



ARTICLE

Lightweight Meta-Learned RF Fingerprinting under Channel Imperfections for 6G Physical Layer Security

Chia-Hui Liu* and Hao-Feng Liu

Department of Electronic Engineering, National Formosa University, Yunlin, Taiwan

*Corresponding Author: Chia-Hui Liu. Email: liuch@nfu.edu.tw

Received: 17 December 2025; Accepted: 11 February 2026; Published: 30 March 2026

ABSTRACT: Artificial Intelligence (AI)-native sixth-generation (6G) wireless networks require data-efficient and channel-resilient physical-layer modeling techniques that learn stable device-specific representations under channel variations and hardware imperfections to support secure and reliable device-level authentication under highly dynamic environments. In such networks, massive device heterogeneity and time-varying channel conditions pose significant challenges, as reliable authentication must be achieved with limited labeled data and constrained edge resources. To address this challenge, this paper proposes an Artificial Intelligence (AI)-assisted few-shot physical-layer modeling framework for channel robust device identification, formulated within the paradigm of Specific Emitter Identification (SEI) based on radio frequency (RF) fingerprinting. The proposed framework explicitly formulates few-shot SEI as a channel-resilient physical-layer modeling problem by integrating a lightweight convolutional neural network and Transformer hybrid encoder with a dual-branch feature decoupling mechanism. Device specific RF fingerprints are separated from channel-dependent factors through orthogonality-constrained learning, which effectively suppresses channel-induced prototype drift and stabilizes metric geometry under channel variations. A meta-learned prototypical inference module is further employed under episodic few-shot training, enabling rapid adaptation to new devices and unseen channel conditions using only a small number of labeled samples. Experimental results on multiple real-world RF datasets, including ORACLE Wi-Fi transmitter measurements and civil aviation ADS-B broadcasts ($D_{\text{Wi-Fi}}$, $D_{\text{ADS-B}}$, and $D_{\text{ADS-B}}^{\text{DF17}}$), demonstrate that the proposed method achieves identification accuracy ranging from 99.1% to 99.8% using only 10 labeled samples per device, while maintaining episode-level performance variance below 0.02. In addition, the proposed model contains approximately 1.45×10^5 trainable parameters, making it suitable for deployment on resource-constrained edge devices. These results indicate that the proposed framework provides a concrete and scalable AI-driven solution for physical-layer security and device-level authentication in AI-native 6G wireless networks.

KEYWORDS: 6G wireless networks; specific emitter identification; RF fingerprinting; few-shot learning

1 Introduction

With the rapid expansion of the Internet of Things (IoT) and industrial wireless communication systems, future network architectures must maintain stable and reliable terminal identification capabilities over long-term operations. Sole reliance on software-level credentials, such as cryptographic keys or serial numbers, is vulnerable to tampering and leakage. Moreover, frequent device replacement, high labeling costs, and constrained edge computing resources highlight the need for lightweight, self-adaptive, and scalable device-level authentication mechanisms. Against this backdrop, Specific Emitter Identification (SEI), also commonly referred to as radio frequency (RF) fingerprinting, exploits device-specific hardware

imperfections introduced during RF manufacturing to identify emitters directly from physical-layer signals, providing a security mechanism independent of software-level information. This capability is particularly critical in AI-native sixth generation (6G) communication systems, where intelligent and self-adaptive signal management and authentication are expected to be achieved through learning-driven models. Early SEI methods based on traditional signal processing predominantly relied on handcrafted features and pipeline-based workflows. Typical approaches extracted RF impairments such as I/Q imbalance, carrier frequency offset (CFO), and phase noise from steady-state or transient signal segments, followed by classification using higher-order statistics and time frequency transformations, including Short-Time Fourier Transform (STFT), Continuous Wavelet Transform (CWT), Fractional Fourier Transform (FrFT), and bispectrum analysis. Although these methods offer strong interpretability, their performance degrades under low signal-to-noise ratio (SNR) and highly dynamic channel conditions, where noise, multipath effects, interference, and frequency drift tend to obscure device-specific characteristics. With the advancement of deep learning, end-to-end feature extraction has been increasingly adopted in SEI. Convolutional Neural Networks (CNNs) and hybrid CNN Transformer architectures have demonstrated improved identification performance by learning hierarchical and global feature representations directly from raw or transformed signals. However, these models typically depend on large-scale labeled datasets and assume distributional consistency, limiting their generalization to new devices or open-set environments. In practical deployments, data scarcity and heterogeneity are inevitable due to privacy constraints and high acquisition costs. To address these challenges, Few-Shot Learning (FSL) combined with Meta-Learning has emerged as a promising solution. Metric-based meta-learning approaches, such as prototypical networks, learn task-invariant embedding spaces by constructing stable geometric structures across episodic training tasks, enabling rapid adaptation to new devices and unseen channel conditions while maintaining robust decision boundaries.

The system models the wireless emission process through an AI-based feature decoupling and prototypical metric framework. Raw I/Q signals collected through transmitter channel receiver paths are normalized and fed into a lite CNN Transformer encoder that learns device-specific and channel-invariant representations. The model serves as a computational surrogate for hardware-level imperfections, enabling device-level authentication and physical-layer security in AI-native 6G and IoT networks. In summary, this study addresses the challenges of entangled device fingerprints and channel effects in SEI tasks, as well as the difficulty of achieving stable identification under few-shot conditions. To this end, we propose the FS-SEI framework, which integrates feature disentanglement with prototypical meta-learning, as illustrated in Fig. 1. In the proposed framework, device-related and channel-related features are explicitly disentangled during the feature extraction stage, with orthogonality constraints imposed to preserve their independence. At the decision stage, a prototypical space is constructed solely from device features to perform metric-based classification, thereby mitigating the impact of channel perturbations on the decision boundaries. Through an episodic meta-learning procedure and a channel parameter sweeping training strategy covering simulated disturbances such as Carrier Frequency Offset (CFO), Additive White Gaussian Noise (AWGN), and I/Q imbalance the model autonomously calibrates the embedding space across varying channel conditions. This enables rapid alignment with new class distributions under limited samples, achieving high stability and reproducibility. Overall, the proposed design not only provides an AI-driven approach for physical-layer modeling and authentication but also aligns with the requirements of AI-native 6G architectures, which emphasize self-adaptive physical-layer modeling and device-level identity verification. This framework offers a stable and scalable foundation for secure modeling in future intelligent communication systems and edge networks. As 6G communication architectures evolve toward AI-native and highly decentralized paradigms, future network nodes will no longer rely on traditional centralized core networks. Instead, they

will be widely deployed across multi-tier infrastructures such as cell-free massive Multiple-Input Multiple-Output (MIMO), Reconfigurable Intelligent Surface (RIS) assisted communications, and integrated satellite terrestrial networks. In such low-latency, highly dynamic wireless environments, base stations, edge nodes, and unmanned aerial vehicle (UAV) relays must rapidly identify whether surrounding devices are trusted nodes, even under unpaired or extremely limited-sample conditions. Therefore, Few-Shot Specific Emitter Identification (FS-SEI), which exhibits few-shot capability and stable feature representation under unknown channel conditions and hardware perturbations, is poised to become a foundational mechanism for device authentication and physical-layer security in 6G systems directly addressing the core requirements of the model proposed in this study. Unlike existing FS-SEI methods that primarily rely on either meta-learning or feature disentanglement in isolation, the proposed framework explicitly couples orthogonality-guided feature decoupling with episodic meta-learning, enabling the learned metric geometry to remain stable under channel variations while preserving rapid few-shot adaptability.

The main contributions of this work are summarized as follows:

- We formulate few-shot Specific Emitter Identification (FS-SEI) as a channel-resilient physical-layer modeling problem, explicitly addressing the challenges of data scarcity, channel variation, and device heterogeneity in AI-native 6G wireless environments.
- We propose a lightweight meta-learned FS-SEI framework that integrates a CNN Transformer hybrid encoder with a dual-branch feature disentanglement mechanism, enabling explicit separation of device-specific RF fingerprints from channel-dependent factors.
- An orthogonality-guided representation stabilization strategy is introduced to suppress feature entanglement and mitigate channel-induced prototype drift, ensuring stable metric geometry for prototype-based inference under dynamic channel conditions.
- A prototypical metric learning module is employed under episodic meta-learning, allowing rapid adaptation to new devices and unseen channel conditions using only a small number of labeled samples, without relying on large-scale pretraining or complex hybrid pipelines.
- Extensive experiments on multiple real-world RF datasets (D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17}) demonstrate that the proposed method achieves high identification accuracy, low variance, and stable convergence under few-shot and cross-channel settings, while maintaining a lightweight model size suitable for edge deployment.

