

BeamID: Domain-Adaptive Radio Fingerprinting with MIMO Beamforming Feedback

Khandaker Foysal Haque[†], Francesca Meneghello[†], Francesco Restuccia[†]

[†] Institute for Intelligent Networked Systems, Northeastern University, United States

Abstract—Client identification at the network control plane has gained increasing attention for supporting monitoring and securing workflows in wireless deployment. Interestingly, device-specific hardware characteristics in the network interface card (NIC) introduce systematic impairments in channel estimation that percolate into multiple-input, multiple-output (MIMO) beamforming feedback. Recent work has shown that beamforming feedback matrices reconstructed from passively captured beamforming feedback information (BFI) frames can be exploited for accurate radio fingerprinting (RFP). However, existing fingerprinting approaches lead to severe performance degradation under domain shifts caused by changes in the propagation channel or deployment configuration. In this paper, we present **BeamID**, a domain-adaptive RFP framework that enables reliable cross-environment identification from beamforming feedback. **BeamID** learns radio-discriminative embeddings in a source domain using supervised contrastive learning and performs few-shot adaptation directly in embedding space to align target-domain representations without full retraining. This design explicitly addresses environment-induced distribution shifts while preserving inter-radio separability. We evaluate **BeamID** using real Wi-Fi measurements collected across nine distinct locations with fifteen different NICs, demonstrating that few-shot adaptation substantially restores cross-domain RFP accuracy. With as few as five labeled samples per client in a new unseen environment, **BeamID** achieves up to $\sim 99\%$ identification accuracy while requiring only ~ 2.3 seconds of adaptation time.

I. INTRODUCTION

Wireless networks have become a fundamental component of modern networking infrastructure, supporting an ever-growing number of heterogeneous devices, including smartphones, laptops, smart appliances, and embedded systems. As the number and diversity of the wireless devices continue to increase, the ability to reliably identify and authenticate devices at the physical layer has gained significant attention, with implications spanning spectrum management, security, and enforcement of network-policies [1], [2]. Motivated by this, radio fingerprinting (RFP) has emerged as a promising alternative to cryptographic authentication, as it exploits hardware-dependent imperfections that wireless signals inherit without requiring the exchange of secret keys [3], [4].

Recent advances in RFP have demonstrated that beamforming feedback contains rich device-specific information that can be leveraged for client or wireless access point (AP) identification using deep learning techniques [3], [5].¹ While these approaches achieve high identification accuracy when

¹Throughout this paper, we use the terms *client*, *radio device*, and *network interface card (NIC)* interchangeably to refer to the physical wireless device participating in the network.




Fig. 1: Domain shift in RFP— models trained in one environment fail to generalize to unseen domains.

training and testing are performed in the same environment, this assumption rarely holds in practice. Changes in location, antenna orientation, or surrounding objects alter the propagation channel, inducing systematic distribution shifts in observed fingerprints even when the underlying device remains unchanged. As a result, models trained in one environment fail when evaluated in previously unseen environments, as depicted in Figure 1.

Preliminary results presented in Section II-C demonstrate that a convolutional neural network (CNN)-based RFP model ([5]) trained and evaluated within the same domain can achieve excellent performance, yet fail when deployed in an unseen domain. This is because domain shifts cause previously learned decision boundaries to become unreliable, leading to severe performance degradation. As a result, existing RFP systems often require extensive retraining or large amounts of labeled data whenever the deployment context changes, limiting scalability and hindering real-world adoption.

In this paper, we move beyond domain-specific fingerprinting and address the problem of *domain-adaptive* wireless client identification. We formulate cross-domain RFP as a few-shot adaptation problem [6], [7] and propose a learning-based framework named **BeamID**. It first learns client-centric representations in a source domain and then aligns them to a target domain through lightweight adaptation in embedding space with only a few labeled samples from the unseen target domain. **BeamID** is designed to operate as a controller-level fingerprinting service, enabling client identification to be

integrated into network monitoring and security workflows, as discussed in Section II-B.

Summary of Contributions

- We propose *BeamID*, a domain-adaptive RFP framework that exploits standard-compliant multiple-input, multiple-output (MIMO) beamforming feedback to enable reliable client identification across heterogeneous locations via few-shot embedding adaptation.
- We introduce an embedding-based few-shot adaptation strategy [8], [9] built upon supervised contrastive learning [10], [11], which aligns target-domain fingerprints with source-domain client-centric representations using only a small number of labeled samples per device. We further provide quantitative embedding-space analysis that explains the effectiveness of the proposed domain adaptation.
- We conduct extensive experiments with a Wi-Fi testbed, collecting beamforming feedback with 15 different NICs (each representing a client) at each of 9 distinct locations, demonstrating that *BeamID* enables reliable cross-domain client identification performance. With as few as 5 target-domain (unseen) samples per client, *BeamID* achieves up to 99% accuracy while requiring only 2.3 seconds of adaptation time. For reproducibility, we open-source our code and dataset at: <https://github.com/Restuccia-Group/BeamID>.

