35	83.47	58.87	70.49
40	**89.51**	61.10	**74.52**
45	85.55	**63.58**	73.95
50	88.63	50.82	68.67

Table 6: Ablation study on the novel local entropy regularization weight λ_3 conducted on the 10-class RFF dataset. Bold indicates the best results, and underlined indicates the selected value.

Conclusion

In this paper, we implement an open-world recognition method framework for radio frequency fingerprinting called OpenRFI. To realize it, we develop the Roinformer model for feature extraction, combining it with data augmentation strategies like noise jittering and signal frame reordering to create an effective pre-trained model. Additionally, we incorporate an instance-level similarity loss into the original objective function of the baseline, which effectively improves the accuracy of known classes. In particular, to address the catastrophic forgetting problem of known classes, we introduce a novel local entropy regularization method and profoundly analyze its mechanism, ultimately achieving a better open-world recognition performance.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (U20B2042).

References

Baldini, G.; and Steri, G. 2017. A survey of techniques for the identification of mobile phones using the physical fingerprints of the built-in components. *IEEE Communications Surveys & Tutorials*, 19(3): 1761–1789.

Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32.

Cao, K.; Brbic, M.; and Leskovec, J. 2022. Open-World Semi-Supervised Learning. arXiv:2102.03526.

Cao, X.; Zheng, X.; Wang, G.; Yu, W.; Shen, Y.; Li, K.; Lu, Y.; and Tian, Y. 2024. Solving the Catastrophic Forgetting Problem in Generalized Category Discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16880–16889.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.

Deng, P.; Hong, S.; Qi, J.; Wang, L.; and Sun, H. 2023. A lightweight transformer-based approach of specific emitter identification for the automatic identification system. *IEEE Transactions on Information Forensics and Security*, 18: 2303–2317.

Fini, E.; Sangineto, E.; Lathuilière, S.; Zhong, Z.; Nabi, M.; and Ricci, E. 2021. A unified objective for novel class discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9284–9292.

Geng, C.; Huang, S.-j.; and Chen, S. 2020. Recent advances in open set recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(10): 3614–3631.

Guan, D.; Xing, Y.; Huang, J.; Xiao, A.; El Saddik, A.; and Lu, S. 2025. S2Match: Self-paced sampling for data-limited semi-supervised learning. *Pattern Recognition*, 159: 111121.

Guo, L.-Z.; Zhang, Z.-Y.; Jiang, Y.; Li, Y.-F.; and Zhou, Z.-H. 2020. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International conference on machine learning*, 3897–3906. PMLR.

Han, K.; Rebuffi, S.-A.; Ehrhardt, S.; Vedaldi, A.; and Zisserman, A. 2021. Autonovel: Automatically discovering and learning novel visual categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6767–6781.

Hua, J.; Sun, H.; Shen, Z.; Qian, Z.; and Zhong, S. 2018. Accurate and efficient wireless device fingerprinting using channel state information. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, 1700–1708. IEEE.

Huang, J.; Fang, C.; Chen, W.; Chai, Z.; Wei, X.; Wei, P.; Lin, L.; and Li, G. 2021. Trash to treasure: Harvesting ood data with cross-modal matching for open-set semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8310–8319.

Huang, R.; Peng, X.; Chai, Z.; Li, M.; Ren, J.; and Yang, X. 2024. Radio frequency fingerprint extraction and authentication towards open set in noisy channels. *Digital Signal Processing*, 146: 104363.

Li, J.; Xiong, C.; and Hoi, S. C. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9475–9484.

Liu, J.; Wang, Y.; Zhang, T.; Fan, Y.; Yang, Q.; and Shao, J. 2023. Open-world semi-supervised novel class discovery. *arXiv preprint arXiv:2305.13095*.

Liu, P.; Yang, P.; Song, W.-Z.; Yan, Y.; and Li, X.-Y. 2019. Real-time identification of rogue WiFi connections using environment-independent physical features. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, 190–198. IEEE.

MacQueen, J.; et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, 281–297. Oakland, CA, USA.

Polak, A. C.; and Goeckel, D. L. 2015. Wireless device identification based on RF oscillator imperfections. *IEEE Transactions on Information Forensics and Security*, 10(12): 2492–2501.

Qian, H.; Tian, T.; and Miao, C. 2022. What makes good contrastive learning on small-scale wearable-based tasks? In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, 3761–3771.

Robinson, J.; and Kuzdeba, S. 2021. Novel device detection using RF fingerprints. In *2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)*, 0648–0654. IEEE.

Shen, G.; Zhang, J.; Marshall, A.; and Cavallaro, J. R. 2022. Towards scalable and channel-robust radio frequency fingerprint identification for LoRa. *IEEE Transactions on Information Forensics and Security*, 17: 774–787.

Shen, G.; Zhang, J.; Marshall, A.; Valkama, M.; and Cavallaro, J. R. 2023. Toward length-versatile and noise-robust radio frequency fingerprint identification. *IEEE Transactions on Information Forensics and Security*, 18: 2355–2367.

Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33: 596–608.

Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised data augmentation for consistency training. *Advances in neural information processing systems*, 33: 6256–6268.

Xie, R.; Xu, W.; Chen, Y.; Yu, J.; Hu, A.; Ng, D. W. K.; and Swindlehurst, A. L. 2021. A generalizable model-and-data driven approach for open-set RFF authentication. *IEEE Transactions on Information Forensics and Security*, 16: 4435–4450.

Yu, Q.; Ikami, D.; Irie, G.; and Aizawa, K. 2020. Multi-task curriculum framework for open-set semi-supervised learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 438–454. Springer.

Zhang, J.; Woods, R.; Sandell, M.; Valkama, M.; Marshall, A.; and Cavallaro, J. 2021. Radio frequency fingerprint identification for narrowband systems, modelling and classification. *IEEE Transactions on Information Forensics and Security*, 16: 3974–3987.

Zhao, B.; and Han, K. 2021. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. *Advances in Neural Information Processing Systems*, 34: 22982–22994.

Zhong, Z.; Zhu, L.; Luo, Z.; Li, S.; Yang, Y.; and Sebe, N. 2021. Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9462–9470.

Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.