Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Reliability-Guided Spectrum Generative Augmentation for Few-Shot Emitter Identification in Mobile Networks

Wankang Zhang
Henan Institute of Science and Technology
Zhengwei Xu
xuzhengweihhu@outlook.com
Henan Normal University
Tao Li
Brunel University of London

Research Article

Keywords: Specific emitter identification, electromagnetic spectrum intelligence, next-generation mobile networks, few-shot learning, generative augmentation, reliability calibration

Posted Date: July 15th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10236582/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Reliability-Guided Spectrum Generative Augmentation for Few-Shot Emitter Identification in Mobile Networks

Wankang Zhang[1], Zhengwei Xu[2*], Tao Li[3]

[1]School of Computer Science and Technology, Henan Institute of Science and Technology, Xinxiang, 453003, China.
[2*]College of Computer and Information Engineering, Henan Normal University, Xinxiang, 453007, China.
[3]Brunel University of London, London, United Kingdom.

*Corresponding author(s). E-mail(s): Zhengwei Xu xuzhengweihhu@outlook.com;
Contributing authors: Wankang Zhang wkzhist@163.com; Tao Li Tao.Li@brunel.ac.uk;

Abstract

With the evolution of next-generation mobile networks, the electromagnetic spectrum is becoming an interactive operating environment that simultaneously carries communication, sensing, and autonomous control functions. In integrated sensing and communication systems, low-altitude mobile networks, and self-organizing wireless infrastructures, reliable recognition of active radio emitters is essential for spectrum awareness and secure coordination. Specific emitter identification (SEI) supports this requirement by extracting hardware-dependent radio-frequency fingerprints from received signals. Nevertheless, dynamic spectrum operation usually provides only a few labeled observations for newly accessed or intermittently visible emitters, which makes conventional deep SEI models prone to unstable decision boundaries. This paper proposes a reliability-guided spectrum generative augmentation framework, named RSGA-SEI, for few-shot SEI in embodied spectrum intelligence. RSGA-SEI maps baseband observations into contour stella image samples, uses spectrum-aware generation to synthesize auxiliary spectrum samples, and adopts reliability-weighted classification to regulate the contribution of generated data. Different from prototype-geometry optimization, the classifier forms prototypes only from measured support samples and weights synthetic samples according to their feature

consistency with the real support set. The resulting design connects spectrum perception, sample expansion, emitter recognition, and adaptive spectrum action in a compact closed-loop workflow. The 5-shot evaluation reaches an accuracy of 87.14% while keeping the online inference path compact.

Keywords: Specific emitter identification, electromagnetic spectrum intelligence, next-generation mobile networks, few-shot learning, generative augmentation, reliability calibration

1 Introduction

Next-generation mobile networks are changing the role of the electromagnetic spectrum from a passive carrier of wireless signals into an active environment where communication, sensing, and intelligent operation are coupled. Integrated sensing and communication studies show how sensing and transmission functions can share spectrum resources [1]. Dynamic spectrum interaction in unmanned aerial vehicle communication further shows that radio perception and network control can be coupled in mobile systems [2]. In such settings, spectrum occupancy, node behavior, propagation dynamics, and resource allocation are no longer independent factors. Instead, they form a feedback process in which perception affects action, and action reshapes the observed radio environment.

This closed-loop nature creates new security and management requirements for mobile systems. Wireless nodes must not only detect whether spectrum resources are available, but also infer which emitters are present and whether their radio fingerprints match trusted identities. SEI is a physical-layer technique for this purpose because it distinguishes transmitters through device-dependent radio-frequency fingerprints (RFFs) caused by manufacturing imperfections and hardware nonlinearities [3]. Lightweight SEI models and dataset-oriented studies further show that RFFs can support spectrum-aware authentication when network entities operate in open or rapidly changing environments [4, 5]. RFF-based recognition also complements protocol identifiers and credential-based access control [6, 7].

Deep learning has substantially improved SEI by enabling feature extraction from raw waveforms, transformed signal representations, and constellation-like images [8]. Large-scale radio signal recognition studies further show that image-like signal representations can capture device-dependent patterns in practical datasets [9]. Representative studies have combined time-frequency analysis with residual learning [10], used dense convolutional models to represent complex-valued signal structures [11], and designed multi-branch networks to jointly capture modulation and device-specific features [12]. However, these methods are usually trained under the assumption that enough labeled samples are available for every emitter. This assumption is often fragile in embodied spectrum operation, where mobile nodes may appear briefly, low-altitude platforms may change channels frequently, and newly joined devices may provide only

a small support set. Under such few-shot conditions, a high-capacity classifier can easily fit incidental noise, channel distortion, or transient environment patterns rather than stable RFF characteristics [13, 14].

Generative augmentation can reduce this sample shortage. By synthesizing additional signal representations, augmentation methods can enlarge the effective training set and provide more support for discriminative learning [15, 16]. Recent studies have explored adversarial and diffusion-based strategies for limited-sample radio-frequency recognition [17, 18]. Image-mapping augmentation has also been studied for CSI-based SEI [19]. Contour stella image (CSI) mapping is relevant here because it converts in-phase and quadrature sampling distributions into visual patterns that preserve amplitude, phase, and density information. Yet generative augmentation also introduces a risk for closed-loop spectrum intelligence because unreliable synthetic samples may be treated as fully trusted observations. The classifier can then learn biased emitter boundaries and feed incorrect identity information into later spectrum decisions.

