

BANDWEAVE: Enhanced Channel Estimation in MIMO Networks with Multi-Band Fusion

Khandaker Foysal Haque*, Francesca Meneghello*, Jonathan Ashdown† and Francesco Restuccia*

*Northeastern University, United States †SpectrumAI Solutions, United States

Abstract—High-throughput wireless communication relies on high-resolution channel estimates to accurately compensate for the distortions imposed by the underlying wireless multipath propagation. However, due to the frequency-dependent superposition of delayed propagation paths, certain components may undergo destructive interference at specific frequencies, thus obscuring their contribution. This leads to a *partial* reconstruction of the channel which ultimately degrades communication performance. As such, we propose BANDWEAVE, a new general-purpose multi-band channel frequency response (CFR) fusion approach that significantly improves channel estimation accuracy by leveraging complementary multipath information across multiple frequency bands. BANDWEAVE enables adaptive, robust and efficient CFR fusion through a three-phase progressive learning algorithm. We experimentally evaluated BANDWEAVE with an IEEE 802.11ac multi-user (MU)-multiple-input multiple-output (MIMO) testbed and a 60 GHz single-user (SU)-MIMO system in three propagation environments with different fading conditions. Experimental results show that BANDWEAVE channel estimation outperforms single-band estimation by achieving more than 16% of throughput improvement and up to 10.2% bit error rate (BER) reduction. BANDWEAVE achieves a gain of more than 7.5× in communication performance compared to state of the art (SOTA) band-merging approaches. Measurements on Jetson Nano, Jetson Orin Nano and Raspberry Pi 5 show that BANDWEAVE reduces inference time and energy consumption by up to 20× and 18×, respectively, compared to SOTA methods. A video demonstration of BANDWEAVE is available¹.

Index Terms—Multi-band fusion, channel estimation, MIMO.

I. INTRODUCTION

The rapid rise of augmented reality (AR), virtual reality (VR) [1], [2], and immersive media streaming [3], [4] has placed unprecedented demands on wireless networks [1]. These applications require ultra-low latency – below 5 ms – high reliability and high spectral efficiency to deliver immersive and seamless user experiences [5].

To address this growing connectivity demand, multiple-input multiple-output (MIMO) is becoming an essential component of next-generation wireless systems. Yet, a key limiting factor for MIMO is the insufficient resolution in channel frequency response (CFR) estimation, which directly affects beamforming precision, interference suppression, and multiplexing gain. For instance, resolving multipath components with sub-meter range accuracy requires a delay resolution in the order of 1.7–3.3 nanoseconds, corresponding to bandwidths of 300–600 MHz [6], [7]. While recent wireless systems – such as the IEEE 802.11be Wi-Fi amendment – can support up to 320 MHz bandwidth, most existing deployments often operate at narrower bands (e.g., 20/40/80 MHz), thus limiting the achievable resolution to 6.25–25 nanoseconds [8], [9]. Additionally, these resolutions are the upper bounds and can be barely achieved in real-world deployment due to

noise and hardware imperfections affecting channel estimation. While recent work [7], [10]–[12] has explored cross-band prediction and band splicing to enhance sensing resolution, these approaches do not focus on improving CFR estimation quality for communication. Consequently, communication performance remains constrained by low-resolution estimates.




Fig. 1: CFR and multipath decomposition across different Wi-Fi channels (155 and 165).

In this paper, we present BANDWEAVE, a novel multi-band CFR fusion framework that enhances channel estimation by coherently combining complementary multipath information across spectrum bands. Our key observation, illustrated in Figure 1, is that while the CFR (on the left plot) is band-specific as it expresses how the different multipath components combine in the frequency domain at the different modulating frequencies, the *multipath components* (on the right plot) are band-independent as they are only linked with the propagation environment. *By collecting CFRs estimated at different frequencies, we are looking at the same physical scene under slightly different viewpoints* (the different frequencies), so that paths that may be hidden in one spectrum band – due to destructive superpositions among them – can appear in another, and thus, *multipath components can be “weaved” together to achieve better channel estimation*. By achieving superior knowledge of the propagation channel, the beamformer obtains a more suitable precoding to improve the stream signal-to-noise-ratio (SNR) toward each beamformee while reducing inter-stream and inter-user interference.

Technical Challenges: Coherent fusion of CFRs across frequency bands is technically challenging due to cross-band inconsistencies. Frequency-dependent path loss introduces amplitude mismatches, while phase offsets caused by hardware non-ideality disrupt temporal coherence. Additionally, delay variations across center frequencies and hardware-induced calibration errors, such as synchronization issues, further degrade CFR alignment. These impairments hinder direct fusion, often resulting in inaccurate or unstable channel estimates. Effective CFR fusion requires a context-aware and adaptive framework that compensates for such distortions.

¹<https://youtu.be/COGgKKtj8SA>

Proposed Approach: BANDWEAVE learning-based multi-band CFR fusion (i) leverages the time-domain representation of the channel information representing the multipath structure for precise delay-domain alignment of the CFRs at different bands, (ii) applies SNR-guided delay and phase compensation to correct frequency-dependent attenuation and oscillator-induced phase errors, and (iii) uses a progressive learning pipeline, consisting of a supervised pretraining, a simulation-in-the-loop fine-tuning and an online feedback-aware adaptation, to handle delay variation across frequencies and calibration mismatches under dynamic conditions. As shown in Figure 2, BANDWEAVE operates in a closed loop fashion: beamformers fuse multi-band CFRs in the time domain and feed back enhanced channel estimates to the beamformer for improved MIMO precoding.




Fig. 2: Overview of BANDWEAVE. ‘Ch.’ stands for channel.