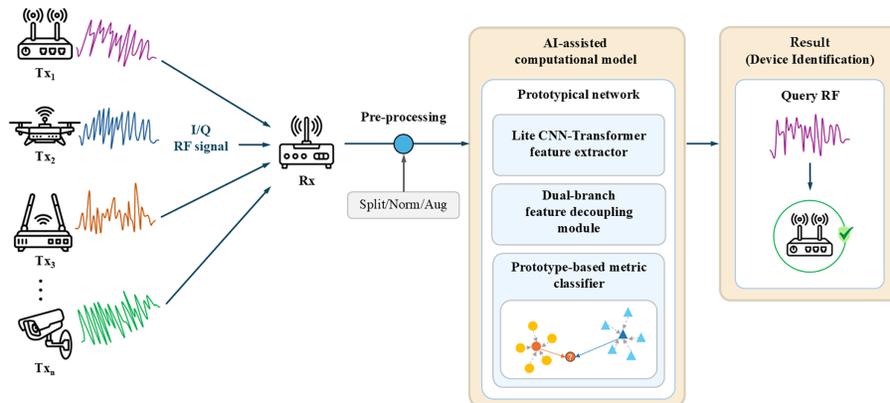


Figure 1: Overall framework of the proposed AI-assisted FS-SEI computational model.

2 Related Work

2.1 Large-Scale Deep Specific Emitter Identification

In the field of SEI, deep learning has become the dominant approach for building high-performance systems due to its strong feature extraction and nonlinear modeling capabilities. Under controlled conditions with large-scale labeled datasets, Convolutional Neural Networks (CNN) effectively suppress channel noise and signal distortion, while Complex-Valued Neural Networks operate directly on complex baseband signals to enhance RF feature representation. These models have established strong performance baselines, and early supervised studies formulated SEI as an end-to-end representation learning problem using standard CNN architectures [1,2]. To address deployment constraints on resource-limited edge devices, subsequent works combined one-dimensional CNNs (1D-CNNs) with Discrete Wavelet Transform (DWT) for dimensionality reduction and complexity mitigation [3], achieving superior accuracy over traditional machine learning methods even under low signal-to-noise ratio (SNR) conditions. Network compression and lightweight design techniques were further introduced to reduce the computational cost of Complex-Valued Neural Networks (CVNNs), enabling near-perfect accuracy and significant latency reduction in high-SNR scenarios [4]. Recent research has focused on improving robustness to domain shifts caused by channel variation and receiver diversity. Representative approaches include multi-receiver collaboration and Long Short-Term Memory (LSTM)-based temporal transfer modeling [5], which enhance cross-environment discriminability through multi-perspective signal decomposition. A domain-adversarial Transformer encoder has been proposed to mitigate performance degradation caused by limited annotations and distribution mismatch [6]. In addition, adaptive strategies combining contrastive learning with triplet loss, entropy minimization, self-supervised pseudo-labeling, and contrastive regularization have been introduced to adapt to receiver characteristics under unlabeled settings [7]. Bispectrum-based feature fusion with CNN decision modules has also been explored to improve stability in distributed sensing scenarios [8], often combined with transfer learning to reduce cold-start effects and convergence time. Despite these advances, deep learning-based SEI fundamentally depends on large-scale, high-quality labeled data. In real-world and dynamic environments, device identification and ground-truth annotation are difficult to obtain, and data collection is further constrained by privacy and regulatory issues. Moreover, distribution shifts caused by channel drift and receiver variability limit the effectiveness of conventional supervised strategies in maintaining stable performance and enabling rapid deployment. This mismatch between data-intensive deep learning models and sparse-data operational scenarios significantly restricts their applicability to few-shot and rapid adaptation tasks. These challenges highlight the growing need for data-efficient and adaptive physical-layer modeling frameworks, particularly in the context of future 6G and large-scale IoT networks. Despite their strong performance under abundant labeled data and controlled conditions, large-scale deep SEI models are inherently ill-suited for few-shot 6G scenarios, as their reliance on data-intensive training and heavyweight architectures limits adaptability under data scarcity and restricts practical deployment on resource-constrained edge devices.

2.2 Few-Shot Learning SEI

Against this background, few-shot learning has emerged as a key viable solution to enhance the adaptability and efficiency of SEI systems in dynamic environments. In particular, for 6G and edge wireless application scenarios, FSL methods that can maintain identification stability under extremely low annotation and highly dynamic channel conditions are of critical importance for enabling AI-assisted device-layer modeling and communication security. The fundamental advantage of this approach lies in its ability to rapidly establish and maintain stable class decision boundaries using only a minimal number of support samples.

1. **Transfer-Based Fine-Tuning:** Among existing Few-Shot Supervised Environmental Inference (FS-SEI) methods, transfer-based FSL is widely adopted. These methods typically follow a two-stage paradigm consisting of self-supervised pretraining and subsequent (semi-)supervised fine-tuning. The pretraining stage aims to learn transferable and domain-invariant representations, which are then adapted to downstream FS-SEI tasks using limited labeled data. Recent studies primarily employ self-supervised learning (SSL) for pretraining. Contrastive self-supervised learning (CSSL) frameworks have been shown to improve robustness to noise and reduce sensitivity to weight initialization [9]. Masked modeling approaches, such as the Asymmetric Masked Autoencoder (AMAE) [10], further enhance annotation efficiency and feature transferability in RF applications. In addition, combining self-supervised pretraining with adversarial augmentation has been proposed to learn more flexible and cross-source discriminative radio frequency fingerprint (RFF) representations [11]. The main advantage of transfer-based methods lies in their high training parallelism and compatibility with existing pretrained models. However, their performance is limited by the domain gap between pretraining data and the target FS-SEI task.
2. **Metric-Based Learning:** In contrast to transfer-based approaches, metric-based FSL learns a discriminative embedding space and performs classification using distance or similarity measures. Instead of fixed classifiers, it relies on prototype-based decision structures to directly model semantic relationships, making it well suited for few-shot scenarios with limited labeled data. The performance of metric-based methods depends on the quality and stability of the embedding space. When class numbers increase or intra-class variation is large, distance reliability and classification accuracy may degrade. To improve robustness, several extensions have been proposed. Metric Adversarial Training (MAT) [12] combines adversarial regularization on unlabeled data with metric constraints on labeled samples, enhancing stability under low annotation ratios. Attention mechanisms such as the Convolutional Block Attention Module (CBAM) adaptively weight time–frequency features [13], while robust representation learning for SEI under class-irrelevant factors has also been explored to improve generalization and robustness in complex environments [14]. These approaches improve robustness to domain shifts and channel perturbations. Overall, metric-based learning provides a lightweight, interpretable, and effective framework for FS-SEI, and is well suited for device identification and physical-layer security in future 6G wireless systems.
3. **Hybrid Models and Frameworks:** As few-shot recognition tasks become more complex, recent studies have adopted hybrid frameworks that integrate multiple learning paradigms, including metric learning, self-supervised learning, transfer learning, and traditional signal processing, to improve FS-SEI robustness and generalization. These frameworks typically employ multi-branch architectures to balance fine-grained feature extraction with global representation stability. For example, the Asymmetric Dual-Path Masked Autoencoder (ADPMAE) framework [15] enhances discriminative capability under limited samples by jointly reconstructing local signal details and modeling global semantic structures. The Knowledge–Data–Model (KDM) driven framework [16] combines expert-crafted features, self-supervised learning, and supervised FS-SEI to achieve strong robustness and high processing efficiency. In addition, a hybrid FS-SEI framework based on feature contrastive reconstruction and virtual adversarial training (VAT) [17] demonstrates effective performance in both self-supervised and semi-supervised scenarios under complex electromagnetic environments. This model adopts a CNN–Transformer architecture to capture temporal and spectral characteristics, enabling a balanced representation of global semantics and local structures. Despite their strong performance, hybrid frameworks often suffer from increased structural complexity, high computational cost, and extended training time, which significantly limit their practicality for deployment. As future 6G and intelligent edge networks impose strict constraints on computational resources, lightweight and scalable AI modeling

strategies are increasingly required. Consequently, designing hybrid FS-SEI frameworks that balance performance and computational efficiency remains a critical research challenge.

As summarized in Table 1, large-scale deep SEI methods, while effective under abundant labeled data, are fundamentally mismatched with few-shot 6G scenarios due to their data-intensive training requirements and heavyweight model architectures, which limit adaptability under data scarcity and hinder deployment on resource-constrained edge devices.

Table 1: Comparison between large-scale deep SEI and representative FS-SEI paradigms.