To address this issue, this paper develops RSGA-SEI, a reliability-guided augmentation framework for few-shot SEI in next-generation mobile networks. The proposed framework is built around an embodied spectrum loop rather than a purely offline image-classification pipeline. First, baseband observations are transformed into CSI samples to describe the radio state perceived by a receiver. Second, the spectrum-aware generative augmentation module (SAGA) generates auxiliary CSI samples by combining local contour priors with global spectrum-context priors. Third, the spectrum-embedded reliability-weighted classifier (SE-RWC) evaluates the reliability of generated samples through prototypes computed only from measured support samples and trains a lightweight classifier with reliability-weighted supervision. The design keeps the deployment path simple because generation and reliability scoring are used during training, while inference only requires the trained classifier.

The main contributions are listed below.

- We formulate few-shot SEI as a spectrum-embedded recognition problem for next-generation mobile systems, linking emitter identification with perception, decision, and action feedback.
- We propose RSGA-SEI, a generative augmentation framework that converts baseband observations into CSI samples and expands the support set through spectrum-aware synthesis.
- We design SE-RWC, a lightweight reliability calibration stage that uses real support prototypes to assign soft weights to generated samples before classifier training.
- We evaluate RSGA-SEI through 5-shot comparison, ablation, multi-shot generalization, signal-to-noise ratio robustness, confusion-matrix analysis, and complexity indicators.

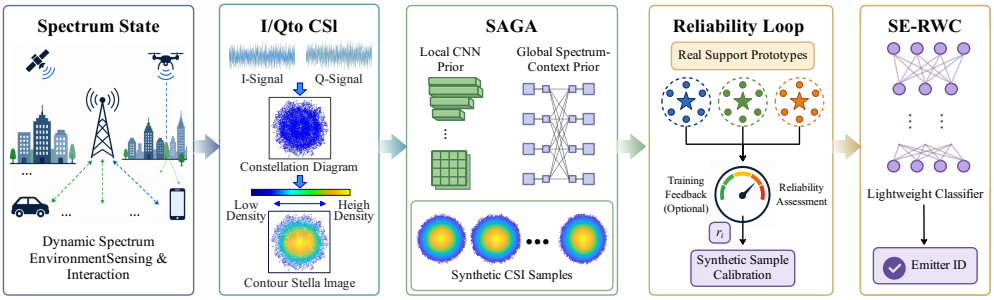


Fig. 1 Closed-loop workflow of RSGA-SEI for few-shot SEI in embodied spectrum intelligence. Spectrum observations are converted into CSI samples, expanded by SAGA, calibrated through a reliability loop, and used by SE-RWC to support emitter identification and adaptive spectrum operation.

2 Related Work

2.1 Emitter Identification in Intelligent Spectrum Systems

SEI supports radio access security by associating received signals with device-specific hardware fingerprints. Early SEI studies relied on expert-designed descriptors extracted from transient responses, bispectrum statistics, and constellation patterns [20, 21]. Transform-domain residual networks and multi-domain feature learning later connected these descriptors with deeper neural representations [10, 11]. These descriptors are interpretable and computationally lightweight, but they often depend on carefully selected signal segments and may become brittle when the channel, receiver, or signal format changes. Deep neural models have therefore become widely used because they can learn hierarchical RFF features from raw waveforms and differential constellation traces [22, 23]. CSI and complex-valued self-supervised models further show that signal images can encode device-dependent patterns for classification [24, 25].

The role of SEI is also changing with the evolution of next-generation mobile networks. In conventional offline recognition tasks, SEI is mainly evaluated as a closed-set classification problem. In intelligent spectrum systems, however, emitter identity can affect spectrum access, interference avoidance, cooperative sensing, and security monitoring. This is especially relevant to integrated sensing and communication and autonomous wireless infrastructures, where spectrum perception and network action form a feedback process [1, 2]. In such environments, recognition errors are not isolated label mistakes; they may influence subsequent resource allocation or coordination decisions. Therefore, SEI methods for embodied spectrum intelligence should not only pursue high accuracy, but also remain robust under limited labels, changing channels, and uncertain sample quality.

2.2 Limited-Sample Learning for RFFs

Few-shot radio-frequency recognition has received increasing attention because newly deployed devices, intermittent emitters, and non-cooperative transmitters may provide only a few labeled observations. Metric-learning methods seek a feature space where samples from the same emitter are compact and different emitters remain separable [14, 26]. Meta-learning and transfer-learning methods attempt to adapt quickly to new emitters or domains by learning reusable initialization or pseudo labels [27, 28]. Task-level adaptation has also been used to improve sample efficiency in SEI [29]. Self-supervised methods reduce dependence on labels by exploiting signal transformations and instance-level consistency [30, 31]. Contrastive and cross-domain approaches further improve representation transfer under limited labels [32, 33].

These strategies improve sample efficiency, but few-shot SEI remains difficult because RFF cues are weak, local, and easily entangled with channel or receiver effects. Class-reconstruction and adversarial training can strengthen the decision boundary when labels are scarce [34], while augmentation-based few-shot learning can enlarge the support set without changing the online classifier [35]. However, the assumptions behind ordinary visual few-shot learning do not transfer perfectly to RFFs. A transformed or synthesized signal sample must preserve physical device characteristics, not only visual similarity. This requirement motivates methods that explicitly consider whether additional training samples are trustworthy enough to affect the classifier.

2.3 Generative Spectrum Augmentation

Generative augmentation offers a direct way to relieve label scarcity by synthesizing additional radio-frequency samples or signal images. Adversarial generators have been used to expand small radiation-source datasets [15, 17]. Diffusion models have recently been explored for radio-frequency recognition and modulation-related tasks [16, 18]. Compared with simple geometric transformations, generative models can provide richer intra-class variation when the measured support set covers only a narrow portion of the signal distribution. At the same time, generative augmentation can introduce distribution shift if the generator produces visually plausible but physically inconsistent samples.