Summary of Novelty

- We propose BANDWEAVE, the first framework to coherently fuse CFRs across both contiguous and non-contiguous bands to enhance channel estimation resolution. Unlike prior sensing-oriented techniques [13]–[16], BANDWEAVE fundamentally improves communication performance by directly refining the quality of multi-band channel estimates;
- We introduce a novel three-phase progressive learning framework to address frequency-dependent path loss and phase misalignments. Unlike standalone signal-processing [10] and learning-based approaches [11], [13], our method integrates: (i) supervised pretraining, (ii) simulation-in-the-loop fine-tuning guided by link-level performance metrics (i.e., bit error rate (BER), throughput), and (iii) online feedback-aware adaptation to guarantee robustness under dynamic channel conditions. The core of BANDWEAVE is a transformer-based fusion that integrates per-band encoders, multi-head self-attention, and positional encoding to capture multipath structures;
- We extensively evaluate BANDWEAVE with two experimental platforms, i.e., a multi-user (MU)-MIMO testbed using commercial IEEE 802.11 Wi-Fi devices, and a 60 GHz single-user (SU)-MIMO system. BANDWEAVE achieves up to 16.0% improvement over the IEEE 802.11 standard, while delivering up to $7.5\times$ and $4.9\times$ gains in throughput and BER, respectively, compared to the state of the art (SOTA) approach HiSAC [10]. We further demonstrate its efficiency on edge platforms Jetson Nano, Jetson Orin Nano, and Raspberry Pi 5. BANDWEAVE achieves up to $20\times$ faster inference and $18\times$ lower energy consumption than SOTA baseline HiSAC [10]. For reproducibility, we open-

sourced our datasets and code at <https://github.com/Restuccia-Group/BandWeave>.

Summary of Impact

BANDWEAVE breaks new ground by providing a *new and unconventional methodology* that substantially improves MIMO networking and paves the way for additional research in multi-band technologies. Moreover, BANDWEAVE operates *independently* from the wireless technology and the spectrum band being used, which makes it applicable in any MIMO network.

II. RELATED WORK

Existing work can be broadly categorized into cross-band CFR estimation and spectrum aggregation for enhanced delay resolution. *Noticeably, prior work is sensing-oriented and focus on target separability. Instead, BANDWEAVE improves estimation accuracy through coherent multi-band fusion, directly improving communication performance.*

Cross-Band CFR Estimation: CLCP [12] exploits multi-view representations from nearby sensors to predict unobserved CFR using a product-of-experts model. OptML [17], FIRE [18], and Deep UL2DL [19] propose efficient CFR prediction in frequency division duplexing systems by avoiding explicit feedback. These methods employ neural networks and generative models to infer downlink CFR from uplink observations. However, they assume band reciprocity or target-specific frequency translations, and do not generalize to multi-band CFR alignment or fusion across contiguous or non-contiguous spectra. Finally, HORCRUX [11] adopts a hierarchical structure that decomposes the uplink channel into distance-based sub-channels, allowing lightweight models to estimate downlink parameters. While promising for cross-band prediction, this work does not implement multiband fusion.

Spectrum Splicing: Efforts to enhance delay resolution through spectrum aggregation include UWB-Fi [13], which asynchronously samples non-contiguous Wi-Fi channels over 4.7 GHz to achieve ultra-wideband sensing. While it uses model-driven deep learning to mitigate channel sampling irregularities, it remains tailored to localization and radio-based respiration monitoring. Similarly, Tian et al. [15] propose a two-stage framework that combines sparse non-uniform discrete Fourier transform inversion with subspace projection to improve range resolution. Li et al. [14] increase sub-channel density within a single band to refine delay estimation, while mmSplicer [16] reconstructs high-resolution channel impulse responses (CIRs) from partial mmWave band measurements. Although both improve ranging accuracy, they are not designed to handle frequency-dependent distortions that undermine cross-band fusion fidelity. Wang et al. [20] improve localization by fusing CFRs at 2.4 GHz, 5.1 GHz, and 5.8 GHz bands using hill-climbing to resolve delay ambiguity, but they overlook amplitude and phase mismatches across bands. Zhuo et al. [21] further highlight that hardware-induced linear and nonlinear phase errors persist across bands, emphasizing the

need for robust cross-band calibration. HiSAC [10] addresses the resolution bottleneck in CFR estimation by leveraging non-contiguous multiband communication signals. It combines geometric priors and frequency diversity to improve temporal resolution but focuses on sensing applications.

III. MULTI-BAND CFR FUSION WITH BANDWEAVE

Multipath propagation is a spatial-temporal phenomenon that persists across frequencies, as multiple bands “observe” the same environment. Hence, *why should we limit channel estimation to a single band when diverse bands offer complementary views of the same propagation landscape?*

Key Observation: As illustrated in Figure 3, the CIRs extracted from different bands with distinct center frequencies and bandwidths (i.e., 80 MHz vs. 160 MHz) exhibit consistent dominant multipath components. These components correspond to dominant scatterers and reflectors that are simultaneously visible across bands, even if attenuated or shifted slightly due to frequency-dependent propagation effects. Notably, channels with similar bandwidth but different center frequencies (e.g., Wi-Fi channel 36 and 163) reveal temporally aligned dominant paths, while wider-band channels offer finer resolution of the same multipath clusters. By leveraging this shared multipath structure, we can fuse CFRs across bands to construct a richer, more accurate, and more resilient channel representation. Based on this insight, we design the signal processing-driven multi-band fusion strategy presented in Section III-B. Next, we introduce the complete BANDWEAVE approach in Section III-C.




Fig. 3: Multipath analysis through CIR across different center frequencies and bandwidths.

A. System Model

We consider a multi-carrier MU-MIMO system using K Orthogonal Frequency Division Multiplexing (OFDM) sub-channels and consisting of one beamformer with M antennas and N_U beamformees with N antennas each. In MU-MIMO, the beamformer uses a precoding matrix to linearly combine the signals to be simultaneously transmitted to the different beamformees. The precoding matrix is derived from the CFR matrices, estimated by each of the beamformees, which describe how the environment modifies the irradiated signals in their path to the receivers. The estimation process is called *channel sounding* and is triggered by the beamformer by periodically broadcasting the pilot frames. Each pilot frame