Method Category	Learning Paradigm	Data Requirement	Model Complexity	Robustness to Channel Variation	Suitability for 6G
Large-Scale Deep SEI	Supervised learning	Large-scale labeled datasets	High	Limited	Not suitable
Transfer-Based FS-SEI	Pretraining and fine-tuning	Moderate to large	High	Moderate	Partially suitable
Proposed FS-SEI	Orthogonality-guided decoupling and meta-learning	Low	Low	Strong	Highly suitable

2.3 Feature Decoupling and Orthogonal De-Correlation

In SEI, device-invariant factors namely stable radio frequency fingerprints (RFFs) caused by hardware imperfections are often entangled with channel-variant factors within the same embedding space. This entanglement leads to prototype drift and unstable decision boundaries. Feature decoupling aims to explicitly separate these two types of factors within the embedding layer, making it particularly suitable for FSL scenarios. By constructing class prototypes primarily based on device-related components, feature decoupling can significantly enhance intra-class compactness and inter-class separability, thereby improving classification stability. Relevant studies have used feature disentanglement to suppress class-irrelevant factors and improve discriminability [18], ensuring that prototype distance calculations are predominantly influenced by key signal peaks, which represent device-specific characteristics, thus reinforcing feature discriminability. Other approaches employ auxiliary tasks such as phase shift prediction [19] during the pretraining stage to learn both phase-relevant and irrelevant signal components. Orthogonality constraints are then applied in downstream FS-SEI tasks to precisely disentangle device-related features from channel-induced variations. Such methods have been shown to significantly outperform various deep learning baseline models on D_{ADS-B} few-shot datasets. They effectively improve the geometric distribution of the prototype space by enhancing intra-class compactness and inter-class separation. Feature decoupling and orthogonal de-correlation provide a clear mechanism to counter domain shift, maintaining representational stability and separability under cross-channel, cross-device, and distributional variation scenarios. These techniques offer a theoretically grounded approach to representation disentanglement for few-shot identification. Their ability to preserve stable representations under channel variation. In summary, SEI techniques have progressively evolved from traditional hand-crafted feature approaches toward data-driven architectures centered on deep learning. This transition aligns with the core demands of AI-native communication system

modeling envisioned in 6G architectures and provides theoretical justification and practical foundation for the FS-SEI framework proposed in this study.

3 Methodology and Framework Design

Under the vision of 6G communication architectures, achieving robust physical-layer modeling and device-level identification under limited supervision is one of the key challenges in intelligent communications. In this context, we propose a Few-Shot Specific Emitter Identification framework, which formulates SEI as a metric-based meta-learning task. The objective is to learn self-adaptive device representations that remain stable under highly dynamic channel conditions.

3.1 Problem Definition: Few-Shot SEI under Multi-Domain Channel Variations

To address the challenges of data scarcity and high variability in specific emitter (SE) signals, this study formulates SEI as a few-shot metric learning problem. Furthermore, SEI is treated as an AI-driven physical-layer signal modeling task aimed at learning robust and generalizable representations under multi-domain channel conditions. Let the set of source entities be denoted as $c = \{1, \dots, N\}$, where only k labeled samples are available per class. Each episodic training instance follows the N-way k-shot, q-query scheme, in which N classes (i.e., devices) are randomly sampled from the training set. For each class, k support samples from the support set S , and q samples from the query set Q . The input consists of I/Q signal segments of fixed length, which are encoded by a feature extractor $Q f_{\theta}$ into two complementary subspace representations: a shared feature z_{shared} and a specific feature z_{spec} . We denote z_{shared} as the channel-related representation capturing environment-consistent factors and z_{spec} as the device-specific representation encoding hardware-induced RF fingerprints. The objective of the task is to achieve accurate and reproducible device-level identification using radio frequency fingerprints (RFFs) under unknown channel conditions and limited labeled data. This problem setting not only simulates the challenges commonly encountered in traditional SEI scenarios but also aligns with the requirements of AI-native communication modeling in 6G and large-scale IoT environments. Specifically, it targets robust physical-layer identification and modeling under few-shot and highly dynamic channel conditions.

3.2 Meta-Learned SEI Framework for AI-Native 6G Communications

This study proposes a Few-Shot FS-SEI framework based on the concept of metric-based meta-learning, as illustrated in Fig. 2. The framework can be viewed as an AI-driven physical-layer modeling pipeline designed to simulate and learn stable representations of device-level RF behavior under channel perturbations. Given input I/Q signal segments, the model first employs a dual-branch feature extractor to explicitly decouple the representations into two components: a device-specific feature z_{spec} , which captures fine-grained hardware-induced variations of the source entity, and a shared feature z_{shared} , which encodes channel-consistent environmental information. An orthogonality constraint is applied between z_{spec} and z_{shared} to enforce approximate linear independence, ensuring that the metric geometry remains unaffected by channel-related shared factors and that feature separability is preserved. In the metric-based inference head, only z_{spec} is used to construct class prototypes, and query inference is performed based on distance metrics. This design focuses the decision boundary on device-level characteristics. The overall training strategy integrates episodic FSL with supervised learning. During meta-iterations, multiple task-specific losses are aggregated to update the model parameters, enabling the framework to generalize across varying domains and device identities. The entire training mechanism is built upon the principle of meta-learning, enabling the model to accumulate cross-device and cross-channel experience across multiple tasks and learn a generalized initial parameter set θ . When encountering a new device or channel condition, the

model can rapidly adapt with only a few samples. This design equips the FS-SEI framework with the capability for self-adaptive communication modeling, aligning with the concepts of online re-training and continual learning emphasized in 6G architectures. Overall, FS-SEI can be regarded as a meta-learned computational model capable of automatically learning stable mappings of device fingerprints under varying channel perturbations. It achieves a tight integration of AI-driven physical-layer modeling and device-level identification, offering a scalable and reproducible modeling foundation for physical-layer security and edge device authentication in future 6G communication systems.

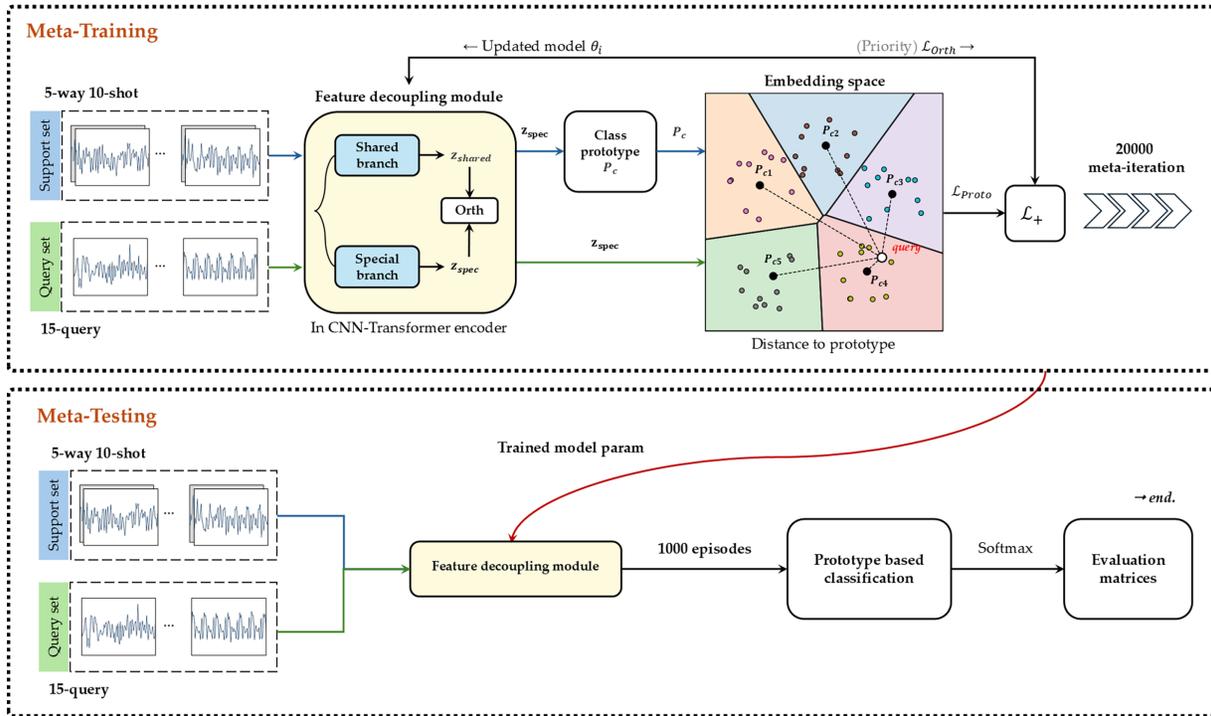


Figure 2: Few-shot SEI workflow.

3.2.1 Feature Disentanglement for Robust Physical-Layer Representation Learning

Prior to prototype generation, the feature extractor module is responsible for extracting and disentangling device-intrinsic characteristics from channel-variant factors embedded in raw radio frequency signals. The primary design objective is to suppress channel perturbations such as path loss, obstruction, multipath effects, and hardware gain discrepancies. By mitigating these disturbances, the module outputs stable and quantifiable device representations to the decision-making stage. The overall framework comprises three interconnected components: (1) a lightweight hybrid backbone, (2) a dual-branch feature decoupling module, and (3) a sequence-level integration and decorrelation unit.

- (1) **Lightweight Hybrid Backbone for I/Q Representation:** I/Q segments from both the support and query sets are first processed by the backbone to obtain intermediate sequential representations. The backbone adopts a hybrid architecture consisting of a front-end convolutional neural network (CNN) and a back-end lightweight Transformer. This design captures both local temporal or time frequency textures and long-range dependencies while effectively suppressing redundant information. The output is a scale-consistent feature tensor that facilitates the subsequent separation of shared and device-specific features. It also serves as the common input to both the decoupling and metric learning modules.