II. BACKGROUND AND PRELIMINARY ANALYSIS

A. Beamforming Feedback for RFP

Modern wireless systems including Wi-Fi 5/6 (IEEE 802.11ac/ax) employ explicit channel sounding as presented in Figure 2 to enable orthogonal frequency-division multiplexing (OFDM) MIMO transmissions [12], [13]. During this process, the AP periodically broadcasts null data packets (NDPs) (step 1 of Figure 2) containing known training symbols [12], [13]. Using these NDPs, each client estimates the channel frequency response (CFR) (step 2 of Figure 2) of the underlying MIMO links between the AP and its receive antennas over the OFDM sub-channels. The resulting channel estimates can be represented as

$$\mathbf{H} \in \mathbb{C}^{K \times M \times N}, \quad (1)$$

where K denotes the number of sub-channels and M and N are the numbers of transmit and receive antennas, respectively. Since the NDP is not beamformed, the resulting channel estimates are not affected by inter-user or inter-stream interference. Rather than feeding back the full CFR, most modern MIMO systems [12], [13] mandate a compression procedure to reduce signaling overhead. For each sub-channel k , the CFR $\mathbf{H}_k$ is decomposed, and a unitary beamforming matrix $\mathbf{V}_k$ is constructed from the dominant spatial directions. This matrix, referred to as *beamforming feedback matrix*, is sufficient for the AP to compute downlink precoding weights.

To further reduce feedback size, $\mathbf{V}_k$ is parameterized using a sequence of phase and rotation angles [12], [13], which




Fig. 2: Leveraging Wi-BFI tool to extract beamforming feedback matrices from standard compliant channel sounding.

are uniformly quantized and transmitted in compressed beamforming feedback information (BFI) frames. Due to the strict latency constraints of MIMO precoding, typically requiring channel updates every ~ 10 ms—BFI frames are transmitted in clear text (step 3 of Figure 2) and can be passively captured by any device monitoring the wireless channel [14]. Leveraging tools such as Wi-BFI [15], we can reconstruct the beamforming feedback matrices $\mathbf{V}_k$ from the captured BFI frames (step 4 of Figure 2). Although $\mathbf{V}_k$ represents a form of compressed CFR, any hardware-specific artifacts, including hardware impairments of the NICs, percolate from CFR to the $\mathbf{V}_k$ matrices. This introduces consistent, client-specific distortions in the reported $\mathbf{V}_k$ matrices that persist across time and environments. We leverage representation learning through embeddings of the $\mathbf{V}_k$ matrices to fingerprint clients, which forms the basis for the domain-adaptive learning framework described in Section IV.

B. Network Integration and Security

We design *BeamID* as a deployable network function rather than a standalone learning-based fingerprinting mechanism. As shown in Figure 3, *BeamID* operates within the network controller and does not require any modification to AP or clients. The network consists of standard-compliant APs and associated clients, while monitoring telemetry derived from passively observed beamforming feedback is made available to the controller through out-of-band monitoring with Wi-BFI [15] tool, without affecting network operation.

Within the controller, *BeamID* runs as a softwareized fingerprinting service that infers device identity and produces confidence-aware identification outputs. These outputs are used by a security engine responsible for client identity verification and access policy evaluation. *BeamID* operates as a fingerprinting service that provides client identity evidence to existing control and management workflows without altering their internal operation. From a control-plane perspective, the interaction between *BeamID* and the network controller follows the workflow illustrated in Figure 3. Reconstructed beamforming feedback matrix, $\mathbf{V}_k$ is provided to the *BeamID*




Fig. 3: Controller-level integration of BeamID for policy-driven wireless client security.

module as input during the inference. In turn, BeamID outputs client identity estimates and confidence indicators that are used by the controller’s security engine to trigger policy evaluation, enforcement actions, or adaptation. Based on the security engine’s decision, control actions are issued to the enforcement plane, including allowing access, restricting connectivity or limiting access. In this role, BeamID provides client-specific identity evidence derived from physical-layer measurements, complementing higher-layer authentication mechanisms by detecting device substitution or misassociation under stable credentials.

Few-shot domain adaptation is handled as a runtime control operation. When confidence degradation or distribution drift is detected, the controller requests a small number of labeled samples to adapt BeamID without retraining from scratch. This enables continuous and environment-aware client identification with low overhead, making BeamID suitable for wireless networks with dynamic deployment conditions.

C. Preliminary Analysis on Domain Adaptability

We first conduct a preliminary analysis to examine how domain shift affects RFP when no explicit adaptation mechanism is employed. To this end, we train a simple CNN consisting of three VGG-style [16] convolutional blocks using beamforming feedback matrices to identify 15 different NICs. The trained model is then evaluated in two different scenarios: (i) testing on data collected in the same location as the training data, and (ii) testing on data collected in a different location.