CSI-based SEI matches this augmentation setting. By mapping in-phase and quadrature samples into two-dimensional density patterns, CSI preserves amplitude, phase, and distribution information while allowing the use of mature image backbones [19]. Existing CSI augmentation methods have shown strong performance under limited samples, but most of them still treat generated examples as ordinary labeled data after synthesis. In closed-loop spectrum intelligence, unreliable generated samples may bias the learned emitter boundary and propagate incorrect identity information to downstream spectrum decisions.

2.4 Reliability-Aware Spectrum Learning

The above studies indicate that data expansion alone is not sufficient for embodied spectrum scenarios. A deployable few-shot SEI model should decide not only what synthetic samples to generate, but also how much each generated sample should be

trusted during classifier learning. Reliability estimation is therefore closely related to prototype learning, sample reweighting, and uncertainty-aware training. Real support samples provide the most credible anchor, while synthetic samples should contribute according to their consistency with this anchor. The proposed RSGA-SEI adopts this reliability-aware strategy by combining SAGA with SE-RWC. Unlike heavy prototype-geometry optimization, the proposed design keeps the deployment path compact and calibrates synthetic samples only during training, which better fits mobile spectrum systems that require fast online emitter identification.

3 Methodology

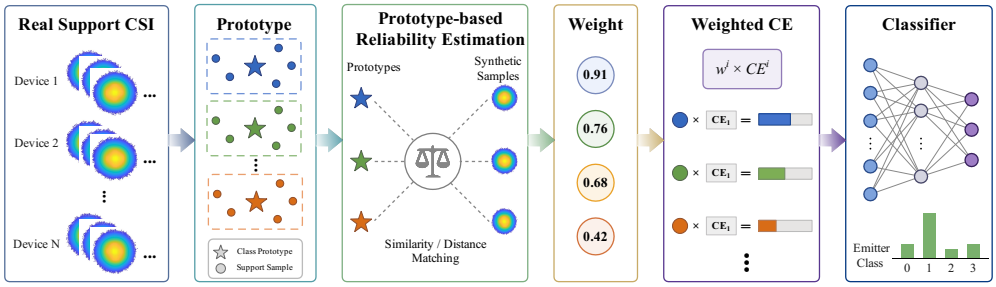


Fig. 2 Reliability weighting in SE-RWC. Real support CSI samples define spectrum prototypes. Synthetic CSI samples generated by SAGA receive soft reliability weights before contributing to weighted cross-entropy training, and the resulting classifier supports downstream spectrum action.

3.1 Embodied Spectrum Learning Formulation

The proposed RSGA-SEI framework is designed for a training scenario in which only a small number of labeled observations are available for each emitter. Let $\mathcal{D}^r = \{(x_i^r, y_i^r)\}$ denote the measured support set, where x_i^r is a radio observation and y_i^r is the emitter label. The objective is to train an SEI model that can classify unseen observations while remaining suitable for a spectrum perception, decision, and action loop. The training process contains CSI visualization, support expansion using SAGA, and reliability-weighted classifier learning using SE-RWC.

3.2 CSI Representation of Spectrum Observations

For a received baseband signal, the in-phase and quadrature components are written as

$$\begin{aligned} I(n) &= A(n)\cos(\varphi(n)), \\ Q(n) &= A(n)\sin(\varphi(n)), \end{aligned} \tag{1}$$

where $A(n)$ is the instantaneous amplitude and $\varphi(n)$ is the instantaneous phase. The resulting samples are projected onto a two-dimensional constellation plane. Let $d(u, v)$

denote the number of projected samples falling into grid cell (u, v), and let Ω_{ij} be the sliding window centered at cell (i, j). The local density value is computed by normalizing the windowed count as

$$\rho(i,j) = \frac{\sum_{(u,v)\in\Omega_{ij}} d(u,v)}{\sum_{i'=W_0}^{W_1} \sum_{j'=H_0}^{H_1} \sum_{(u,v)\in\Omega_{i'j'}} d(u,v) + \epsilon_\rho}, \tag{2}$$

where ϵ_ρ avoids division by zero. The normalized density map is color-coded to form a CSI sample. This representation allows the classifier to observe geometric distribution, local density, and phase-related patterns in a unified image form.

3.3 SAGA Module

SAGA expands the measured support set by generating auxiliary CSI samples for each emitter. It contains a spectrum-context prior encoder (SCPE), neighbor-view diversity regularization (NVDR), and dual-prior consistency discrimination (DPCD).

3.3.1 SCPE module

The SCPE module provides the feature basis for the generative stage. It contains a local branch and a context branch. The local branch extracts contour, density, and neighborhood texture cues from CSI samples. The context branch summarizes broader spatial relationships that reflect the global shape of the perceived spectrum state. For an input CSI image X, the prior features are

$$q_l = P_l(F_l(X)), \quad q_c = P_c(F_c(X)), \tag{3}$$

where $F_l(\cdot)$ and $F_c(\cdot)$ denote the local and context extractors, and $P_l(\cdot)$ and $P_c(\cdot)$ are projection heads. The fused prior embedding is

$$e(X) = P_e([q_l; q_c]). \tag{4}$$

The fusion is implemented by concatenating the local and context embeddings and then applying the projection head P_e. No attention block or prototype alignment is introduced in this fusion step.

3.3.2 NVDR term

The purpose of NVDR is to prevent the generator from creating redundant CSI images while keeping the generated samples near the same emitter identity. Given an emitter label y, a latent code $z \sim \mathcal{N}(0, I)$, and a small disturbance $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$, the class-conditional generator creates two neighboring views as

$$A = T_a(G(z, y)), \quad B = T_b(G(z + \epsilon, y)), \tag{5}$$

where T_a and T_b denote mild image-plane transformations, such as small shifts, weak rotations, and intensity perturbations. Their embeddings are used to compute

a normalized cross-correlation matrix as

$$C_{ij} = \frac{\sum_b e(A)_{b,i} e(B)_{b,j}}{\sqrt{\sum_b e(A)_{b,i}^2} \sqrt{\sum_b e(B)_{b,j}^2}}. \tag{6}$$

The regularization term is

$$\mathcal{L}_{\mathrm{NVD}} = \sum_i (1 - C_{ii})^2 + \sum_i \sum_{j \neq i} C_{ij}^2. \tag{7}$$

The diagonal part keeps the paired views semantically aligned, and the off-diagonal part discourages repeated feature dimensions from carrying the same response.