- (2) **Dual-Branch Feature Decoupling for Shared and Specific Components:** This module is embedded within the CNN layers of the backbone to generate the shared feature z_{shared} and the device-specific feature z_{spec} . The shared branch is designed with a moderate receptive field and controlled downsampling rate, focusing on learning rhythmic patterns and signal envelopes that remain stable across devices and environments. By progressively expanding the effective receptive field, the model captures long-range, channel-invariant structures without significantly increasing computational overhead, while remaining insensitive to transient spikes and stochastic noise. In contrast, the specific branch adopts fine-grained filtering and cross-channel reorganization strategies to enhance responsiveness to device-level attributes such as process-induced distortions, modulation residues, I/Q imbalance, and scenario-specific interferences. Redundancy and latency are reduced through parameter decomposition and per-channel operations. Both branches operate with aligned temporal strides and downsampling rates to ensure statistical consistency in scale, thereby facilitating subsequent metric learning and decorrelation procedures. The outputs of the two branches are concatenated along the channel dimension, resulting in a unified feature tensor $H = [H_{shared} \parallel H_{spec}] \in \mathbb{R}^{B \times 192 \times T/6}$.
- (3) **Orthogonality-Guided Representation Stabilization:** A lightweight Transformer is employed to enhance long-range temporal dependencies while suppressing redundancy along the channel dimension. Through an attention mechanism, dynamic weighting is applied to both shared and specific features, enabling global context aggregation at low computational cost. The integrated representation is subsequently projected back into two complementary subspaces to obtain the final outputs z_{shared} and z_{spec} .

After feature disentanglement, an orthogonality constraint is imposed *during the training phase* to explicitly suppress residual statistical dependence between the two subspaces. This constraint is applied exclusively to the batch-normalized representations z_{shared} and z_{spec} . Produced by the dual-branch module, rather than to the raw input signals or class prototypes. Its purpose is to prevent channel-related shared factors from leaking into the device-specific embedding space, thereby stabilizing the geometric structure used for subsequent metric inference. During the training phase, an orthogonality loss is imposed on the batch-normalized representations to enforce near-orthogonality between the two subspaces in terms of their inner product. This constraint reduces mutual interference and promotes independent feature encoding.

$$\mathcal{L}_{Loth} = \frac{1}{N \times k} \sum_{i=1}^{N \times k} \|z_{shared}^T z_{spec}\|_F^2 \quad (1)$$

The orthogonality constraint is applied uniformly across all experiments and datasets and is jointly optimized with the classification objective, without dataset-dependent weighting or tuning. This constraint maintains linear independence and suppresses feature entanglement. By enforcing orthogonality at the representation level, channel-related variations are prevented from perturbing the device-specific subspace, ensuring that subsequent prototype construction relies on stable and discriminative device-level attributes. In the subsequent decision-making stage, only z_{spec} is utilized for prototype generation, thereby anchoring class discrimination on device-level attributes. This strategy enhances intra-class compactness and inter-class separability while mitigating the risk of prototype drift. Overall, this module offers strong robustness to cross-channel variations under controllable parameter and computational costs. It enables stable and reproducible few-shot identification performance in real-world wireless environments, and provides a practical structural design for feature disentanglement and physical-layer representation learning in 6G-oriented physical layer modeling.

The lightweight nature of the proposed framework is not achieved merely by reducing model size, but is jointly constrained by architectural and training design choices. On the feature extraction side, a

lite CNN Transformer hybrid encoder is adopted, explicitly avoiding global self-attention with quadratic complexity in sequence length, resulting in near-linear scaling during training and inference. On the decision side, prototype-based metric inference is employed, eliminating large learnable classifier weights and concentrating parameters on the shared backbone and lightweight disentanglement modules. Compared to hybrid FS-SEI frameworks that rely on self-supervised pretraining, multi-branch reconstruction, contrastive learning, or VAT, the proposed method requires only a single episodic supervised meta-learning procedure, substantially reducing training complexity and deployment overhead. Quantitatively, the proposed model contains approximately 1.45×10^5 trainable parameters, which is significantly lower than typical CNN/Transformer-based architectures. Under the most demanding $D_{\text{ADS-B}}^{\text{DF17}}$ configuration, inference latency is measured on an NVIDIA RTX 3060 GPU with batch size is 1. The average forward inference latency per query sample is approximately 0.084 ms, while the total inference time per episode is approximately 6.323 ms, indicating that the proposed framework is suitable for real-time or near-real-time deployment on resource-constrained edge devices. We also acknowledge a scalability limitation under extremely large N-way settings. Prototype-based inference requires computing distances between each query and all class prototypes, leading to inference cost that grows approximately linearly with N. When N becomes very large, additional strategies such as coarse-to-fine hierarchical prototypes, candidate set reduction, approximate nearest neighbor search, or prototype compression may be required, which are left for future work.

3.2.2 Prototype-Guided Metric Module for Few-Shot Device-Level Identification

This study adopts a metric-based inference head as the final prediction module. The core principle aligns with that of the Prototypical Network (ProtoNet), relying on class prototypes and metric distance calculations. Rather than depending on a large number of learnable classification weights, the model performs inference through a measurable geometric structure in the embedding space. This approach better conforms to the modeling philosophy of AI-driven physical-layer representation learning. In each N-way k-shot episodic task, the device-specific features z_{spec} of each class in the support set S are aggregated to form prototype vectors. These prototypes are then used to evaluate the similarity between query samples and each class via a metric distance function. As a result, the model performs classification without relying on large-scale learnable classification weights.

Formally, for the support subset S_c corresponding to class c , the prototype vector is defined as:

$$P_c = \frac{1}{|S_c|} \sum_{x \in S_c} z_{spec}(x), c = 1, \dots, N \quad (2)$$

This step aggregates the device-specific features of a limited number of support samples into class prototypes via mean pooling. Under conditions of scarce data and varying channel environments, this operation helps mitigate the impact of sample-specific noise on the decision geometry, providing a stable anchor for subsequent metric evaluation.

For any query sample x_q , its proximity to each prototype is measured using a temperature-scaled Euclidean distance as Eq. (3):

$$d_c(x_q) = \frac{\|z_{spec}(x) - P_c\|_2^2}{\tau} \quad (3)$$

where $d(\cdot)$ denotes the distance metric function, and $\tau > 0$ is the temperature scaling factor. The temperature parameter τ is fixed across all experiments rather than tuned per dataset. This design maintains a consistent metric scale for prototype-based inference. A smaller τ sharpens the decision boundary by increasing the

model's sensitivity to small distance variations, whereas a larger τ smooths the output scores, which helps reduce overconfidence and improve calibration.

The distance is transformed into a probability by applying the Softmax function to the negative distance, thereby obtaining the posterior probability P that a query sample belongs to class c as Eq. (4).

$$P(y = c|x_q) = \frac{\exp(-d_c(x_q))}{\sum_{k=1}^N \exp(-d_k(x_q))} \quad (4)$$

The prototype classification loss is defined as Eq. (5). During training, we compute the average cross-entropy loss over the query set Q as the prototype classification loss \mathcal{L}_{Proto} , which is jointly optimized with the orthogonal decorrelation regularization from the disentanglement stage to yield the final objective function:

$$\mathcal{L}_{Proto} = -\frac{1}{|Q|} \sum_{(x_q, y_q) \in Q} \log p(y_q|x_q) \quad (5)$$

where $p(y_q|x_q)$ denotes the predicted probability of query sample x_q belonging to class y_q .

The overall training objective is given by:

$$\mathcal{L} = \mathcal{L}_{Proto} + \lambda \mathcal{L}_{Loth} \quad (6)$$

The parameter λ serves as a weighting coefficient to balance the contributions of classification performance and representation independence. The overall model minimizes the expected value of Eq. (6) through an episodic meta-learning strategy: Decoupling and Prototypical Metric Network for Few-Shot SEI, enabling the embedding space and metric function to better align with few-shot application scenarios. This design maintains a stable decision geometry under conditions of sample scarcity and channel variability, aligning with the metric-based, data-efficient modeling requirements of future sixth-generation (6G) communication systems. Furthermore, it provides an interpretable and extensible decision foundation for physical-layer security and device-level authentication in 6G networks.

3.2.3 Hybrid Meta-Learning Strategy for Cross-Channel Robustness

Under few-shot SEI episodic training, prototype estimation often exhibits high variance and unstable convergence due to task sampling variability and channel perturbations, resulting in gradient oscillation and decision boundary drift. To address these issues, this study proposes a supervised hybrid meta-learning framework to improve training stability and cross-channel generalization. The training follows a bi-level optimization architecture. At the inner level, N -way k -shot, q -query episodic tasks are used to preserve few-shot adaptation characteristics. At the outer level, losses from multiple episodic tasks are aggregated, and a single meta-update is performed to control convergence behavior. Within each episode, support and query samples are processed by a dual-branch encoder to extract shared and distinctive features. Metric-based classification is performed using only the distinctive features to construct class prototypes, ensuring that learning focuses on device-specific characteristics while remaining robust to channel variations. After aggregating losses across tasks, model parameters are updated once using a unified learning rate and a positive step size, which stabilizes prototype geometry and mitigates gradient direction drift. Training and validation both follow a fixed episodic protocol, with statistical performance monitoring used for stability assessment and early stopping. This approach enables stable few-shot training without requiring an additional pretraining stage, while maintaining strong adaptability and cross-channel robustness. Furthermore,

the proposed pipeline is reproducible, scalable, and consistent with self-adaptive learning principles in sixth-generation (6G) communication systems, providing a practical foundation for physical-layer modeling and device-level authentication tasks.