Figure 4 summarizes the resulting identification accuracy across five locations (domains). When training and testing are performed within the same location, the model achieves consistently high accuracy, approaching near-perfect performance. This result confirms that beamforming feedback indeed carries strong client-specific signatures that can be captured by a simple deep learning model under the same domain. However, when the same model is evaluated on data collected in a different location, accuracy drops dramatically across all tested locations. Despite the underlying NIC remaining unchanged,




Fig. 4: Learning based RFP performance trained and tested in the same vs different environment

variations in propagation conditions and the surrounding environment significantly alter the observed fingerprints, causing the learned decision boundaries to fail.

This preliminary experiment demonstrates that domain shift is a fundamental challenge in RFP that motivates the need for a domain-adaptive fingerprinting framework, which we address in the following sections.

III. RELATED WORKS

Over the past decade, prior work has investigated a broad spectrum of fingerprinting approaches across different wireless technologies, signal representations, and learning paradigms. In this section, we review the most relevant lines of work and highlight how existing approaches differ from BeamID in terms of signal primitives, scalability, and robustness to domain shifts.

Classical RFP with hand-crafted features: Early RFP methods relied on manually designed features extracted from raw radio measurements, including transient behavior, constellation statistics, and other signal-domain descriptors [17]–[19]. While effective in controlled settings, these approaches require substantial domain expertise and careful feature tuning, and they typically struggle to scale with the number of clients and to generalize across different propagation conditions and network configurations.

Deep neural network (DNN) based RFP: Motivated by the limitations of manual feature engineering, a broad line of work has proposed DNN-based fingerprinting frameworks that learn discriminative representations jointly with the identification task [20]–[24]. For low-power IoT technologies, prior work has shown that DNN models can fingerprint ZigBee and LoRa radios directly from received signal representations [23], [24]. In Wi-Fi, ORACLE [21] and related schemes [22] introduce controlled transmitter-side impairments to amplify device-specific signatures. However, this strategy can increase bit error rate (BER) if impairments are not compensated, and it requires the ability to manipulate the transmitter waveform. A few complementary work focuses on mitigating channel effects via learned compensation. In particular, Restuccia and Melodia et al. [20], [25] propose augmenting DNNs with channel-compensation components like finite impulse response filtering to reduce the distortion induced by the propagation channel and improve robustness.

MIMO-assisted signal reconstruction for RFP: A few recent work leverages spatial diversity to reduce channel-induced distortion before fingerprint extraction [26], [27]. These approaches exploit multiple antennas and space-time diversity to reconstruct a less-distorted estimate of the transmitted signal based on receiver-side channel estimation, and then perform classification on the reconstructed signal. While promising, these designs have been primarily evaluated via simulations [26], [27], leaving their real-world applicability and sensitivity to implementation artifacts unclear. Moreover, they rely on specific diversity-oriented coding structures (e.g., Alamouti/Tarokh), which limit generality across commodity Wi-Fi devices.

Beamforming-feedback-based wireless client fingerprinting: The work most closely related to the proposed approach is DeepCSI [3] and its extension DeepCSiv2 [5]. These studies demonstrate that beamforming feedback contains stable, client-specific artifacts that can be exploited for accurate device identification. These methods achieve high accuracy when trained and deployed in the same propagation conditions, but their performance degrades sharply under domain shift caused by changes in the environment and deployment configuration.

Novelty of BeamID: *In contrast to prior work that (i) depends on hand-crafted features [17]–[19], (ii) assumes controllable transmitter impairments [21], [22], and (iii) relies on diversity-coded signal reconstruction evaluated mainly in simulation [26], [27], BeamID targets domain-adaptive fingerprinting from passively captured, standard-compliant beamforming feedback. By learning class (client) centric embeddings via supervised contrastive learning and performing lightweight few-shot alignment in embedding space, BeamID explicitly addresses environment-induced distribution shifts without full retraining, enabling practical cross-domain device identification.*

IV. DOMAIN-ADAPTIVE FINGERPRINTING FRAMEWORK

A. System Overview

The proposed BeamID framework performs RFP by learning client-discriminative representations directly from beamforming feedback matrices $\mathbf{V}_k$. As described in Section II-A, during standard channel sounding, each client periodically reports BFI, from which the corresponding $\mathbf{V}_k$ matrices can be reconstructed. These matrices constitute the fundamental input to the BeamID fingerprinting system. As illustrated in Figure 5, BeamID stems from five main steps presented as follows—

Step 1: Data aggregation: BeamID framework begins by aggregating the reconstructed $\mathbf{V}_k$ matrices across time into a fixed-size tensor representation. Each tensor corresponds to a short observation window and captures the spatial response characteristics of a single client. This step forms a fixed-size input by stacking beamforming feedback collected over a fixed observation window.