3.3.3 DPCD strategy

The DPCD strategy constrains generated CSI samples with both local and context-aware discriminators. Let D_l and D_c be the two discriminator heads. For real samples, q_l^X and q_c^X are obtained from the measured CSI image X. For generated samples, q_l^G and q_c^G are obtained from $G(z, y)$ through the same SCPE branches. Their hinge loss is

$$\begin{aligned} \mathcal{L}_D = & \mathbb{E}_{(X,y)} \left[\max(0, 1 - D_l(q_l^X, y)) \right] \\ & + \mathbb{E}_{(X,y)} \left[\max(0, 1 - D_c(q_c^X, y)) \right] \\ & + \mathbb{E}_{z,y} \left[\max(0, 1 + D_l(q_l^G, y)) \right] \\ & + \mathbb{E}_{z,y} \left[\max(0, 1 + D_c(q_c^G, y)) \right], \end{aligned} \tag{8}$$

where superscripts X and G denote real and generated CSI samples, respectively. The adversarial objective for the generator is

$$\mathcal{L}_G = -\mathbb{E}_{z,y} \left[D_l(q_l^G, y) + D_c(q_c^G, y) \right]. \tag{9}$$

To reduce disagreement between the two prior views, the consistency penalty is

$$\mathcal{L}_{\mathrm{PC}} = \mathbb{E}_{z,y} \left(D_l(q_l^G, y) - D_c(q_c^G, y) \right)^2. \tag{10}$$

The generator is trained by

$$\mathcal{L}_{SAGA} = \mathcal{L}_G + \lambda_{\mathrm{nvd}} \mathcal{L}_{\mathrm{NVD}} + \lambda_{\mathrm{pc}} \mathcal{L}_{\mathrm{PC}}. \tag{11}$$

3.4 SE-RWC Module

The SE-RWC stage turns generated samples into calibrated training evidence. Instead of forcing all samples into a fixed geometric prototype arrangement, SE-RWC uses measured support samples as the trust anchor. Synthetic samples are helpful only when their features remain compatible with this real support structure. The module is illustrated in Fig. 2.

A feature extractor $f_\theta(\cdot)$ maps a CSI image into an embedding. For class c, the prototype computed from measured support samples is

$$\mu_c = \frac{1}{|\mathcal{D}_c^r|} \sum_{(x_i^r, y_i^r) \in \mathcal{D}_c^r} f_\theta(x_i^r), \tag{12}$$

where $\mathcal{D}_c^r$ is the support subset belonging to class c. The support dispersion is estimated as

$$\sigma_c^2 = \frac{1}{|\mathcal{D}_c^r|} \sum_{(x_i^r, y_i^r) \in \mathcal{D}_c^r} \|f_\theta(x_i^r) - \mu_c\|_2^2 + \delta, \tag{13}$$

where δ is a small constant.

For a generated sample (x_i^s, y_i^s), the reliability value is

$$r_i = \exp\left(-\frac{\left\| f_\theta(x_i^s) - \mu_{y_i^s} \right\|_2^2}{\gamma \sigma_{y_i^s}^2} \right), \tag{14}$$

where γ controls the smoothness of the decay. The final training weight is bounded by

$$w_i^s = \min(1, \max(r_{\min}, r_i)). \tag{15}$$

Measured support samples keep the full weight $w_i^r = 1$. For a mini-batch $\mathcal{B}$ drawn from the union of measured and generated CSI samples, the classifier is optimized with

$$\mathcal{L}_{\text{SE-RWC}} = -\frac{1}{\sum_{(x_i, y_i) \in \mathcal{B}} w_i} \sum_{(x_i, y_i) \in \mathcal{B}} w_i \log p_\theta(y_i | x_i). \tag{16}$$

This loss retains the simplicity of supervised classification while allowing the training process to distrust weak synthetic evidence.

Algorithm 1 Training Procedure of RSGA-SEI

Require: Measured few-shot set $\mathcal{D}^r$, class number N, generated samples per class K.
Ensure: Lightweight emitter classifier $p_\theta(y|x)$.

```
Transform radio observations into CSI samples.
Update D_l and D_c with L_D, and update G and SCPE with L_SAGA.
Generate D^s = {(G(z_{c,k}, c), c) | c = 1, ..., N; k = 1, ..., K}.
Warm up f_θ using measured support CSI samples.
Calculate support prototypes {μ_c}_{c=1}^N and dispersions {σ_c^2}_{c=1}^N.
for each generated sample (x_i^s, y_i^s) ∈ D^s do
    Estimate r_i and obtain the clipped weight w_i^s.
end for
Train p_θ(y|x) with L_SE-RWC on D^r ∪ D^s. return p_θ(y|x).
```

4 Experimental Evaluation

4.1 Dataset and Evaluation Setting

The evaluation uses the NJUPT-RFF dataset [5], which contains radio-frequency signals from seven emitters collected by laboratory instruments. The signals were acquired in a controlled radio-frequency environment using dedicated transmitting and receiving instruments, and each raw observation contains 200,000 sampling points. Following the few-shot protocol used in the comparison records, each emitter contributes 5 labeled samples for training and 50 samples for testing. Therefore, each trial contains 350 test decisions, and one correctly classified sample corresponds to an accuracy step of 0.2857%. All radio observations are converted into CSI images and resized to 224×224. The online recognition backbone is ResNet-18. The generator and reliability scoring modules are used only during training; the inference path contains only the feature extractor and classifier. For multi-shot and stability analyses, the same support-query construction is repeated at each shot number so that changes in accuracy come from the support size rather than from a different test protocol.