3.2.4 Geometric Modeling of Device-Specific Prototypes

The core objective of the proposed prototype space design is to learn a geometrically stable and semantically consistent device representation space, rather than merely increasing classifier capacity. To this end, prototypes are constructed exclusively from device-specific features, without incorporating channel-related shared components. This design prevents prototype positions from being distorted by varying channel conditions, allowing each prototype to focus on capturing intrinsic hardware characteristics of the device, thereby enhancing stability across different channels and tasks. Compared to standard Prototypical Networks that compute prototypes using holistic embeddings, the proposed method explicitly restricts prototype construction to a device-related subspace through feature disentanglement and orthogonality constraints, conceptually redefining the semantic meaning and geometric role of prototypes. Under this setting, distance measurements in the embedding space possess clear geometric interpretation; therefore, Euclidean distance is adopted as the metric function to directly reflect the actual displacement between samples and their corresponding prototypes. The orthogonality constraint further suppresses mutual interference between device-specific and channel-related features, encouraging tighter intra-class clustering and clearer inter-class separation of prototypes within the embedding space. As a result, prototype drift under few-shot conditions and channel variations is effectively reduced, leading to improved overall decision stability.

3.3 Channel Perturbation Modeling, Simulation Design, and Evaluation Protocol

To consistently characterize multi-domain channel variation in highly dynamic wireless environments during both training and evaluation, this study applies controlled channel perturbations to the input I/Q segments during episode construction. A parameter sweep simulation is employed to systematically cover common front-end imperfections and channel conditions. The configurations and ranges of all perturbations are aligned with those used in the experimental section (see [Table 1](#)), and the same perturbation scheme is applied across the train, validation, and test subsets to ensure reproducibility and fair comparison.

(A) Unified Channel Perturbation Model and Parameter Ranges:

- Carrier Frequency Offset (CFO): Introduces a linear phase rotation in the time domain or a frequency shift in the frequency domain to simulate local oscillator drift and Doppler effects. The perturbation range is set to $\{-30, 50\}$.
- Additive White Gaussian Noise (AWGN): Gaussian white noise is added to the modulated signal to simulate various signal-to-noise ratio (SNR) conditions and evaluate robustness under weak signal scenarios. The perturbation range is set to $\{10, 30\}$.
- I/Q Imbalance: Both amplitude and phase asymmetries are introduced to emulate analog front-end imperfections, thereby reducing the model's reliance on memorized fixed gain or phase patterns. This encourages the distinctive branch to focus on subtle distortions related to device manufacturing variations. The perturbation ranges are set to amplitude $\{-10, 10\}$ and phase $\{-5\%, 5\%$.

The aforementioned perturbations are consistently applied to both the support and query sets during episode generation. A grid-based parameter sweep is employed to cover multiple combinations, thereby forming an *in-silico* population of channel variations. This design fulfills the AI-native communication modeling requirements emphasized in 6G and large-scale IoT systems by enabling the learning of channel-resilient representations under multi-domain variations.

- (B) **Task-Oriented Configuration and Data Partitioning:** The dataset is partitioned into training, validation, and testing subsets with a ratio of 7:2:1. Each episode follows a fixed N-way, 10-shot, 15-query configuration, simulating a continuous frame-level monitoring scenario. For additional few-shot analysis, varying values of $k \in \{5, 10, 20, 30, 40\}$ are also evaluated. To prevent information leakage, no consecutive time segments are allowed to span across different data subsets. The distribution of source entities (classes) and channel condition ratios is maintained consistently across the training, validation, and test sets. All segments are normalized using z-score standardization. Channel perturbations are applied uniformly to all three subsets following the principles outlined in Section (A).
- (C) **Evaluation Protocol and Metrics:** During the testing phase, a fixed model with parameters θ is evaluated using a Monte Carlo-based protocol over 1000 episodes. Only forward inference is performed, without backpropagation, while maintaining the same sampling and prototypical inference procedure as in training. The primary evaluation metrics are accuracy and macro-averaged F1 score, reported with mean and standard deviation. For auxiliary analysis, confusion matrices and t-SNE visualizations are employed to assess intra-class compactness and inter-class separability in the learned feature space. This protocol simultaneously reflects few-shot adaptability and cross-channel robustness, forming, together with the aforementioned numerical simulation design, a reproducible AI-driven physical-layer modeling validation pipeline.

4 Experiments and Results

4.1 Dataset Preprocessing and Experimental Setup

In this study, the proposed Few-Shot FS-SEI method is evaluated on three open-source Radio Frequency Fingerprinting (RFF) datasets: D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} . The Wi-Fi dataset D_{Wi-Fi} is sourced from the publicly available ORACLE Wi-Fi dataset released by Sankhe et al. [20]. The experimental setup includes 16 identical USRP X310 software-defined radios (SDRs) deployed as transmitters to emulate a real-world scenario involving multiple homogeneous devices. The transmitted signals are generated using the MATLAB WLAN System Toolbox in compliance with the IEEE 802.11a standard. Each packet contains a random payload and identical address fields, and is streamed to the SDR transmitters. At the receiver side, another SDR captures the transmitted signals at a center frequency of 2.45 GHz with a sampling rate of 5 MS/s. The raw baseband I/Q complex signals and associated metadata are stored using the Signal Metadata Format (SigMF) structure. Each device provides approximately 20 million samples, with measurement distances ranging from 2 to 62 feet (approximately 0.6 to 18.9 m), covering a variety of distance and channel conditions. The ADS-B dataset D_{ADS-B} was released by Zhang [21], where the system captures 1090 MHz ADS-B broadcast signals in open airspace using a Signal Hound SM200B SDR. The signals are sampled at 50 MHz with a bandwidth of 10 MHz and a gain setting of 30 dB, collected using an omnidirectional antenna in an unobstructed environment. All signals are automatically decoded according to the DO-260B standard to obtain aircraft identifiers, which are labeled using the corresponding ICAO codes. The dataset includes 26,613 long signal instances across 1661 classes and 167,234 short signal instances across 1713 classes. Long signals consist of 6000 samples and short signals of 3000 samples, with each signal segmented into 4800 sample sequences. After filtering, 530 classes of long signals and 198 classes of short signals are retained in the final dataset. The D_{ADS-B}^{DF17} dataset, published by Zhang, contains Mode-S Extended Squitter (DF = 17) non-cooperative radar response signals. It comprises I/Q complex signals from 216 aircraft, sampled at 40 MHz. For the unprocessed raw packet format, alignment is first performed at the packet level using the synchronization sequence as an anchor. Subsequently, fixed length sliding windows are used to extract I/Q segments containing 6000 samples. This segment length is designed based on two principles: (1) to include statistical context before and after the packet, thereby enabling stable learning of slowly varying structures;

and (2) to balance information density and computational cost to ensure training efficiency. To mitigate distributional shifts caused by multiple devices and varying distances, balanced sampling is employed in the training set. Specifically, a maximum number of samples is set for each device distance configuration, and 3000 segments are uniformly selected across distance intervals. All extracted segments undergo z-score normalization to eliminate amplitude scale variations introduced by different time periods or hardware links, allowing the learned representations to focus on waveform shape and phase relationships rather than absolute energy. These preprocessing steps are consistently applied during episode construction across the train, validation, and test sets to prevent data leakage and ensure comparability of results.

Few-shot SEI is highly susceptible to distributional shifts arising from source entity variations and environmental diversity, such as distance, occlusion, multipath effects, and mobility. When support samples are limited, prototype estimation is prone to bias, which may destabilize the metric geometry. To mitigate this issue, a fixed data preprocessing protocol tightly coupled with the FS-SEI framework is established prior to training:

- **Fixed-Length Slicing:** I/Q segments of fixed length (as specified in [Table 1](#)) are extracted using the packet synchronization sequence as an anchor. This ensures sample homogeneity and enhances information density.
- **Z-Score Normalization:** Each segment is standardized using z-score normalization to eliminate variations in energy scale and gain. This emphasizes waveform shape and phase relationships over absolute amplitude, thereby reducing interference from scale differences in orthogonal feature decoupling and prototype-based metric learning.
- **Balanced Sampling:** Samples are uniformly drawn across distance intervals to maintain consistent proportions of devices and channel conditions. This mitigates distributional shift and ensures consistency across training, validation, and testing phases.

After the aforementioned preprocessing, the input distribution is aligned with the training objective, enabling the formation of discriminative and generalizable class geometries within the prototypical metric space.

Channel perturbation modeling and episodic evaluation protocols strictly follow the unified design described in [Section 3.3](#).