Step 2: Embedding encoder: The aggregated tensor is processed by an embedding encoder $f_\theta(\cdot)$, implemented as




Fig. 5: BeamID system overview

a CNN. The encoder maps each input tensor to a low-dimensional embedding vector $\mathbf{z} \in \mathbb{R}^d$, which serves as a compact representation of the client's fingerprint. As presented in Figure 6, the encoder consists of three stacked 1D convolutional blocks that progressively extract client-relevant patterns from the input tensor of $\mathbf{V}_k$, followed by global average pooling to produce a fixed-length embedding. The resulting representation preserves client-specific separability while remaining amenable to alignment under domain shift.




Fig. 6: Encoder architecture for $\mathbf{V}_k$ Embedding

Step 3: Source-domain representation learning: In the source-domain training phase, the encoder in Figure 6 is trained using labeled data collected in a reference environment (source-domain). Supervised contrastive learning is employed to shape the embedding space such that samples originating from the same clients form compact, class-centric clusters, while embeddings corresponding to different clients are well separated. A classification head $g(\cdot)$ is trained jointly to enable radio identification from the learned embeddings. Further details of the training objective and optimization procedure are provided in Section IV-B.




Fig. 7: Cosine similarity distributions of source-target embedding pairs before and after few-shot adaptation

Step 4: Few-shot domain adaptation: When the trained model is deployed in a new, unseen environment, changes in propagation conditions cause target-domain embeddings to deviate from their source-domain counterparts. To address this mismatch, BeamID performs few-shot domain adaptation using only a small number of labeled samples per device collected in the target domain. Adaptation is carried out through lightweight fine-tuning in embedding space, aligning target-domain samples with their corresponding source-domain clusters while preserving inter-client separation.

Step 5: Inference: After adaptation, inference is performed using the fine-tuned embedding encoder and classifier without further model retraining. Client identities are inferred directly from passively observed beamforming feedback at the controller, enabling continuous identification during normal network operation. This design allows BeamID to provide timely and low-overhead identity evidence that is directly used by *Enforcement Plane* to perform control and security functions in dynamic network deployments.

B. Source-Domain Representation Learning

We formalize source-domain training as a supervised representation learning problem over beamforming feedback. As described in Section II-A, the reconstructed beamforming matrices $\mathbf{V}_k$ are aggregated over time to form fixed-size input tensors. Let $\mathbf{x}_i$ denote such an aggregated tensor derived from beamforming feedback, and let $y_i \in \{1, \dots, C\}$ denote its corresponding client identity. A lightweight convolutional encoder $f_\theta(\cdot)$ maps each input $\mathbf{x}_i$ to a low-dimensional embedding vector $\mathbf{z}_i$, which represents the learned fingerprint of the corresponding client,

$$\mathbf{z}_i = f_\theta(\mathbf{x}_i), \quad \mathbf{z}_i \in \mathbb{R}^d. \quad (2)$$

Rather than relying solely on a classification objective, we explicitly shape the geometry of the embedding space using *supervised contrastive learning*, which directly enforces intra-class compactness and inter-class separation-properties that are essential for robustness under domain shift. Formally, given a mini-batch $\mathcal{B}$, the supervised contrastive loss for a sample i is defined as

$$\mathcal{L}_{\text{sup}}^{(i)} = -\frac{1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(\mathbf{z}_i^\top \mathbf{z}_p / \tau)}{\sum_{a \in \mathcal{B} \setminus \{i\}} \exp(\mathbf{z}_i^\top \mathbf{z}_a / \tau)}, \quad (3)$$

where $\mathcal{P}(i) = \{p \in \mathcal{B} | y_p = y_i, p \neq i\}$ is the set of positive samples sharing the same client identity as i , and τ is a temperature parameter. The temperature parameter τ controls the sharpness of the similarity distribution, with smaller values enforcing tighter intra-class clustering in the embedding space. This objective pulls embeddings of the same client together while pushing apart embeddings of different clients, independently of their temporal or spectral variability. To enable device identification during training, the contrastive objective is combined with a standard cross-entropy loss applied to a classification head $g(\cdot)$,

$$\mathcal{L}_{\text{CE}} = -\sum_{i \in \mathcal{B}} \log p(y_i | \mathbf{z}_i), \quad (4)$$

where $p(y_i | \mathbf{z}_i)$ denotes the softmax output of the classifier. The overall training objective is

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda \mathcal{L}_{\text{CE}}, \quad (5)$$

where λ controls the relative contribution of the classification loss. While the classifier is used to assign client identities during inference, the primary role of source-domain training is to construct a structured, class-centric embedding space. This embedding space captures client/NIC-dependent characteristics of beamforming feedback while suppressing variations caused by channel dynamics. Once source-domain training is complete, the learned embedding geometry serves as a stable reference for target-domain adaptation, allowing new samples to be aligned using only a small number of labeled observations, as discussed next (Section IV-C).