4.2 Implementation Configuration

Table 1 summarizes the configuration adopted for the experimental setting. The generative stage uses Adam optimization and follows the same training duration as the augmentation baseline. For SE-RWC, the reliability smoothness is set to $\gamma = 1.0$, and the minimum generated-sample weight is fixed to $r_{\min} = 0.2$. The generative stage is trained before classifier learning, after which 100 synthetic CSI samples are produced for each emitter and combined with the real support set. This setting keeps the online model unchanged while the training process uses reliability-calibrated synthetic evidence. The reliability parameters are selected from a preliminary validation sweep and kept fixed in all subsequent comparisons.

4.3 Comparison With Few-Shot SEI Methods

Table 2 summarizes the main 5-shot comparison. The cited baseline values are reported under the same NJUPT-RFF few-shot protocol. The CSI baseline, SAGA alone, and RSGA-SEI use the 5-shot support and 50-sample query setting described above. The RSGA-SEI result is 87.14%, corresponding to 305 correctly classified samples among 350 test decisions. The margin over CSI-DA is small because SE-RWC changes the training weights of generated samples rather than adding an extra online inference module.

Compared with the non-augmented CSI baseline, RSGA-SEI improves the 5-shot accuracy by 13.71 percentage points, indicating that spectrum-aware sample expansion is necessary when only a handful of emitter observations are available. Compared with SAGA alone, the gain is 3.43 percentage points, which indicates that generated CSI samples should not be treated as equally reliable training samples. The improvement over CSI-DA is smaller, at 0.28 percentage points, but it is obtained without adding any online inference module. The remaining gap to the sufficient-sample upper

Table 1 Hyperparameter configuration of RSGA-SEI

Parameter	Value
SAGA: spectrum-aware generation	
Learning rate (G/D)	2×10^{-4}
Latent dimension	64
Training duration	25 kimg
Neighbor-view weight λ_{nvd}	5×10^{-3}
Prior-consistency weight λ_{pc}	1.0
SE-RWC: reliability-weighted learning	
Synthetic samples per class	100
Classifier learning rate	2×10^{-4}
Training epochs	100
Reliability smoothness γ	1.0
Minimum synthetic weight $r_{\min}$	0.2

Table 2 Few-shot SEI comparison on the NJUPT-RFF dataset

Method	Training principle	Accuracy (%)
DMEL [26]	Metric ensemble learning	61.71
CRCN-AT [34]	Reconstruction and adversarial training	67.71
DAFS [35]	Augmentation-aided few-shot learning	69.71
CSI baseline	CSI input with ResNet-18	73.43
SAGA alone	Spectrum-aware generation with cross-entropy	83.71
CSI-DA [19]	CSI augmentation with classifier	86.86
RSGA-SEI	Reliability-guided generated-sample calibration	87.14
Upper bound	Sufficient-sample classifier	95.71

bound is 8.57 percentage points, indicating that the few-shot setting still leaves room for stronger sample-efficient representation learning.

4.4 Ablation Study

The ablation study evaluates whether generated samples should be used directly or calibrated before classifier training. As shown in Fig. 3, direct SAGA augmentation improves the CSI baseline from 73.43% to 83.71%. In this variant, all generated samples are assigned unit cross-entropy weights. Hard reliability filtering discards generated samples whose reliability is below $r_{\min}$ and assigns unit weights to the remaining generated samples, which raises the accuracy to 84.86%. Soft reliability weighting keeps all generated samples and applies the clipped reliability weight w_i^s, reaching 85.71%. The full RSGA-SEI setting further uses balanced mini-batches so that measured support samples and generated samples from each emitter are both

represented during classifier training, obtaining the best result of 87.14%. The ablation has two implications. First, the large gain from CSI to CSI+SAGA confirms that the main bottleneck is the scarcity of labeled support samples. Second, the gradual improvement from hard filtering to soft weighting and then to the complete SE-RWC indicates that reliability should be used as a graded training signal rather than as a binary sample-selection rule.

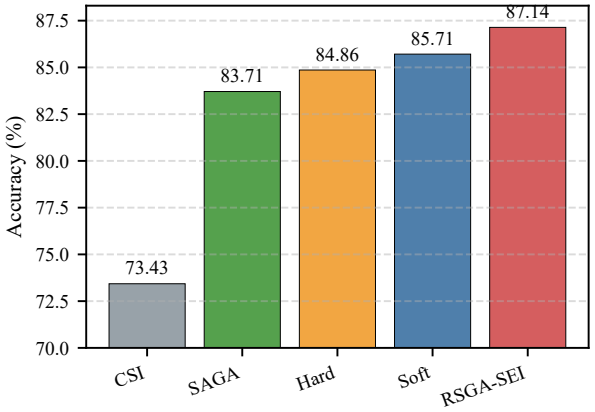


Fig. 3 Ablation analysis of RSGA-SEI under the 5-shot setting.

4.5 Multi-Shot Generalization and Stability

Fig. 4 extends the evaluation from 5-shot to 30-shot settings. All methods improve as more labeled support samples become available. Compared with the plain CSI baseline, the gain is largest in the low-shot region, where synthetic support samples directly compensate for missing observations. Compared with CSI-DA, the margin becomes clearer from the 15-shot to 25-shot settings because the real support prototypes become more stable and reliability weighting can better distinguish consistent synthetic samples from misleading ones. At 30-shot, the accuracy reaches 94.86%, still below the sufficient-sample upper bound of 95.71%.