Algorithm 1: Signal Processing-based Multi-Band CFR Fusion

```

1: Input:  $\{\text{CFR}_i(f)\}_{i=1}^{N_B}$  (CFRs from  $N_B$  bands)
2: Output: Fused CFR,  $\text{CFR}_{\text{fused}}(f)$ 
3: Step 1: CIR Extraction
4: for  $i = 1$  to  $N_B$  do
5:    $\text{CIR}_i(t) \leftarrow \text{IFFT}(\text{CFR}_i(f))$ 
6: Step 2: CIR Normalization
7: for  $i = 1$  to  $n$  do
8:    $\text{CIR}_i^{\text{norm}}(t) \leftarrow \frac{\text{CIR}_i(t)}{\|\text{CIR}_i(t)\|}$ 
9: Step 3: Delay and Phase Alignment
10: Select reference CIR:  $\text{CIR}_{\text{ref}}^{\text{norm}}(t) \leftarrow \text{CIR}_1^{\text{norm}}(t)$ 
11: for  $i = 2$  to  $N_B$  do
12:   Estimate delay offset using Equation 1
13:   Refine  $\Delta\tau_i$  using parabolic interpolation
14:   Delay-align using  $\text{CIR}_i^{\text{aligned}}(t) \leftarrow \text{CIR}_i^{\text{norm}}(t + \Delta\tau_i)$ 
15:   Estimate phase offset as
16:    $\phi_i(t) \leftarrow \angle \text{CIR}_{\text{ref}}^{\text{norm}}(t) - \angle \text{CIR}_i^{\text{aligned}}(t)$ 
17:   Apply phase correction with Equation 2
18: Step 4: SNR-Based Component-wise Fusion
19: for all components  $t$  do
20:   for  $i = 1$  to  $N_B$  do
21:      $P_i(t) \leftarrow |\text{CIR}_i^{\text{aligned}}(t)|^2$ 
22:   Normalize the weights as  $w_i(t) \leftarrow \frac{P_i(t)}{\sum_{i=1}^{N_B} P_i(t)} \quad \forall i$ 
23:   Obtain the fused CIR using Equation 3
24: Step 5: CFR Reconstruction
25:  $\text{CFR}_{\text{fused}}(f) \leftarrow \text{FFT}(\text{CIR}_{\text{fused}}(t))$ 
26: Return:  $\text{CFR}_{\text{fused}}(f)$ 

```

contains sequences of bits the decoded version of which is known to the beamformees. Upon receiving the pilot frame, each beamformee estimates the $K \times M \times N$ CFR matrix $\mathbf{H}(f)$ describing the channel impairments. Next, the CFR is compressed to reduce the channel overhead and fed back to the beamformer for precoding.

B. Signal Processing-based Multi-Band CIR Fusion

The signal processing-based multi-band fusion approach is summarized in Algorithm 1. Given CFR measurements from N_B frequency bands, denoted as $\{\text{CFR}_i(f)\}_{i=1}^{N_B}$, the process begins with transforming each CFR into its corresponding CIR representation through IFFT (step 1, lines 3-5 of Algorithm 1) [22], [23]. The resulting CIRs, $\text{CIR}_i(t)$, capture fine-grained multipath structures that are critical for effective CFR fusion. To eliminate amplitude scale differences, thus ensure consistency across bands, each CIR is normalized by its magnitude $\|\text{CIR}_i(t)\|$ (step 2, lines 6-8 of Algorithm 1). Delay and phase alignment are then performed to have CIRs from different bands with a common reference frame (step 3, lines 9-16 of Algorithm 1). Without loss of generality, we consider the CIR from the first band as the reference, denoted as $\text{CIR}_{\text{ref}}^{\text{norm}}(t)$ (line 10 of Algorithm 1). For each of the remaining CIRs associated with the other center frequencies, the relative delay $\Delta\tau_i$ is estimated by maximizing the correlation between this CIRs and the reference as

$$\Delta\tau_i = \arg \max_k \sum_t |\text{CIR}_{\text{ref}}^{\text{norm}}(t) \cdot \text{CIR}_i^{\text{norm}}(t - k)|, \quad (1)$$

and parabolic interpolation is applied to refine $\Delta\tau_i$ with fractional delay resolution [24] (lines 12-13 of Algorithm 1). After aligning the delays (line 14 of Algorithm 1), the CIRs are further phase-synchronized. The phase difference at each component is obtained as $\phi_i(t) = \angle \text{CIR}_{\text{ref}}^{\text{norm}}(t) - \angle \text{CIR}_i^{\text{aligned}}(t)$, and used to apply complex exponential compensation (lines 15-16 of Algorithm 1):

$$\text{CIR}_i^{\text{aligned}}(t) \leftarrow \text{CIR}_i^{\text{aligned}}(t) \cdot e^{j\phi_i(t)}. \quad (2)$$

Hence, the CIRs are fused at the component level based on their instantaneous signal power. For each component t , we compute the power of the aligned CIR as $P_i(t) = |\text{CIR}_i^{\text{aligned}}(t)|^2$, and derive the corresponding normalized fusion weight as $w_i(t) = P_i(t) / \sum_{i=1}^{N_B} P_i(t)$ (step 4, lines 17-21 of Algorithm 1). The final fused CIR is then synthesized as a weighted sum (line 22):

$$\text{CIR}_{\text{fused}}(t) = \sum_{i=1}^{N_B} w_i(t) \cdot \text{CIR}_i^{\text{aligned}}(t). \quad (3)$$

Finally, the fused CIR is transformed back to the frequency domain using a conventional FFT operation (step 5, lines 23-25 of Algorithm 1).

We primarily evaluate Algorithm 1 in a raytracing-based simulation environment considering a network of one access point (AP) and one station (STA) operating on Wi-Fi channels 36 and 163 (160 MHz of bandwidth) using BER as a key performance metric. The top plot in Figure 4 compares the BER obtained when the AP uses the original and the fused CFRs (with Algorithm 1) to transmit data on Wi-Fi channel 36 across 300 communication rounds. Additionally, we assessed BANDWEAVE performance on a mmWave SU-MIMO experimental testbed deployed in an office environment (see Section IV-B), where we fuse two 1 GHz bands centered at 57 GHz and 58 GHz and use the fused CFR to transmit data on the 57 GHz band. The results are reported in the bottom subplot of Figure 4. While the simulation yields a consistent BER improvement of 5.39%, the real-world experiment demonstrates a reduced gain of 2.57% alongside an increased occurrence of rounds (points marked as red crosses) where fusion leads to performance degradation. This increased variance and degraded reliability in real scenarios underscore the sensitivity of the signal processing-based fusion method to hardware non-idealities, temporal drift, and frequency-dependent distortions. To address these limitations, we design the learning-based BANDWEAVE framework described next.