C. Few-Shot Domain Adaptation in Embedding Space

After source-domain learning, the learned embedding space captures client-discriminative structure under the reference environment (source-domain). When deployed in a new environment, environment-induced distribution shifts in the beamforming feedback cause target-domain (unseen environment) embeddings to misalign, leading to a pronounced drop in identification accuracy (Section II-C).

To address this challenge, we formulate cross-domain RFP as a *few-shot adaptation* problem performed directly in embedding space. Specifically, we assume access to a small number of labeled target-domain samples per device, which can be collected with minimal overhead during deployment. Rather than retraining the model from scratch, we leverage these

samples to align target-domain embeddings with the class-centric structure learned in the source-domain.

Let $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i)\}$ denote the labeled source-domain data used to train the embedding encoder $f_\theta(\cdot)$, and let $\mathcal{D}_t^{(K)} = \{(\mathbf{x}_j^t, y_j)\}_{j=1}^K$ denote the few-shot labeled target-domain samples, with K samples per client. During adaptation, the encoder parameters are fine-tuned using $\mathcal{D}_t^{(K)}$ while preserving the geometric structure of the embedding space learned in the source-domain.

Adaptation is performed by optimizing the same supervised contrastive objective used during source training, combined with a lightweight classification loss. By reusing the contrastive objective, target-domain embeddings corresponding to the same client are explicitly pulled toward each other and toward their corresponding source-domain clusters, while embeddings of different clients remain separated.

Figure 7 illustrates cosine similarity distributions between source and target domain embeddings in three different settings— same domain without adaptation, cross-domain without adaptation and cross-domain with few-shot adaptation. In the same-domain setting (no adaptation), embeddings from the same NICs (clients) exhibit consistently higher similarity than those from different NICs, yielding a clear separation. In cross-domain deployment without adaptation, this separation collapses— similarities of same-NIC and different-NIC pairs become closely overlapped, resulting in a small median gap of $\Delta_{\text{median}} = 0.074$. After few-shot adaptation, the similarity distributions shift in opposite directions, with same-NIC pairs moving toward higher similarity and different-NIC pairs toward lower similarity. This restores a pronounced separation ($\Delta_{\text{median}} = 0.234$), indicating improved cross-domain alignment while maintaining inter-device discrimination.

These distributions provide direct quantitative evidence that few-shot adaptation operates by realigning target-domain embeddings with their source-domain structure. The observed recovery of separation explains why high identification accuracy can be achieved with only a small number of labeled target samples.

V. EXPERIMENTAL EVALUATION

A. Experimental Setup

The experimental evaluation of BeamID was conducted using the WiSEC testbed, a large-scale, multi-NICs, and multi-band platform designed for testing Wi-Fi networks. WiSEC is equipped with an ASUS B250 motherboard featuring 15 PCIe slots, each hosting an Intel AX210 Wi-Fi NIC implementing the IEEE 802.11ax standard. The WiSEC testbed is depicted in Figure 8 (bottom), illustrating the hardware setup. For BeamID evaluation, the 15 WiSEC NICs were used as clients whereas, the AP was implemented through an ASUS RT-AX86U (AX5700) router.

A linux machine empowered with an Intel AX200 NIC and Wi-BFI tool [15] has been used as the monitor device to collect BFI frames and reconstruct the $\mathbf{V}_k$ matrices for fingerprinting. We collected the data for training the BeamID algorithm by considering a single-user MIMO network where only one




Fig. 8: WiSEC testbed and measurement setup. Measurements were collected for each 15 NIC across 9 spatially separated locations.

client (NIC) at a time was connected to the AP. However, the same data can be collected in a multi-user MIMO system. To generate traffic that triggers the transmission of beamforming feedback from the client to the AP, we established TCP *iperf* sessions from the client to the AP, simulating typical data transmissions. *To eliminate antenna and cable induced variability, all NICs were connected to the same antenna pair using identical cables.* As presented in Figure 8 (top), for each of the 15 NIC, we collect the data at 9 different locations with a spacing of 4 inches between each consecutive location. The whole data collection campaign is carried out through transmission on channel 36 for each of the NICs across the different experimental scenarios, considering 80 MHz of bandwidth for every instance. In each configuration (different NIC and location pair), approximately 1500 BFI frames were collected, providing sufficient samples for both training and evaluation.

B. BeamID Performance without few-shot learning (FSL)

We first evaluate BeamID without applying few-shot domain adaptation in order to quantify the effect of domain shift. The model is trained using data collected from a single location and evaluated either in the same environment or in a different, unseen environment.