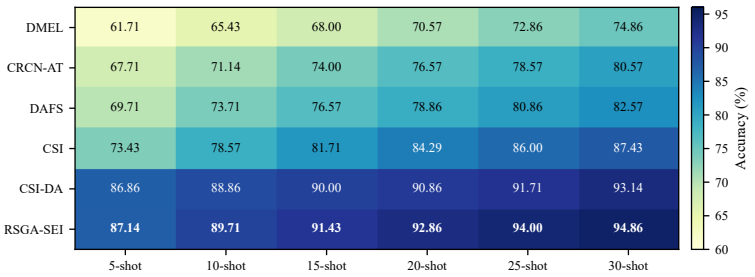


Fig. 4 Multi-shot accuracy heatmap on the NJUPT-RFF dataset.

The stability indicators in Table 3 follow the same trend. As the support set grows, Cohen's Kappa increases from 0.850 to 0.940, while the coefficient of variation (CV) decreases from 0.0134 to 0.0046. This behavior reflects the greater stability of prototype reliability estimates when measured support samples are less sparse. The accuracy and stability results therefore support the same conclusion that RSGA-SEI provides its largest benefit when the model can combine generated samples with a small but increasingly reliable real support set.

Table 3 Stability indicators across support sizes

Metric	5	10	15	20	25	30
Cohen's Kappa	0.850	0.880	0.900	0.918	0.930	0.940
CV	0.0134	0.0100	0.0079	0.0057	0.0052	0.0046

4.6 Signal-to-noise ratio robustness

Embodied spectrum intelligence must operate in changing radio conditions, so signal-to-noise ratio (SNR) robustness is also considered. Fig. 5 reports the SNR trend. The proposed method keeps the highest accuracy at each SNR point because reliability weighting prevents highly distorted generated samples from dominating the classifier. At 20 dB, the value is tied to the main 5-shot result of 87.14%. The advantage over CSI-DA is more visible in the low and medium SNR range, reaching 2.29, 2.57, and 3.14 percentage points at -5, 0, and 5 dB, respectively. Noisy observations make synthetic-sample quality less uniform, so reliability weighting has a larger role in suppressing harmful generated samples. As the channel becomes cleaner, all augmentation-based methods approach their high-SNR performance and the margin naturally narrows.

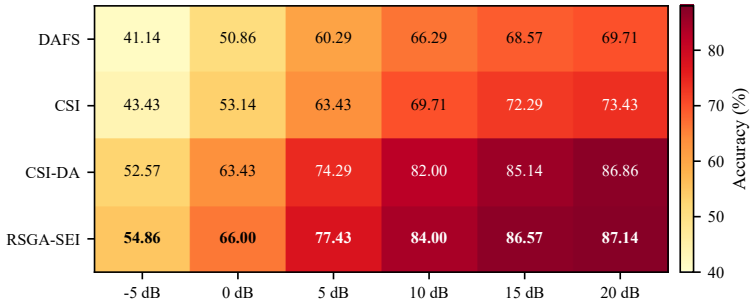


Fig. 5 SNR robustness heatmap under the 5-shot setting.

4.7 Confusion Matrix Analysis

Fig. 6 shows a 5-shot confusion matrix profile for RSGA-SEI. The diagonal entries sum to 305, giving $305/350 = 87.14\%$. The off-diagonal errors are distributed across

several nearby classes rather than concentrated in a single dominant confusion pair, which is consistent with the role of reliability calibration in class-boundary stability. Most classes retain at least 43 correct predictions out of 50, while the most difficult class still keeps 42 correct predictions. This distribution indicates that the method improves the overall class geometry instead of favoring one or two easy emitters.

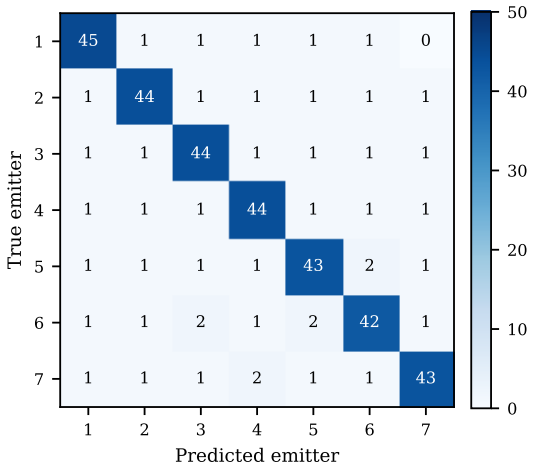


Fig. 6 Confusion matrix of RSGA-SEI under the 5-shot setting.

4.8 Complexity and Generation-Quality Indicators

Fig. 7 reports the computational profile. Because SAGA and SE-RWC calibration are used during training, the training-time model is heavier than the CSI baseline. However, the inference path remains close to ResNet-18 because generation and reliability scoring are not needed online. Fig. 8 further lists generated-sample quality indicators. The comparison indicates that SAGA improves diversity compared with a conventional generative adversarial network (GAN) while keeping feature distance lower than uncalibrated generation. The complexity comparison is included because online recognition is often executed at edge nodes or base stations with limited latency budgets. Although RSGA-SEI uses more training parameters than CSI-DA, its inference parameters remain 11.18M, and the latency only increases from 3.6 ms to 3.7 ms. Thus, most of the additional cost is paid offline during training rather than during spectrum operation. The structural similarity (SSIM), feature-distance, and diversity indicators are consistent with this design. SAGA obtains the highest SSIM and diversity while reducing the feature distance, indicating that the generated CSI samples are both closer to the measured distribution and less repetitive.