Fig. 4: BER performance evaluation with Algorithm 1.

C. BANDWEAVE Walkthrough

To address the limitations of standalone signal-processing fusion under practical impairments, we develop BANDWEAVE, a learning-driven framework that extends Algorithm 1 to adaptively align and fuse CFRs across both contiguous and non-contiguous frequency bands, irrespective of their bandwidth.

Figure 5 illustrates the BANDWEAVE closed-loop workflow from channel estimation through sounding to data precoding for beamforming, where each step corresponds to a stage in the system pipeline. The process begins with multi-band channel sounding by the beamformer using pilot frames across all N_B available bands (step 1 of Figure 5). Each beamformee then estimates the CFR, $\mathbf{H}(f)$, for each band and extracts the corresponding delay-domain CIR representation, $\mathbf{H}(t)$ (step 2 of Figure 5). The CIRs, $\mathbf{H}(t)$ from each band is passed through the three-phase BANDWEAVE learning algorithm and the fused CIR, $\mathbf{H}_{\text{fused}}(t)$ is obtained (step 3 of Figure 5, detailed in Section III-D). The fused CIR, $\mathbf{H}_{\text{fused}}(t)$, is transformed back to the frequency domain as $\mathbf{H}_{\text{fused}}(f)$ using a standard FFT operation (step 4 of Figure 5). BANDWEAVE obtains the beamforming feedback from fused CFR, $\mathbf{H}_{\text{fused}}(f)$ following IEEE 802.11 standard [8], [9], [25], which is then fed back to the beamformer (step 5 of Figure 5). The beamformer then derives the steering matrix from the beamforming feedback and applies it for MU-MIMO precoding before data transmission (step 6 of Figure 5), completing the BANDWEAVE communication round.

D. BANDWEAVE Learning Algorithm

The learning process begins with supervised pre-training using labeled simulation-based data fused with Algorithm 1, allowing the model to learn reliable fusion priors (Phase 1 in Figure 5). This is followed by simulation-in-the-loop fine-tuning, which adapts the model to realistic wireless impairments in a controlled environment (Phase 2 in Figure 5). Once deployed, the model enters an online feedback-aware adaptation phase, where it adjusts the trained weights based on real-time feedback from the beamformees to maintain optimal communication performance under dynamic conditions (Phase 3 in Figure 5). Together, these three phases define an end-to-end learning framework that progressively enhances fusion generalization and resilience.

Phase 1: Supervised Pretraining for N_B -Band CIR Fusion

The supervised pretraining phase of BANDWEAVE is designed to learn a generalized fusion policy that synthesizes high-resolution CFR representations from multiple frequency bands. This pretraining serves as the foundational initialization for downstream adaptive learning phases. Let $\{\mathbf{C}^{(i)}\}_{i=1}^{N_B}$ denote the CIRs from N_B different bands, where each $\mathbf{C}^{(i)} \in \mathbb{R}^{T \times M \times N \times 2}$ represents the delay-domain representation containing T components across M transmit and N receive antennas, with 2 channels for real and imaginary parts. Each CIR is first reshaped into a sequence of antenna tokens of size $A = M \cdot N$, resulting in $\mathbf{C}^{(i)} \in \mathbb{R}^{T \times A \times 2}$. Each CIR $\mathbf{C}^{(i)}$ is projected into a higher-dimensional embedding space using a shared linear encoder. The embedding for the i -th band is




Fig. 5: BANDWEAVE Walkthrough. The numbers in the figure indicate the steps of the processing.

computed as: $\mathbf{E}^{(i)} = \text{ReLU}(\mathbf{C}^{(i)} \cdot \mathbf{W}_{\text{enc}} + \mathbf{b}_{\text{enc}})$, for $i = 1, 2, \dots, N_B$, where $\mathbf{W}_{\text{enc}} \in \mathbb{R}^{2 \times d_m}$ is a learnable weight matrix that projects the real and imaginary components of each token into a d_m -dimensional embedding space, and $\mathbf{b}_{\text{enc}} \in \mathbb{R}^{d_m}$ is the corresponding bias vector. To retain temporal and spatial structure, two types of positional information are added to each embedding. A sinusoidal positional encoding matrix $\mathbf{P} \in \mathbb{R}^{T \times 1 \times d_m}$ encodes delay ordering, while a learnable antenna embedding $\mathbf{E}_{\text{ant}} \in \mathbb{R}^{1 \times A \times d_m}$ provides spatial identity across antenna pairs. These are added to each token in $\mathbf{E}^{(i)}$, with broadcasting over the component dimension, $\mathbf{E}^{(i)} \leftarrow \mathbf{E}^{(i)} + \mathbf{P} + \mathbf{E}_{\text{ant}}$. The embedded representations from all bands are then concatenated to form a unified token sequence:

$$\mathbf{Z}_0 = [\mathbf{E}^{(1)}; \mathbf{E}^{(2)}; \dots; \mathbf{E}^{(N_B)}] \in \mathbb{R}^{(N_B \cdot T \cdot A) \times d_m}. \quad (4)$$

This sequence is processed by a transformer encoder [26] consisting of two layers of multi-head self-attention and feedforward blocks to learn joint cross-band representations, as illustrated in Figure 6. To obtain a unified output, we introduce $T \cdot A$ learnable fusion tokens that attend over all embedded tokens from N_B bands via a cross-attention module. The resulting fused representation $\mathbf{Z}_{\text{fused}} \in \mathbb{R}^{T \cdot A \times d_m}$ is then decoded by a two-layer multilayer perceptron (MLP) to predict the fused CIR as:

$$\hat{\mathbf{C}}_{\text{flat}} = \text{MLP}(\mathbf{Z}_{\text{fused}}) \in \mathbb{R}^{T \cdot A \times 2}, \quad (5)$$

$$\hat{\mathbf{C}}_{\text{fused}} = \text{reshape}(\hat{\mathbf{C}}_{\text{flat}}) \in \mathbb{R}^{T \times M \times N \times 2}. \quad (6)$$

Training is first performed using simulated data, adopting a supervised loss computed against a pseudo-ground-truth fused CIR $\mathbf{C}^*$ obtained through Algorithm 1. Specifically, the loss ($\mathcal{L}_{\text{MSE}}$) is defined as the mean squared error between the predicted and target CIRs:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{T \cdot M \cdot N} \sum_{t,x,y} \left\| \hat{\mathbf{C}}_{\text{fused}}[t, x, y, :] - \mathbf{C}^*[t, x, y, :] \right\|_2^2. \quad (7)$$

This pretraining step enables the model to learn frequency-invariant features of multipath propagation and provides a strong initialization for fine-tuning in more complex, real-world scenarios.