Same-domain evaluation. Figure 9 reports the identification accuracy when training and testing are performed with data collected in the same location. Both BeamID and DeepCSIv2 achieve consistently high accuracy across all 9 locations. State-of-the-art (SOTA) approach DeepCSIv2 attains accuracy between 98.49% and 99.84% (average 99.12%), while BeamID




Fig. 9: BeamID (w/o few-shot adaptation) performance while training and testing on the same environment




Fig. 10: Cross-environment identification performance of BeamID without few-shot adaptation. Models are trained on the source location (x-axis) and evaluated on unseen locations.

achieves comparable performance ranging from 98.78% to 99.62% (average 99.09%).

Cross-domain evaluation. The performance drops sharply when models trained in one domain are evaluated in a different domain, as presented in Figure 10. DeepCSiv2 degrades to accuracy values between 6.01% and 7.35% (average 6.68%), indicating near-random classification. In contrast, BeamID consistently outperforms DeepCSiv2 across all locations even without FSL, achieving accuracy between 10.49% and 16.15% (average 13.01%), corresponding to an average absolute improvement of 6.33% and up to a $2.6\times$ relative gain. Although BeamID exhibits improved robustness to domain shift compared to DeepCSiv2, its accuracy remains insufficient for reliable cross-environment identification without FSL, motivating the need for the few-shot domain adaptation.

C. Cross-Domain Performance with FSL

1) Accuracy vs. Number of Target Samples (k -shot Curve):

Figure 11 shows the cross-domain identification accuracy of BeamID as a function of the number of labeled target-domain samples per device (k), averaged over all unseen locations for different source locations. With only $k = 5$ target samples, BeamID achieves high cross-environment accuracy ranging from approximately 97.4% to 98.8% across source locations. Increasing k to 15 raises accuracy above 99% for all the cases, indicating rapid adaptation with minimal supervision. Further increasing k yields diminishing returns—accuracy improves to about 99.3–99.5% at $k = 25$ and converges to 99.6–99.75% for $k \geq 45$. The performance variance across different source




Fig. 11: Cross-environment adaptation accuracy versus k , with models trained on different source locations (as indicated in the legend) and evaluated on the remaining unseen locations.




Fig. 12: Pairwise cross-domain transfer accuracy at $k = 15$. Each cell represents the accuracy obtained when training on a given source domain and testing on a target domain.

locations decreases sharply as k increases, dropping from approximately 1.4% at $k = 5$ to below 0.2% beyond $k = 35$. These results indicate that $k = 15$ is an optimal operating point, as it consistently achieves 99% accuracy with minimal additional benefit from larger k .

2) *Pairwise Cross-Domain Performance:* Figure 12 reports the pairwise cross-domain identification accuracy of BeamID at $k = 15$, where each entry corresponds to training on a given source domain and adapting to a target domain. The diagonal entries represent same-domain training and testing, while off-diagonal entries capture cross-domain transfer performance. Across all source-target pairs, BeamID achieves consistently high accuracy, with most off-diagonal values exceeding 98% and many approaching or reaching 99%. The limited variation across rows and columns indicates that adaptation performance is largely independent of the choice of source location, demonstrating stable and symmetric domain-transfer behavior.

3) *Accuracy vs. Number of Source Domains:* We next study how increasing the number of source (training) domains af-

fects cross-domain identification performance on unseen *target* (testing) domains. We consider scenarios $S2 - S8$, where S_k denotes training on k distinct source locations and evaluating on the remaining unseen target locations. The corresponding train-test domain splits are summarized in Table I. Reported results are averaged over all unseen target domains.

TABLE I: Experimental Scenarios with Training and Testing Locations

Experimental Scenario	Train loc.	Test loc.
S2	1, 2	3, 4, 5, 6, 7, 8, 9
S3	1, 2, 3	4, 5, 6, 7, 8, 9
S4	1, 2, 3, 4	5, 6, 7, 8, 9
S5	1, 2, 3, 4, 5	6, 7, 8, 9
S6	1, 2, 3, 4, 5, 6	7, 8, 9
S7	1, 2, 3, 4, 5, 6, 7	8, 9
S8	1, 2, 3, 4, 5, 6, 7, 8	9

As shown in Figure 13, increasing the number of source domains in BeamID provides limited gains in generalization to unseen target domains. For DeepCSiv2, mean accuracy improves modestly when increasing the number of source domains from $S2$ to $S4$ (from 11.0% to 22.4%). However, this trend is not consistent and remains highly unstable across scenarios. These results indicate that while exposure to multiple environments can partially improve generalization, source-domain diversity alone is insufficient to overcome environment-induced distribution shifts in channel information. In contrast, BeamID maintains consistently high accuracy on unseen target domains across all scenarios, achieving above 99% mean accuracy. This demonstrates that BeamID’s cross-domain performance is largely independent of the number of source domains used during training, owing to its embedding-based FSL mechanism.




Fig. 13: Accuracy versus number of source environments under domain shift. S_k denotes training on k distinct source environments, with performance averaged over all unseen target environments.