All experiments were conducted on a Windows platform with a 13th Gen Intel Core i5-13400 CPU, an NVIDIA GeForce RTX 3060 GPU (12 GB), and 32 GB RAM, using Python 3.10 and PyTorch 2.5.x. Model evaluation followed a fixed N-way K-shot episodic learning protocol. During testing, multiple independent episodes were randomly generated, and performance was reported as the mean \pm standard deviation of query-set accuracy across episodes to reduce stochastic sampling bias and capture episode-level uncertainty. Balanced sampling and z-score normalization were applied to ensure fair comparison across devices and distance conditions by mitigating signal energy scaling and gain variations. Channel effects were simulated by synchronously applying CFO, AWGN, and I/Q imbalance to both support and query samples within each episode, avoiding intra-episode distribution mismatch. Perturbation parameters were selected from discrete ranges and combined via a grid-based sweep, with perturbation configurations randomly sampled per episode and consistently applied across training, validation, and testing to ensure fairness and reproducibility. The CFO, SNR, and I/Q imbalance ranges were designed to emulate common wireless hardware and channel non-idealities, including oscillator offsets, Doppler effects, noise variations, and RF front-end amplitude and phase mismatches, and are summarized in [Table 2](#). We assume consistent receiver configurations within each dataset, with dataset-specific packet alignment performed when required (e.g., synchronization-based alignment for $D_{\text{ADS-B}}^{\text{DF17}}$). Accordingly, the evaluation focuses on robustness to channel

perturbations and RF front-end imperfections under a fixed receiver assumption, while receiver hardware heterogeneity is left for future work. Numerical stability issues were mitigated through normalization, controlled perturbation ranges, and Monte Carlo averaging across independent episodes. Overall, the proposed experimental design combines consistent data partitioning, controlled channel perturbation modeling, and Monte Carlo episodic evaluation (i.e., averaging performance over a large number of independently sampled episodes) to reduce data leakage and sampling bias, providing a reproducible and interpretable assessment despite the inherent limitations of simulated channel conditions. To facilitate cross-dataset comparison and improve clarity, Table 3 summarizes the key characteristics of the D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} datasets used in this study.

Table 2: Simulation and training parameters.

Simulation Parameter	Value
Signal slice length in D_{Wi-Fi} and D_{ADS-B}^{DF17}	6000
Signal slice length in D_{ADS-B}	4800
Dimension of samples	$2 \times \pi$
Number of categories in D_{Wi-Fi}	16
Number of categories in D_{ADS-B}	19
Number of categories in D_{ADS-B}^{DF17}	216
CFO (f_{Δ})	$\{-30, 50\}$
AWGN	$\{10, 30\}$
I/Q imbalance (Amplitude)	$\{-10, 10\}\%$
I/Q imbalance (Phase)	$\{-5, 5\}\%$
Number of samples per categories	$\{5, 10, 20, 30, 40\}$
Meta-iterations	20,000
Meta-batch size	8
Optimizer	AdamW
Activation function	ReLU
Learning rate	1×10^{-3}
Weight decay	1×10^{-2}
Warmup	1500 (iter)

Table 3: Summary of dataset characteristics.

	D_{Wi-Fi}	D_{ADS-B}	D_{ADS-B}^{DF17}
Center Frequency	2.45 GHz	1090 MHz	1090 MHz
Sample Rate	5 MS/s	50 MHz	40 MHz
Standard/Modulation	IEEE 802.11a	DO-260B	Mode-S Extended Squitter (DF = 17)
Devices/Classes	16 hardware devices	Long: 530 aircraft Short: 198 aircraft	216 aircraft
Signal Length	~20 M samples/device	Long: 6000 samples Short: 3000 samples	33,600–1,392,000 samples

(Continued)

Table 3 (continued)

	D_{Wi-Fi}	D_{ADS-B}	D_{ADS-B}^{DF17}
Hardware/Receiver	USRP X310 SDRs	Antennas and Signal Hound SM200B SDR	Mode-S Extended Squitter
Note	Distance: 2–62 ft	Bandwidth = 10 MHz; Baseband/IF = 0; Gain: 30 dB	Baseband/IF \approx 0

4.2 Experimental Results and Analysis

To evaluate the robustness and modeling capability of the proposed FS-SEI framework under multi-channel and few-shot conditions, this section presents and analyzes its overall performance on three open-source radio frequency datasets: D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} . The evaluation covers various k-shot task configurations and controlled channel perturbation scenarios to assess the model's convergence speed, classification stability, and generalization performance under cross-domain channel variations. The results demonstrate that the FS-SEI model, which integrates feature decoupling with prototypical metric learning, achieves high accuracy, low variance, and stable convergence across multiple real-world radio frequency datasets. These findings validate its feasibility and reproducibility as an AI-assisted physical-layer modeling framework. The model's performance under various datasets and channel perturbation conditions is illustrated in Figs. 3 and 4, which include confusion matrices, t-SNE visualizations, and convergence curves to depict classification accuracy, feature space geometry, and training stability. Under the 10-shot setting, the training curves indicate that the model converges rapidly across all datasets, with accuracy reaching saturation at an early stage. For the D_{Wi-Fi} dataset, convergence is slightly slower, which can be attributed to its higher intra-class signal variability and more pronounced channel condition differences. This requires the model to spend more time disentangling and aligning device-specific and channel-related components. In contrast, the D_{ADS-B} and D_{ADS-B}^{DF17} datasets exhibit simpler signal structures and higher channel consistency, resulting in flatter curves and consistently low loss early in training, reflecting the model's rapid adaptation to stable radio frequency features.

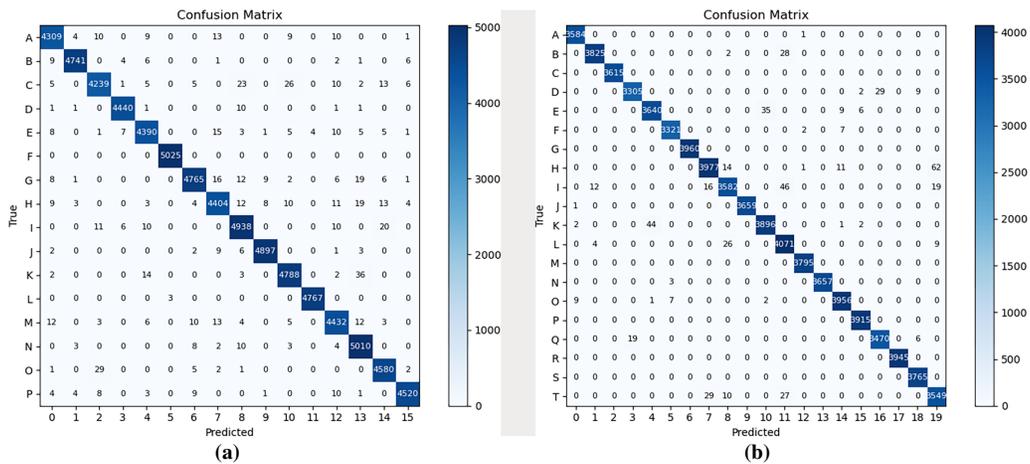


Figure 3: Confusion matrices under the 10-shot setting during testing: (a) D_{Wi-Fi} ; (b) D_{ADS-B} .

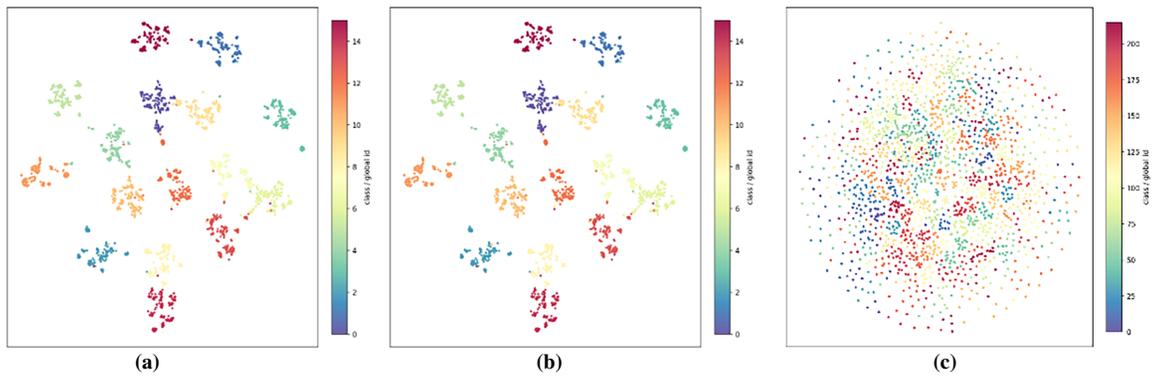


Figure 4: t-SNE visualization of RFF features under the 10-shot setting: (a) D_{Wi-Fi} ; (b) D_{ADS-B} ; (c) D_{ADS-B}^{DF17} .

The visualization reflects the intermediate embedding space prior to prototype-based metric inference. Classification accuracy reported in Table 4 is evaluated in the meta-learned prototypical decision space rather than in this low-dimensional projection.

Table 4: Test results of the proposed method on D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} under varying k-shot conditions.