Phase 2: Fine-tuning through Simulation in the Loop

While the supervised pretraining phase provides a strong initialization, it remains agnostic to downstream system-level




Fig. 6: Transformer-based architecture for supervised pretraining with N_B bands.

metrics such as throughput or BER. To address this, we introduce a simulation-in-the-loop fine-tuning phase that bridges data-driven learning and physical-layer dynamics by directly optimizing the fusion model using feedback from an end-to-end communication simulation. Specifically, we fine-tune the pretrained fusion model $f_{\theta}(\cdot)$, parameterized by weights θ , using BER feedback obtained from full-stack communication simulations driven by a ray-tracer, in a closed-loop fashion. The procedure begins by feeding CIRs from N_B different bands into the pretrained model from Phase 1, which outputs a fused CIR, denoted as $\hat{\mathbf{C}}_{\text{fused}}$. This fused CIR is then passed through the simulator implementing the full-stack communication standard to evaluate BER under a configurable link setup, which includes modulation and coding scheme (MCS), antenna configurations, carrier frequency, and SNR. Our goal is to fine-tune the model parameters θ such that the BER is minimized. For this, we use a simulation-driven loss function defined as $\mathcal{L}_{\text{sim}}(\theta) = \left\| \mathcal{B}(\hat{\mathbf{C}}_{\text{fused}}) \right\|_2^2$, where $\mathcal{B}(\cdot)$ represents the BER obtained using the fused CIR to design the precoding in the end-to-end communication simulator.

Overall, the simulation-in-the-loop framework consists of two components: the learning model (from Phase 1) and the full-stack communication simulation module. Critically, the latter is non-differentiable due to discrete physical-layer operations such as modulation, quantization, and decoding, along with ray tracing-based channel modeling, which together prevent gradient flow back to the learning model [27], [28]. As a result, standard gradient-based optimizers such as SGD or Adam [29] are inapplicable for minimizing the simulation-

driven loss $\mathcal{L}_{\text{sim}}(\theta)$. To address this, we adopt the covariance matrix adaptation evolution strategy (CMA-ES) [30], a robust derivative-free optimization algorithm well-suited for high-dimensional, non-convex, and noisy search spaces, which enables efficient fine-tuning guided by simulation outcomes. CMA-ES maintains a Gaussian distribution over model parameters and iteratively updates its mean and covariance based on BER feedback from sampled candidates. This process continues until convergence or early stopping.

Phase 3: Online Feedback Aware Adaptation:

To bridge the gap between simulation-guided optimization and real-world deployment, BANDWEAVE integrates a lightweight online adaptation phase, detailed in Algorithm 2, which refines the pretrained fusion model (from Phase 2) using real-time feedback from its deployment. In contrast to Phase 2, which relies on a closed-loop simulator, this phase leverages actual performance measurements—such as throughput or BER—collected during live operation. The adaptation is entirely self-supervised and does not require ground-truth CIR labels.

Let $\left(\left\{\mathbf{C}_i^{(k)}\right\}_{k=1}^{N_B}, T_{\text{base}}^{(i)}\right)_{i=1}^D$ denote the deployment-time dataset comprising D samples, where $\mathbf{C}_i^{(k)}$ is the CIR from band k for sample i , and $T_{\text{base}}^{(i)}$ is the throughput achieved with standard MIMO precoding (single-band). Given the fine-tuned parameters from Phase 2, θ_0 , our goal is to adapt the model f_θ such that the fused CIR $\hat{\mathbf{C}}_i = f_\theta\left(\left\{\mathbf{C}_i^{(k)}\right\}_{k=1}^{N_B}\right)$ yields better throughput than the baseline. We decompose the adaptation into two performance-conditioned strategies, namely gradient-free exploration and teacher-guided fine-tuning via surrogate loss, as detailed next.

Gradient-Free Exploration (lines 5-7 of Algorithm 2): If $T_{\text{model}}^{(i)} \geq T_{\text{base}}^{(i)} + \delta$, where $T_{\text{model}}^{(i)}$ is the measured throughput and $\delta > 0$ is a user-defined margin, the learning algorithm is considered to outperform the baseline. In this case, a perturbed version of the current model parameters is generated as $\tilde{\theta} = \theta + \epsilon$, where ϵ is sampled from a zero-mean Gaussian distribution, $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. The perturbed model $f_{\tilde{\theta}}$ is then evaluated to assess whether the modification leads to improved performance. If $T_{\text{perturbed}}^{(i)} > T_{\text{model}}^{(i)}$, the perturbation is accepted, otherwise, it is rejected. The perturbation procedure is iteratively repeated until either a maximum number of trials K_{max} is reached or a predefined patience threshold P (i.e., consecutive failed updates) is exceeded. This gradient-free, reward-driven updates allow BANDWEAVE to refine the fusion policy directly from communication performance feedback without executing backpropagation, which is more computationally demanding.