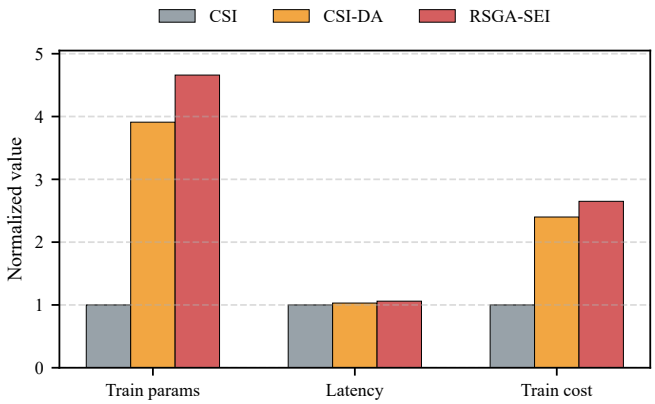


Fig. 7 Complexity comparison using normalized training and online indicators.

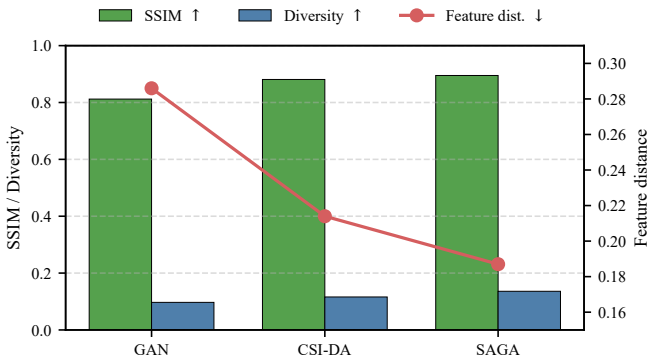


Fig. 8 Generated-sample quality indicators for fidelity, feature distance, and diversity.

5 Conclusion

This paper presents RSGA-SEI, a reliability-guided spectrum generative augmentation framework for few-shot SEI in next-generation mobile networks. The method connects CSI-based spectrum observation, spectrum-aware sample generation, support-prototype reliability calibration, and lightweight emitter classification. By replacing heavy geometric prototype optimization with SE-RWC, RSGA-SEI better matches embodied spectrum intelligence scenarios where training data are scarce and online inference must remain compact. The framework treats emitter recognition as part of a closed spectrum interaction loop, where recognition outputs can support spectrum access, resource coordination, and performance monitoring in adaptive mobile systems. During deployment, generative expansion and reliability calibration are removed, and only the lightweight feature extractor and classifier are retained. The experimental results show that reliability-guided calibration improves the 5-shot setting while keeping the deployment cost close to the CSI baseline. The method still

has limitations, since reliability weighting can suppress harmful synthetic samples but cannot fully recover missing spectrum patterns when the support set is extremely small or the generator misses important physical modes. Future work should therefore emphasize repeated-split testing, low-SNR analysis, and cross-shot evaluation.

References

[1] Lin, Y., Tu, Y., Dou, Z.: An improved neural network pruning technology for automatic modulation classification in edge devices. IEEE Transactions on Vehicular Technology **69**(5), 5703–5706 (2020)

[2] Lin, Y., Wang, M., Zhou, X., Ding, G., Mao, S.: Dynamic spectrum interaction of uav flight formation communication with priority: A deep reinforcement learning approach. IEEE Transactions on Cognitive Communications and Networking **6**(3), 892–903 (2020)

[3] Ding, L., Wang, S., Wang, F., Zhang, W.: Specific emitter identification via convolutional neural networks. IEEE communications letters **22**(12), 2591–2594 (2018) https://doi.org/10.1109/LCOMM.2018.2871465

[4] Xu, Z., Han, G., Zhu, H., Liu, L.: A lightweight specific emitter identification model for iiot devices based on adaptive broad learning. IEEE Transactions on Industrial Informatics **19**(5), 7066–7075 (2023) https://doi.org/10.1109/TII.2022.3220840

[5] Peng, Y., Liu, P., Wang, Y., Gui, G., Adebisi, B., Gacanin, H.: Radio frequency fingerprint identification based on slice integration cooperation and heat constellation trace figure. IEEE Wireless Communications Letters **11**(3), 543–547 (2022) https://doi.org/10.1109/LWC.2021.3135137

[6] Shen, G., Zhang, J., Marshall, A., Cavallaro, J.R.: Towards scalable and channel-robust radio frequency fingerprint identification for lora. IEEE Transactions on Information Forensics and Security **17**, 774–787 (2022) https://doi.org/10.1109/TIFS.2022.3152404

[7] Lin, Y., Zhao, H., Ma, X., Tu, Y., Wang, M.: Adversarial attacks in modulation recognition with convolutional neural networks. IEEE Transactions on Reliability **70**(1), 389–401 (2020)

[8] Tu, Y., Lin, Y., Hou, C., Mao, S.: Complex-valued networks for automatic modulation classification. IEEE Transactions on Vehicular Technology **69**(9), 10085–10089 (2020)

[9] Tu, Y., Lin, Y., Zha, H., *et al.*: Large-scale real-world radio signal recognition with deep learning. Chinese Journal of Aeronautics **35**(9), 35–48 (2022)

[10] Pan, Y., Yang, S., Peng, H., Li, T., Wang, W.: Specific emitter identification

based on deep residual networks. IEEE Access **7**, 54425–54434 (2019) https://doi.org/10.1109/ACCESS.2019.2913759

[11] Li, R., Hu, J., Li, S., Ai, W.: Specific emitter identification based on multi-domain features learning. In: 2021 IEEE International Conference on Artificial Intelligence and Industrial Design (AIID), pp. 178–183 (2021). IEEE

[12] Ying, S., Huang, S., Chang, S., He, J., Feng, Z.: Amscn: A novel dual-task model for automatic modulation classification and specific emitter identification. Sensors **23**(5), 2476 (2023)