Clients	D_{Wi-Fi}					D_{ADS-B}					D_{ADS-B}^{DF17}				
	5	10	20	30	40	5	10	20	30	40	5	10	20	30	40
Loss	0.067	0.028	0.040	0.093	0.031	0.036	0.027	0.027	0.028	0.030	0.008	0.003	0.008	0.006	0.007
Accuracy	0.977	0.991	0.990	0.983	0.988	0.993	0.994	0.993	0.992	0.993	0.997	0.998	0.997	0.998	0.998
Precision	0.977	0.992	0.990	0.983	0.988	0.993	0.994	0.993	0.993	0.994	0.997	0.999	0.998	0.998	0.994
Recall	0.977	0.992	0.990	0.983	0.988	0.993	0.994	0.993	0.993	0.994	0.997	0.999	0.998	0.998	0.998
Marco F1	0.977	0.992	0.990	0.983	0.988	0.993	0.994	0.993	0.993	0.994	0.997	0.999	0.998	0.998	0.998

Fig. 4 presents the t-SNE visualization of RFF features under the 10-shot setting, reflecting the feature distributions across different datasets. Subfigure (a) illustrates the feature distribution of the D_{Wi-Fi} dataset, where samples from different classes form distinct and well-separated clusters. This indicates that in relatively stable environments with consistent signal patterns, the model can effectively extract highly discriminative device fingerprints. Subfigure (b) shows the results for the D_{ADS-B} dataset. Although the feature distribution appears slightly more dispersed, samples from the same class still exhibit clear clustering behavior. This suggests that under cross-protocol conditions, the RFF features retain a meaningful degree of discriminative capability. However, in subfigure (c), corresponding to the D_{ADS-B}^{DF17} dataset, the feature distribution is heavily affected by significant distributional shifts and highly variable communication conditions. Although the number of data points increases substantially, samples from different classes are highly entangled with blurred cluster boundaries. However, in subfigure (c), corresponding to the D_{ADS-B}^{DF17} dataset, the t-SNE visualization appears visually entangled with blurred cluster boundaries. This phenomenon is primarily attributed to the extreme multi-class setting (216 classes), the large number of samples, and the projection of high-dimensional embeddings into a two-dimensional space. It should be noted that this visualization reflects the intermediate embedding space prior to prototype-based metric inference and is intended for qualitative analysis only. The final classification accuracy reported in Table 4 is computed in the meta-learned prototypical decision space, where device-specific features are aggregated into class prototypes and queried via distance-based inference. Therefore, the visually entangled appearance in the t-distributed Stochastic Neighbor Embedding (t-SNE) projection does not imply poor separability in the final decision space, and is

not contradictory to the near-perfect accuracy observed for the D_{ADS-B}^{DF17} dataset. Error analysis on the D_{ADS-B}^{DF17} dataset indicates that misclassifications are not uniformly distributed but are concentrated in a small number of confusion pairs and dominant confusion targets, reflecting near-neighbor or centralized confusion rather than systematic failure. This behavior is further supported by the Top-k results, where accuracy reaches near-perfect performance at k is 5, suggesting that most errors correspond to local prototype ambiguity. In addition, misclassifications are mainly associated with tail classes, while only a few high-confidence failures are observed, which is consistent with the overall good calibration behavior of the model.

Due to its lightweight architectural design, the proposed model contains only 144,712 parameters, resulting in minimal increases in FLOPs. The computational load is significantly lower than that of typical convolutional neural networks, transformers, and other deep architectures [9,14]. This demonstrates that the model maintains both high identification performance and computational efficiency under few-shot conditions.

Table 2 summarizes the quantitative results across different values of k . When k increases from 5 to 10, there is a marked improvement in accuracy and a reduction in variance, indicating more stable prototype estimation and sharper decision boundaries. As k further increases to the range of 20 to 40, the overall accuracy begins to saturate, suggesting that under the current perturbation and intra-class variation conditions, the 10-shot setting is sufficient to capture the essential prototype features, with diminishing returns from additional samples. This trend is consistently observed across all three datasets, demonstrating that the proposed framework achieves rapid convergence and stable generalization even under extremely low-sample conditions.

Based on the results in Table 2, the proposed FS-SEI model demonstrates high accuracy and stability across all three open-source radio frequency datasets. Overall, the 10-shot setting offers the best trade-off between performance and computational cost, achieving high accuracy rates of 0.991, 0.994, and 0.998 on the D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} datasets, respectively. Episode-level fluctuations remain below 0.02, indicating stable convergence and reproducibility. Under the low-resource 5-shot condition, the model still maintains high precision, confirming its practical viability and robustness against channel perturbations in few-shot scenarios. As the number of support samples increases to 20–40, the improvements in accuracy become marginal or even slightly decline, suggesting that 10-shot is sufficient to capture the core prototype features, with limited marginal benefits from additional samples. Across all three datasets, Precision, Recall, and Macro-F1 closely align with Accuracy, with differences less than 0.001. This indicates balanced classification performance with compact intra-class clustering and clear inter-class separation. Additionally, the loss is consistently lowest under the 10-shot setting, reflecting the most efficient convergence at this configuration. Overall, the integration of feature decoupling and prototypical metric learning effectively suppresses decision boundary drift induced by channel effects, preserving a stable class geometry across domain shifts. As a result, FS-SEI achieves both high accuracy and stability in data-constrained SEI scenarios, establishing a reproducible and cross-channel resilient performance benchmark for physical-layer security and device-level authentication in future 6G systems. In comparison with existing FS-SEI approaches, the proposed method demonstrates consistent performance improvements across various settings. Specifically, under the 10-shot scenario on the D_{Wi-Fi} dataset, as shown in Fig. 5, our method achieves an average accuracy of 99.16 percent, outperforming ADPMAE. Even in the 20-shot condition, it maintains a 0.38 percentage point advantage over the same baseline, indicating that the geometry of the prototype space remains more stable and the decision boundaries less affected by episode sampling and channel condition variations. This improvement is attributed to the tight coupling of orthogonal feature decoupling and prototypical metric learning, reinforced by episodic aggregation and channel-aware enhancement mechanisms. Together, these reduce variance in prototype estimation and mitigate cross-channel drift. The combined gains in accuracy and stability confirm

the method's superior ability to maintain intra-class compactness and inter-class separability under few-shot conditions. In summary, under few-shot and cross-channel conditions, the proposed method consistently outperforms recent state-of-the-art FS-SEI approaches within a unified evaluation protocol. It achieves lower result variance and higher reproducibility, further supporting the effectiveness of the proposed design in providing more reliable prototype geometry and decision quality when confronted with source heterogeneity and environmental variability.

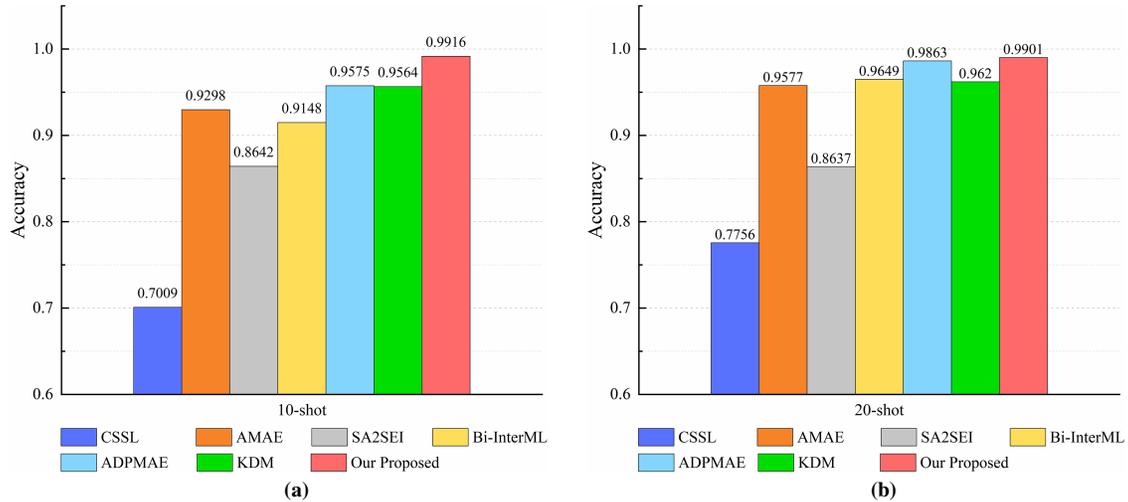


Figure 5: Comparison with recent RFF FS-SEI methods in D_{Wi-Fi} : (a) at 10-shot; (b) at 20-shot.

On the other hand, simulation results demonstrate the high value of the proposed architecture for classifying RFF device samples. The method integrates feature decoupling, lightweight temporal fusion, and a prototype-based decision module, achieving a balance between identification performance and computational efficiency. In contrast to common approaches that use short windows combined with deep stacking, which introduces additional computational cost, or rely on computationally expensive CNN Transformer architectures, the proposed method avoids the quadratic complexity of global self-attention with respect to sequence length. As a result, both training and inference exhibit computational complexity that scales more linearly. This architecture offers high information density and low resource demand on edge devices, aligning with the energy efficiency and deploy ability requirements of sixth-generation communication systems.

This study has several validity-related limitations that should be acknowledged. Regarding internal validity, the proposed model is evaluated under episodic few-shot learning with controlled experimental settings, which may not capture all sources of randomness in real-world deployments. In terms of construct validity, although real-world RF datasets are used, channel impairments such as CFO, AWGN, and IQ imbalance are applied in a controlled *in-silico* manner to enable reproducible analysis of channel robustness. With respect to external validity, the current evaluation assumes a fixed receiver hardware chain, and cross-receiver heterogeneity is not explicitly modeled. These factors may limit direct generalization to heterogeneous receiver environments and motivate future extensions of the proposed framework.