Teacher-Guided Fine-Tuning via Surrogate Loss (lines 9-11 of Algorithm 2): If $T_{\text{model}}^{(i)} < T_{\text{base}}^{(i)}$, BANDWEAVE initiates a corrective fine-tuning phase to address the observed performance degradation. In this phase, the fusion model f_θ , parameterized by θ , is updated using a surrogate target CIR $\hat{\mathbf{C}}_{\text{teacher}}$ generated via Algorithm 1. Given the set of

Algorithm 2: Online Feedback-Aware Adaptation

Require: Model f_θ , deployment data $(\{\mathbf{C}_i^{(k)}\}_{k=1}^{N_B}, T_{\text{base}}^{(i)})_{i=1}^D$

- 1: Set threshold δ , noise scale σ
- 2: **for** each round $i = 1$ to D **do**
- 3: Predict fused CIR: $\hat{\mathbf{C}}_i = f_\theta(\{\mathbf{C}_i^{(k)}\})$
- 4: Measure throughput $T_{\text{model}}^{(i)}$
- 5: **if** $T_{\text{model}}^{(i)} \geq T_{\text{base}}^{(i)} + \delta$ **then**
- 6: $\tilde{\theta} = \theta + \mathcal{N}(0, \sigma^2 \mathbf{I})$
- 7: **if** $T_{\text{perturbed}}^{(i)} > T_{\text{model}}^{(i)}$ **then**
- 8: $\theta \leftarrow \tilde{\theta}$
- 9: **else**
- 10: Generate teacher $\hat{\mathbf{C}}_{\text{teacher}}$ via Algorithm 1
- 11: Minimize $\mathcal{L}_{\text{MSE}}(\theta)$ defined in Equation 8

input multi-band CIRs $\{\mathbf{C}_i^{(k)}\}_{k=1}^{N_B}$, the model is fine-tuned to minimize the mean squared error between its predicted output and the surrogate label using the loss

$$\mathcal{L}_{\text{MSE}}(\theta) = \left\| f_\theta\left(\left\{\mathbf{C}_i^{(k)}\right\}_{k=1}^{N_B}\right) - \hat{\mathbf{C}}_{\text{teacher}} \right\|_2^2. \quad (8)$$

To prevent overfitting to the surrogate CIR, BANDWEAVE performs teacher-guided fine-tuning for a limited number of iterations, set to $N_{\text{steps}} = 5$. This low value ensures that the model benefits from corrective supervision without drifting too far from its self-supervised trajectory. After completing these updates, the model's performance is re-evaluated using the same feedback-driven condition defined in the gradient-free exploration phase. If the updated model satisfies the improvement criterion, i.e., $T_{\text{model}}^{(i)} \geq T_{\text{base}}^{(i)} + \delta$, the system transitions back to the gradient-free exploration phase. This structured alternation between surrogate-guided correction and reward-driven adaptation allows BANDWEAVE to remain responsive under real-world conditions.

IV. EXPERIMENTAL SETUP

We evaluate the performance of BANDWEAVE using a MU-MIMO IEEE 802.11ac network operating in the 5 GHz unlicensed band, and a mmWave (60 GHz) SU-MIMO network. The testbeds are deployed in a conference room, a laboratory space, and an anechoic chamber. We consider line-of-sight (LoS), and non-line-of-sight (NLoS) scenarios, and include experiments in dynamic environments (human movement) as depicted in Figure 7.

A. IEEE 802.11ac MU-MIMO Setup

We use commercial off the shelf (COTS) Asus RT-AC86U 802.11ac Wi-Fi routers as both AP and STAs as presented in Figure 7. Four antennas are enabled at the AP while the two STAs have one antenna each. To implement BANDWEAVE, we had to reverse-engineer the channel sounding procedure on the Asus RT-AC86U devices and modify its functionality, following the approach we adopted in [31], [32]. We implemented the necessary firmware modifications using Nexmon [33] (see [31]) at one of the STAs, while the AP and the other STA are kept unmodified. Following the IEEE 802.11 standard [8], [9], and since commercial hardware does not expose raw




Fig. 7: Experimental scenarios in the anechoic chamber, classroom, and conference room, using both the IEEE 802.11ac MU-MIMO and the mmWave SU-MIMO testbeds.

CFR, we reconstruct the $\mathbf{V}$ matrices from the beamforming feedback extracted from the Broadcom chipset at the modified STAs. The $\mathbf{V}$ matrices from multiple bands serve as input to our learning algorithm. *iperf* sessions between the AP and each of the connected STAs are established to evaluate the BANDWEAVE effectiveness. UDP packets (1500 bytes-long) are transmitted to saturate channel capacity while we collect throughput directly from the *iperf* sessions.

B. mmWave SU-MIMO Setup

We use a fully digital 8×8 mmWave MIMO system capable of fine-grained channel measurements. It comprises two Zynq UltraScale+ RFSoc software-defined radios (SDRs), each integrated with Pi-Radio mmWave front-ends, acting respectively as the beamformer (AP) and beamformee (STA). Each SDR supports up to 8 transmit and 8 receive chains, enabling full 8×8 digital MIMO operation with OFDM-based waveform transmission. The system operates in the 57–64 GHz band, and supports up to 1 GHz bandwidth per band, with tunable center frequencies and fully synchronized clocks for deterministic channel sounding. We operate in SU-MIMO across multiple bands. An illustration of the experimental setup is shown in Figure 7. In our experiments, we collect CFRs from multiple 1 GHz bands (e.g., centered at 57 GHz and 58 GHz), using repeated sounding frames and precise timing control to ensure stable multipath measurements. The extracted CFR matrices serve as input to the BANDWEAVE fusion pipeline, which produces the fused CFR used for precoding and performance evaluation, with BER used as the feedback metric in Phase 3 of the learning algorithm (Algorithm 2).

V. PERFORMANCE EVALUATION

To establish rigorous signal-processing baselines, we adapt HiSAC [10] and UWB-Fi [13] to our multi-band fusion

task for communication. HiSAC [10] is a multi-band sensing strategy originally designed for high-resolution passive ranging. For our comparison, we use its phase offset compensation and OMP-based fusion mechanism to reconstruct CIRs from non-contiguous subbands, enabling a direct and fair comparison with BANDWEAVE in terms of fusion fidelity and temporal resolution. UWB-Fi [13] is a learning-based framework designed for multiband CFR aggregation in localization tasks. In our evaluation, we repurpose its FilterNet and Super-Resolution Network (SRNet) modules for CIR-level reconstruction, shifting the focus from spatial localization to temporal fusion performance.