4) *Ablation Study—Fine-tuning With and Without Encoder:* We conduct an ablation study to assess the effect of adapting (fine-tuning) the encoder during few-shot domain adaptation. We compare two settings: (i) fine-tuning only the classifier while keeping the encoder fully frozen, and (ii) fine-tuning the classifier together with the *last layer of the encoder*. All results are reported for $k = 15$ target samples per device.

Figure 14 presents the results across different source-domain scenarios ($S2-S8$). When only the classifier is fine-tuned, the




Fig. 14: Accuracy versus number of source environments under domain shift. S_k denotes different training and testing scenarios as presented in Table I.




Fig. 15: Cross-domain identification performance of BeamID with and without fine-tuning the encoder.

model achieves lower and more variance in accuracy, with mean performance between 98.2% and 98.9% across scenarios. Enabling adaptation of the last encoder layer consistently improves performance, yielding stable accuracy above 99% in all cases and reducing variance. The gain from fine-tuning the last encoder layer ranges from approximately 0.8% to 1.1%, indicating that even limited feature-level adaptation is effective in compensating for domain shift. Figure 15 further examines this effect across individual source locations. Models are trained on the source location (x-axis of Figure 15) and evaluated on the rest of the unseen locations. For all nine source domains, fine-tuning only the classifier results in noticeable performance degradation and higher variance across unseen target domains. Allowing the last encoder layer to adapt improves mean accuracy by 1.3–2.2% and consistently reduces variance, demonstrating more reliable cross-domain generalization.

TABLE II: Computation Overhead of different fingerprinting strategies

Fine-tuning Type	Inference (ms)	Adaptation (s)	Peak Memory (MB)
DeepCSiv2	227	–	344.21
BeamID (No adapt)	133	–	26.59
BeamID (Finetune Classifier)	133	2.21	34.96
BeamID (Finetune Class.+Enc.)	133	2.32	35.15

5) *Inference Latency and Overhead of Adaptation:* Table II summarizes the computational overhead of BeamID in terms of inference latency, adaptation time, and memory usage. At inference time, BeamID requires 133 ms per sample, significantly lower than DeepCSiv2 (227 ms), and this latency remains unchanged regardless of whether adaptation is enabled. Moreover, the cost of few-shot adaptation is also modest. Fine-tuning only the classifier requires 2.21 s, while

additionally adapting the last encoder layer increases the adaptation time marginally to 2.32 s. Memory overhead of BeamID is also low—without adaptation, it uses 26.59 MB, and adaptation increases peak memory usage to 34.96 MB and 35.15 MB for classifier-only and classifier-plus-encoder fine-tuning, respectively, compared to 344.21 MB for DeepCSiv2. These results show that BeamID enables efficient few-shot adaptation with minimal time and memory overhead, making it suitable for real-world deployment.

VI. CONCLUSION

This paper presents BeamID, a domain-adaptive RFP framework that leverages standard-compliant MIMO beamforming feedback to enable reliable identification under environment-induced domain shifts. BeamID learns client-centric embeddings in a source environment using supervised contrastive learning and performs lightweight few-shot adaptation in embedding space to align target-domain fingerprints without full retraining. By incorporating few-shot adaptation, BeamID effectively mitigates the cross-domain degradation and enables reliable identification under environment changes. Using only a small number of labeled target-domain samples per client (as few as 5 samples), BeamID achieves up to 99% cross-domain identification accuracy. Moreover, BeamID incurs low deployment overhead, maintaining 133 ms inference latency with only ~ 2.3 s adaptation time, making it practical for real-world deployments. We note that our evaluation focuses on a fixed frequency band and a single AP platform. Extending BeamID to heterogeneous AP hardware and frequency bands is an important direction for future work.

VII. ACKNOWLEDGMENT

This work has been supported by the National Science Foundation under grants CCF-2218845, ECCS-2229472 and ECCS-2329013; by the Air Force Office of Scientific Research under grant FA9550-23-1-0261; and by the Office of Naval Research under grant N00014-23-1-2221.