5 Conclusion

This study addresses the problem of SEI under highly dynamic channel conditions by proposing a few-shot modeling framework that integrates feature disentanglement, prototypical metric learning, and meta-learning. The proposed architecture adopts a CNN Transformer hybrid backbone with a dual-branch

disentanglement module, enabling explicit separation of device-invariant and channel-variant factors at the feature level. Orthogonality constraints are applied to enforce their independence, allowing the model to stably capture the distinctiveness of Radio Frequency Fingerprints under conditions of severe channel variability and data scarcity. Through controlled channel perturbations and parameter sweep-based numerical simulations, the model autonomously learns stable representations under varying channel conditions, demonstrating the feasibility of AI-assisted physical-layer modeling. Experimental results show that FS-SEI achieves high accuracy, low variance, and stable convergence across multiple real-world radio frequency datasets including D_{Wi-Fi} , D_{ADS-B} , and D_{ADS-B}^{DF17} , confirming the framework's cross-channel robustness and reproducibility. Furthermore, the model maintains computational efficiency under a lightweight design, making it suitable for deployment on edge devices and practically relevant for meeting the energy efficiency and low latency requirements of sixth-generation communication environments. Comprehensive quantitative and geometric analyses confirm that the feature disentanglement and prototypical metric mechanisms effectively suppress decision drift caused by channel effects, preserve discriminative geometry with tight intra-class clustering and clear inter-class separation, and enable stable convergence with strong generalization. For future sixth-generation network architectures, the capabilities of FS-SEI extend beyond conventional device identification and are poised to serve as a core physical-layer security component within AI-native 6G systems. In scenarios such as cell-free massive multiple-input multiple-output or reconfigurable intelligent surface-assisted environments, edge nodes must rapidly verify hundreds or even thousands of IoT devices to prevent unauthorized or impersonated nodes from entering the system. In dynamic channel conditions involving satellites, unmanned aerial vehicles, or terrestrial relays, nodes must further maintain reliable identity recognition despite severe channel fluctuations. The FS-SEI architecture proposed in this study consistently maintains robust identification performance under few-shot settings, cross-channel perturbations, and high-variability scenarios, indicating its direct applicability to real-world sixth-generation deployment contexts. It provides a strong foundation for the development of next-generation lightweight physical-layer authentication mechanisms. Overall, this research establishes a reproducible, interpretable, and channel-stable computational modeling benchmark for few-shot SEI, and demonstrates its potential as a core enabling technology for physical-layer modeling and security in 6G systems. The proposed approach offers a stable, scalable, and self-adaptive solution for communication modeling and device-level identity authentication in future intelligent wireless and edge networks.

Channel perturbations in this study are constructed in a controlled *in-silico* manner (CFO/AWGN/IQ imbalance), which improves reproducibility but does not fully capture more complex real-world propagation effects and receiver-side hardware heterogeneity; therefore, the conclusions mainly reflect few-shot SEI behavior under a fixed receiver-chain assumption. In addition, prototype-based inference requires computing distances between each query and all class prototypes, leading to inference cost that grows approximately linearly with N-way and may affect real-time scalability when the number of classes becomes very large. Future work will focus on: (i) modeling receiver heterogeneity and enabling online adaptation; (ii) extending channel modeling and/or conducting hardware-in-the-loop validation; and (iii) improving large-scale inference efficiency via hierarchical prototypes, candidate set reduction, or approximate nearest neighbor retrieval.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: The authors confirm contribution to the paper as follows: study conception and design: Chia-Hui Liu; data collection: Chia-Hui Liu and Hao-Feng Liu; analysis and interpretation of results: Chia-Hui Liu; draft manuscript preparation: Chia-Hui Liu. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The datasets used in this study are publicly available open-source RFF datasets, including the ORACLE Wi-Fi dataset [20] and ADS-B datasets.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Ding L, Wang S, Wang F, Zhang W. Specific emitter identification via convolutional neural networks. *IEEE Commun Lett.* 2018;22(12):2591–4. doi:10.1109/LCOMM.2018.2871465.
2. Merchant K, Revay S, Stantchev G, Nousain B. Deep learning for RF device fingerprinting in cognitive communication networks. *IEEE J Sel Top Signal Process.* 2018;12(1):160–7. doi:10.1109/JSTSP.2018.2796446.
3. Xie F, Wen H, Wu J, Chen S, Hou W, Jiang Y. Convolution based feature extraction for edge computing access authentication. *IEEE Trans Netw Sci Eng.* 2020;7(4):2336–46. doi:10.1109/TNSE.2019.2957323.
4. Wang Y, Gui G, Gacanin H, Ohtsuki T, Dobre OA, Poor HV. An efficient specific emitter identification method based on complex-valued neural networks and network compression. *IEEE J Sel Areas Commun.* 2021;39(8):2305–17. doi:10.1109/JSAC.2021.3087243.
5. He B, Wang F. Cooperative specific emitter identification via multiple distorted receivers. *IEEE Trans Inf Forensics Secur.* 2020;15:3791–806. doi:10.1109/TIFS.2020.3001721.
6. Li D, Yao B, Sun P, Li P, Yan J, Wang J. Specific emitter identification based on ensemble domain adversarial neural network in multi-domain environments. *EURASIP J Adv Signal Process.* 2024;2024(42):1–15. doi:10.1186/s13634-024-01138-y.
7. Yang J, Zhu S, Wen Z, Li Q. Cross-receiver radio frequency fingerprint identification: a source-free adaptation approach. *Sensor.* 2025;25(14):4451. doi:10.3390/s25144451.
8. Liu M, Chai Y, Li M, Wang J, Zhao N. Transfer learning-based specific emitter identification for ADS-B over satellite system. *Remote Sens.* 2024;16(12):2068. doi:10.3390/rs16122068.
9. Zhao D, Yang J, Liu H, Huang K. A complex-valued self-supervised learning-based method for specific emitter identification. *Entropy.* 2022;24(7):851. doi:10.3390/e24070851.
10. Yao Z, Fu X, Guo L, Wang Y, Lin Y, Shi S, et al. Few-shot specific emitter identification using asymmetric masked auto-encoder. *IEEE Commun Lett.* 2023;27(10):2657–61. doi:10.1109/LCOMM.2023.3312669.
11. Liu C, Fu X, Wang Y, Guo L, Liu Y, Lin Y, et al. Overcoming data limitations: a few-shot specific emitter identification method using self-supervised learning and adversarial augmentation. *IEEE Trans Inf Forensics Secur.* 2024;19:500–13. doi:10.1109/TIFS.2023.3324394.
12. Fu X, Peng Y, Liu Y, Lin Y, Gui G, Gacanin H, et al. Semi-supervised specific emitter identification method using metric-adversarial training. *IEEE Internet Things J.* 2023;10(12):10778–89. doi:10.1109/JIOT.2023.3240242.
13. Mu S, Zu Y, Chen S, Yang S, Feng Z, Zhang J. Few-shot metric learning with time-frequency fusion for specific emitter identification. *Remote Sens.* 2024;16(24):4635. doi:10.3390/rs16244635.
14. Zhou Z, Li G, Wang T, Zeng D, Li X, Wang Q. A robust open-set specific emitter identification for complex signals with class-irrelevant features. *IEEE Trans Inf Forensics Secur.* 2025;20:6058–73. doi:10.1109/TIFS.2025.3576573.
15. Yang S, Yu S, Li Q, Xia K, Zhu H. Few-shot specific emitter identification via asymmetric dual-path masked autoencoder. *Digit Signal Process.* 2025;163:105201. doi:10.1016/j.dsp.2025.105201.
16. Sun M, Teng J, Liu X, Wang W, Huang X. Few-shot specific emitter identification: a knowledge, data, and model-driven fusion framework. *IEEE Trans Inf Forensics Secur.* 2025;20:3247–59. doi:10.1109/TIFS.2025.3550080.
17. Sun M, Wei L, Yu C, Qiu Z, Teng J. Self-supervised and semi-supervised learning for few-shot specific emitter identification using CNN-transformer with virtual adversarial training. *Appl Intell.* 2025;55:799. doi:10.1007/s10489-025-06645-5.
18. Hu N, Li X, Zhao Z, Fang M, Liu J. Multi-binary network with feature disentanglement for open-set specific emitter identification. *Signal Process.* 2026;240:110344. doi:10.1016/j.sigpro.2025.110344.

19. Xu L, Fu X, Wang Y, Zhang Q, Zhao H, Lin Y, et al. Enhanced few-shot specific emitter identification via phase shift prediction and decoupling. *IEEE Trans Cogn Commun Netw.* 2025;11(1):145–55. doi:10.1109/TCCN.2024.3435886.
20. Sankhe K, Belgiovine M, Zhou F, Riyaz S, Ioannidis S, Chowdhury K. ORACLE: optimized radio classification through convolutional neural networks. *IEEE Infocom.* 2019;2019:370–8. doi:10.1109/INFOCOM.2019.8737463.
21. Zhang X. ADS-B real-world data set (DF = 17). *Sci Data Bank.* 2023. doi:10.57760/sciencedb.o00009.00481.