Note that we consider the throughput as a metric in the IEEE 802.11ac MU-MIMO testbed as it is built on commercial off-the-shelf Wi-Fi routers that allow throughput logging via *iperf*. In contrast, BER measurements are collected for the mmWave SU-MIMO setup as it is implemented with SDRs that enable access to lower layer metrics including BER. We set the throughput margin $\delta = 3$ Mbps and the early stopping patience $P = 3$ iterations in all online adaptation experiments.






Fig. 8: Communications performance using BANDWEAVE, HiSAC, and UWB-Fi under varying environments and channel conditions. The fused CFR is used for transmission in each of the considered bands and the average performance is reported.

A. Performance Comparison with Contiguous Bands

We evaluate the communication performance in LoS, NLoS conditions, including dynamic (human movement) scenarios in two distinct indoor environments—an anechoic chamber and a classroom. As shown in Figure 8a, BANDWEAVE consistently achieves the highest throughput improvements over the standard with a gain up to 6.90%, and 14.27% in the anechoic chamber and classroom, respectively. These improvements stem from BANDWEAVE’s ability to effectively fuse contiguous bands while compensating for channel impairments such as multipath fading and mobility. HiSAC, which relies on super-resolution techniques without fine-tuned adaptation for communication performance, yields marginal throughput gains ($\leq 3.11\%$) and lacks robustness under dynamic conditions.



(a) Throughput gain with two 80 MHz bands: contiguous (Wi-Fi channels 42 and 58) and non-contiguous (channels 42 and 155).



(b) BER reduction with contiguous (centered at 57 & 58 GHz) and non-contiguous bands (centered at 57 & 59 GHz).

Fig. 9: Performance of contiguous and non-contiguous band fusion in dynamic environment (with human movement).

UWB-Fi, despite leveraging a wider virtual bandwidth via spectral stitching, suffers from incoherent band alignment and uncorrected phase distortions, resulting in degraded throughput – down up to -4.40% . BER results for the mmWave testbed follow similar trends as shown in Figure 8b: BANDWEAVE achieves the lowest error rates, while HiSAC offers marginal gains and UWB-Fi leads to performance degradation. Overall, BANDWEAVE outperforms HiSAC by up to $7.5\times$ in throughput gain and up to $4.9\times$ in BER reduction, while consistently surpassing UWB-Fi, which shows degrading performance compared to the standard across all scenarios.

B. Performance Comparison with Non-Contiguous Bands

To evaluate the robustness of BANDWEAVE, we extend our analysis to non-contiguous band fusion as shown in Figure 9. The results demonstrated that BANDWEAVE's superiority persists even in non-contiguous scenarios. In the classroom and conference room environments, BANDWEAVE achieves up to 13.11% throughput gain and 7.1% BER reduction with non-contiguous band fusion, showing only modest drops from the contiguous case (15.15% and 10.2% , respectively). In contrast, HiSAC and UWB-Fi, which offer moderate to null gains in the contiguous setting, exhibit significant degradation in non-contiguous conditions with UWB-Fi dropping throughput by up to 6.92% and worsening BER by 10.2% . In particular, we believe that HiSAC exhibits limited or negative gains, as it targets offline band merging for sensing, whereas our scenario requires online fusion, which is prevented by the time required for Orthogonal Matching Pursuit (OMP) execution.

C. Generalization to Heterogeneous Bandwidth Fusion

We evaluate BANDWEAVE by fusing contiguous and non-contiguous bands with heterogeneous bandwidths (40 MHz and 80 MHz) as presented in Figure 10. In both contiguous and non-contiguous settings, BANDWEAVE consistently outperforms HiSAC and UWB-Fi. Specifically, for contiguous bands, BANDWEAVE achieves more than 16.70% gain when the fused band is used with the 40 MHz band and 13.40%






Fig. 10: Throughput gain fusing bands with heterogeneous bandwidths (40 MHz and 80 MHz). Wi-Fi channels 42 and 58 form the contiguous band pair, while channels 42 and 155 are used for the non-contiguous bands.



(a) Throughput gain comparison between 2-band (Wi-Fi channel 42 and 58) and 3-band (Wi-Fi channel 42, 58, and 155) fusion.



(b) BER gain comparison between 2-band (centered at 57 and 58 GHz) and 3-band (centered at 57, 58 and 59 GHz) fusion.

Fig. 11: BANDWEAVE's scalability from 2-band to 3-band fusion in dynamic environments (with human movement).

with the 80 MHz band. Similar trends are observed in the non-contiguous case as well, indicating that including a wider band in the fusion yields greater performance gain when the fused CFR is applied to a narrower band, as this compensates for the limited sub-channel granularity of the latter.

D. Scalability of BANDWEAVE with Multi-Band Fusion

We investigate how BANDWEAVE behaves when increasing the number of bands considered in the fusion from two to three, evaluating throughput and BER across classroom and conference room environments. As shown in Figure 11, while BANDWEAVE continues to outperform HiSAC and UWB-Fi in both settings, the performance gain from fusing three bands over two is relatively small. For instance, BER gain improves from 10.2% to 12.1% in the classroom, and from 9.2% to 11.3% in the conference room when we switch to three-band fusion from two. However, the computational cost increases significantly due to the transformer-based architecture's (Figure 6 in Section III-D) nonlinear scaling with the number of bands. Hence, the use of more bands in the fusion algorithm should be decided based on a trade-off between performance improvement and computational complexity.

E. Ablation Study: Impact of the Learning Algorithm and Online Adaptation

We begin our ablation study by quantifying the effect of BANDWEAVE's online adaptation strategy. The test is




Fig. 12: Impact of BANDWEAVE online adaptation on throughput and BER across communication rounds in conference room with human movement.




Fig. 13: Impact of learning algorithm in BANDWEAVE on throughput and BER.

performed in the conference room with human movement to introduce dynamic channel variation across 90 communication rounds. We compare three configurations: (i) standard MIMO operation, (ii) BANDWEAVE without online adaptation, and (iii) BANDWEAVE with online adaptation. As shown in Figure 12a, BANDWEAVE without online adaptation improves throughput by 7.2% on average over the standard approach, whereas enabling online adaptation boosts the throughput by an additional 7.3%, resulting in a total improvement of 15% over the baseline. A similar trend is evident in the BER reduction (Figure 12b). The results highlight the key role of the online adaptation (see Algorithm 2) in the BANDWEAVE learning algorithm, as it enables the model to adapt to the deployment conditions.