[13] Peng, S., Jiang, H., Wang, H., Alwageed, H., Zhou, Y., Sebdani, M.M., Yao, Y.-D.: Modulation classification based on signal constellation diagrams and deep learning. IEEE transactions on neural networks and learning systems **30**(3), 718–727 (2018)

[14] Wang, C., Fu, X., Wang, Y., Gui, G., Gacanin, H., Sari, H., Adachi, F.: Interpolative metric learning for few-shot specific emitter identification. IEEE Transactions on Vehicular Technology **72**(12), 16851–16855 (2023)

[15] Tang, Z., Zhang, T., Du, Y., Su, J.: Individual identification method of little sample radiation source based on sgdcgan+ dcnn. IET Communications **17**(3), 253–264 (2023)

[16] Wang, W., Dou, Z., Fu, J., Lin, Y.: A diffusion model-based open set identification method for specific emitters. In: 2023 IEEE/CIC International Conference on Communications in China (ICCC), pp. 1–5 (2023). IEEE

[17] Liu, C., Fu, X., Wang, Y., Guo, L., Liu, Y., Lin, Y., Zhao, H., Gui, G.: Overcoming data limitations: A few-shot specific emitter identification method using self-supervised learning and adversarial augmentation. IEEE Transactions on Information Forensics and Security **19**, 500–513 (2023)

[18] Zhao, C., Hu, W., Cai, M., Wu, Z.: Data augmentation aided automatic modulation recognition using diffusion model. IET Communications **19**(1), 70111 (2025)

[19] Han, G., Wang, W., Xu, Z., Li, A.: Data augmentation and individual identification method for emitters using contour stella image mapping. IEEE Transactions on Cognitive Communications and Networking **10**(4), 1253–1262 (2024) https://doi.org/10.1109/TCCN.2024.3378214

[20] Shieh, C.-S., Lin, C.-T.: A vector neural network for emitter identification. IEEE Transactions on Antennas and Propagation **50**(8), 1120–1127 (2002)

[21] Yang, Z., Wu, Z., Yin, Z., Quan, T., Sun, H.: Hybrid radar emitter recognition based on rough k-means classifier and relevance vector machine. Sensors **13**(1),

848–864 (2013)

[22] Merchant, K., Revay, S., Stantchev, G., Nousain, B.: Deep learning for rf device fingerprinting in cognitive communication networks. IEEE journal of selected topics in signal processing **12**(1), 160–167 (2018)

[23] Peng, L., Zhang, J., Liu, M., Hu, A.: Deep learning based rf fingerprint identification using differential constellation trace figure. IEEE Transactions on Vehicular Technology **69**(1), 1091–1095 (2019)

[24] Lin, Y., Tu, Y., Dou, Z., Chen, L., Mao, S.: Contour stella image and deep learning for signal recognition in the physical layer. IEEE Transactions on Cognitive Communications and Networking **7**(1), 34–46 (2021)

[25] Zhao, D., Yang, J., Liu, H., Huang, K.: A complex-valued self-supervised learning-based method for specific emitter identification. Entropy **24**(7), 851 (2022)

[26] Wang, Y., Gui, G., Lin, Y., Wu, H., Yuen, C., Adachi, F.: Few-shot specific emitter identification via deep metric ensemble learning. IEEE Internet of Things Journal **9**(24), 24980–24994 (2022) https://doi.org/10.1109/JIOT.2022.3194967

[27] Yang, N., Zhang, B., Ding, G., Wei, Y., Wei, G., Wang, J., Guo, D.: Specific emitter identification with limited samples: A model-agnostic meta-learning approach. IEEE Communications Letters **26**(2), 345–349 (2021)

[28] Ling, Q., Yan, W., Zhang, Y., Yu, K., Wang, C.: Transfer learning method for specific emitter identification based on pseudo-labelling and meta-learning. IET Radar, Sonar & Navigation **18**(9), 1460–1473 (2024)

[29] Li, T., Wen, Z., Jin, C., Su, J., Li, J., Hong, Z., Zhang, X., He, S.: Meta-rff: Meta-task adaptive based few-shot open-set incremental learning for rf fingerprint recognition. IEEE Transactions on Cognitive Communications and Networking (2025)

[30] Davaslioglu, K., Boztaş, S., Ertem, M.C., Sagduyu, Y.E., Ayanoglu, E.: Self-supervised rf signal representation learning for nextg signal classification with deep learning. IEEE Wireless Communications Letters **12**(1), 65–69 (2022)

[31] Liu, D., Wang, P., Wang, T., Abdelzaher, T.: Self-contrastive learning based semi-supervised radio modulation classification. In: Milcom 2021-2021 Ieee Military Communications Conference (milcom), pp. 777–782 (2021). IEEE

[32] Wu, Z., Wang, F., He, B.: Specific emitter identification via contrastive learning. IEEE Communications Letters **27**(4), 1160–1164 (2023)

[33] Zhang, J., Liu, Q., An, S., He, Y.: Cross-domain few-shot specific emitter identification via contrastive self-supervised learning. IEEE Communications Letters

(2025)

[34] Liu, C., Fu, X., Ge, Y., Wang, Y., Lin, Y., Gui, G., Sari, H.: A robust few-shot sei method using class-reconstruction and adversarial training. In: 2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall), pp. 1–5 (2022). https://doi.org/10.1109/VTC2022-Fall57202.2022.10012716

[35] Zhang, X., Wang, Y., Zhang, Y., Lin, Y., Gui, G., Ohtsuki, T., Sari, H.: Data augmentation aided few-shot learning for specific emitter identification. In: 2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall), pp. 1–5 (2022). https://doi.org/10.1109/VTC2022-Fall57202.2022.10012809