REFERENCES

- [1] Y. Zhang, Y. Lin, Z. Dou, M. Wang, and W. Li, "Monitoring and identification of wifi devices for internet of things security," in *2019 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–5, IEEE, 2019.
- [2] I. Palamà, A. Amici, G. Bellicini, F. Gringoli, F. Pedretti, and G. Bianchi, "Attacks and vulnerabilities of wi-fi enterprise networks: User security awareness assessment through credential stealing attack experiments," *Computer Communications*, vol. 212, pp. 129–140, 2023.
- [3] F. Meneghello, M. Rossi, and F. Restuccia, "DeepCSI: Rethinking Wi-Fi radio fingerprinting through MU-MIMO CSI feedback deep learning," in *2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS)*, pp. 1062–1072, IEEE, 2022.
- [4] A. Al-Shawabka, F. Restuccia, S. D'Oro, T. Jian, B. C. Rendon, N. Soltani, J. Dy, S. Ioannidis, K. Chowdhury, and T. Melodia, "Exposing the Fingerprint: Dissecting the Impact of the Wireless Channel on Radio Fingerprinting," in *Proc. of IEEE INFOCOM*, 2020.
- [5] F. Meneghello, K. F. Haque, and F. Restuccia, "Radio fingerprinting of wi-fi devices through mimo compressed channel feedback," in *IEEE INFOCOM 2025-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1–6, IEEE, 2025.
- [6] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM computing surveys (csur)*, vol. 53, no. 3, pp. 1–34, 2020.
- [7] H.-J. Ye, H. Hu, D.-C. Zhan, and F. Sha, "Few-shot learning via embedding adaptation with set-to-set functions," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8808–8817, 2020.
- [8] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell, "Meta-learning with latent embedding optimization," *arXiv preprint arXiv:1807.05960*, 2018.
- [9] K. F. Haque, M. Zhang, F. Meneghello, and F. Restuccia, "Si-fi: Learning the beamforming feedback for simultaneous multi-subject sensing," *Computer Networks*, p. 111907, 2025.
- [10] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18661–18673, 2020.
- [11] F. Graf, C. Hofer, M. Niethammer, and R. Kwitt, "Dissecting supervised contrastive learning," in *International Conference on Machine Learning*, pp. 3821–3830, PMLR, 2021.
- [12] IEEE, "IEEE Standard for Information Technology–Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 4: Enhancements for Very High Throughput for Operation in Bands below 6 GHz," *IEEE Std 802.11ac-2013 (Amendment to IEEE Std 802.11-2012)*, 2013.
- [13] IEEE, "IEEE Standard for Information Technology–Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High-Efficiency WLAN," *IEEE Std 802.11ax-2021 (Amendment to IEEE Std 802.11-2020)*, 2021.
- [14] M. S. Gast, *802.11 ac: A Survival Guide: Wi-Fi at Gigabit and Beyond*. "O'Reilly Media, Inc.", 2013.
- [15] K. F. Haque, F. Meneghello, and F. Restuccia, "Wi-BFI: Extracting the IEEE 802.11 Beamforming Feedback Information from Commercial Wi-Fi Devices," in *Proc. of ACM WINTeCH*, 2023.
- [16] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [17] T. D. Vo-Huu, T. D. Vo-Huu, and G. Noubir, "Fingerprinting Wi-Fi Devices Using Software Defined Radios," in *Proc. of ACM WiSec*, 2016.
- [18] L. Peng, A. Hu, J. Zhang, Y. Jiang, J. Yu, and Y. Yan, "Design of a Hybrid RF Fingerprint Extraction and Device Classification Scheme," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 349–360, 2019.
- [19] T. Zheng, Z. Sun, and K. Ren, "FID: Function Modeling-based Data-Independent and Channel-Robust Physical-Layer Identification," in *Proc. of IEEE INFOCOM*, 2019.
- [20] F. Restuccia, S. D'Oro, A. Al-Shawabka, M. Belgiovine, L. Angioloni, S. Ioannidis, K. Chowdhury, and T. Melodia, "DeepRadioID: Real-Time Channel-Resilient Optimization of Deep Learning-based Radio Fingerprinting Algorithms," in *Proc. of ACM MobiHoc*, 2019.
- [21] K. Sankhe, M. Belgiovine, F. Zhou, S. Riyaz, S. Ioannidis, and K. Chowdhury, "ORACLE: Optimized Radio Classification through Convolutional neural networks," in *Proc. of IEEE INFOCOM*, 2019.
- [22] S. Riyaz, K. Sankhe, S. Ioannidis, and K. Chowdhury, "Deep Learning Convolutional Neural Networks for Radio Identification," *IEEE Communications Magazine*, vol. 56, no. 9, pp. 146–152, 2018.
- [23] K. Merchant, S. Revay, G. Stantchev, and B. Nossain, "Deep Learning for RF Device Fingerprinting in Cognitive Communication Networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 160–167, 2018.
- [24] R. Das, A. Gadre, S. Zhang, S. Kumar, and J. M. Moura, "A Deep Learning Approach to IoT Authentication," in *Proc. of IEEE ICC*, 2018.
- [25] S. D'Oro, F. Restuccia, and T. Melodia, "Can You Fix My Neural Network? Real-Time Adaptive Waveform Synthesis for Resilient Wireless Signal Classification," in *Proc. of IEEE INFOCOM*, 2021.
- [26] N. Basha, B. Hamdaoui, K. Sivanesan, and M. Guizani, "Channel-resilient deep-learning-driven device fingerprinting through multiple data streams," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 118–133, 2023.
- [27] B. Hamdaoui, N. Basha, and K. Sivanesan, "Deep learning-enabled zero-touch device identification: Mitigating the impact of channel variability through MIMO diversity," *IEEE Communications Magazine*, vol. 61, no. 6, pp. 80–85, 2023.