We further evaluate BANDWEAVE without any learning algorithm, relying solely on signal processing-based fusion (Algorithm 1). The results are presented in Figure 13. Signal processing-based fusion achieves improvements in the range of 2.3%–3.8% and 1.6%–2.6% for throughput and BER respectively, where it performs worst in scenarios with human movements. Indeed, while signal processing-based fusion offers baseline improvements, it lacks adaptability in dynamic or complex environments. In contrast, adding BANDWEAVE’s three-phase learning algorithm leads to nearly 4× improvements in both throughput and BER, underscoring its critical role in ensuring robust performance in real-world deployment.




Fig. 14: Inference time and energy consumption of the fusion algorithms implemented on Jetson Nano, Jetson Orin Nano, and Raspberry Pi 5.

F. Analysis on Inference Time and Energy Consumption

We assess the computational complexity of BANDWEAVE, HiSAC, and UWB-Fi in terms of inference time and energy usage by implementing them on a Jetson Nano, a Jetson Orin Nano, and a Raspberry Pi 5 (RPi 5). As shown in Figure 14a, UWB-Fi achieves the lowest inference latency across all devices. BANDWEAVE follows closely with 9.4 ms on Jetson Orin Nano and 26.3 ms on RPi 5, while HiSAC takes 20× more time on RPi 5. BANDWEAVE consumes 57–210 mJ across platforms for each inference, significantly outperforming HiSAC, which requires up to 18× more energy on RPi 5. Instead, UWB-Fi consumes the least energy, ranging from 38–130 mJ across platforms. However, overall, BANDWEAVE offers the best trade-off between performance and efficiency, making it well-suited for resource-constrained edge platforms.

VI. CONCLUSIONS

We introduced BANDWEAVE, the first learning-based framework for coherent multi-band CFR fusion designed to enhance channel estimation resolution in MIMO systems. At the heart of BANDWEAVE is a three-phase progressive learning pipeline that integrates supervised pretraining, simulation-guided fine-tuning, and online feedback-aware adaptation. This design enables robust and adaptive CFR fusion in dynamic environments, including LoS and NLoS setups, effectively mitigating cross-band issues – such as delay misalignment, phase offsets, and frequency-dependent distortions – and generalizes across contiguous and non-contiguous frequency bands. Evaluations on IEEE 802.11ac MU-MIMO and mmWave SU-MIMO testbeds show that BANDWEAVE consistently outperforms SOTA baselines, achieving up to 16% higher throughput and lower BER. The authors have provided public access to their code and data at <https://github.com/Restuccia-Group/BandWeave>.

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] C. Creed, M. Al-Kalbani, A. Theil, S. Sarcar, and I. Williams, "Inclusive AR/VR: Accessibility Barriers for Immersive Technologies," *Universal Access in the Information Society*, vol. 23, no. 1, pp. 59–73, 2024.
- [2] D. K. Murala, "METAEDUCATION: State-of-the-art Methodology for Empowering Feature Education," *IEEE Access*, vol. 12, pp. 57992–58020, 2024.
- [3] M. K. Sharma, I. Farhat, C.-F. Liu, N. Sehad, W. Hamidouche, and M. Debbah, "Real-time Immersive Aerial Video Streaming: A Comprehensive Survey, Benchmarking, and Open Challenges," *IEEE Open Journal of the Communications Society*, 2024.
- [4] K. Khan, "Advancements and Challenges in 360 Augmented Reality Video Streaming: A Comprehensive Review," *International Journal of Computing*, vol. 13, no. 1, pp. 1–20, 2024.
- [5] X. Wang, Q. Guo, Z. Ning, L. Guo, G. Wang, X. Gao, and Y. Zhang, "Integration of Sensing, Communication, and Computing for Metaverse: A Survey," *ACM Computing Surveys*, vol. 56, no. 10, pp. 1–38, 2024.
- [6] J. Wang, J. Chuang, S. Semper, and N. Golmie, "Super-Resolution Localization and Tracking in WiFi Sensing," in *2024 33rd International Conference on Computer Communications and Networks (ICCCN)*, pp. 1–9, IEEE, 2024.
- [7] J. P. Van Marter, M. Ben-Shachar, Y. Alpert, A. G. Dabak, N. Al-Dhahir, and M. Torlak, "A Multi-Channel Approach and Testbed for Centimeter-Level WiFi Ranging," *IEEE Journal of Indoor and Seamless Positioning and Navigation*, 2024.
- [8] I. of Electrical and E. E. (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)*, pp. 1–425, 2013.
- [9] Institute of Electrical and Electronics Engineers (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*, pp. 1–767, 2021.
- [10] J. Pegoraro, J. O. Lacruz, M. Rossi, and J. Widmer, "HiSAC: High-Resolution Sensing with Multiband Communication Signals," in *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, pp. 549–563, 2024.
- [11] A. Banerjee, X. Zhao, V. Chhabra, K. Srinivasan, and S. Parthasarathy, "HORCRUX: Accurate Cross Band Channel Prediction," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, pp. 1–15, 2024.
- [12] K. W. Cho, M. Cominelli, F. Gringoli, J. Widmer, and K. Jamieson, "Scalable Multi-modal Learning for Cross-link Channel Prediction in Massive IoT Networks," in *Proceedings of the Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, pp. 221–229, 2023.
- [13] X. Li, H. Wang, Z. Chen, Z. Jiang, and J. Luo, "UWB-Fi: Pushing Wi-Fi Towards Ultra-Wideband for Fine-Granularity Sensing," in *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*, pp. 42–55, 2024.
- [14] X. Li, Z. Tian, Z. Li, and K. Liu, "A Ranging Solution with Multi-channel Splicing in the Single Frequency Band," in *2023 IEEE 11th Asia-Pacific Conference on Antennas and Propagation (APCAP)*, pp. 1–2, IEEE, 2023.
- [15] Z. Tian, Y. Ren, Z. Li, and X. Cheng, "Super-Resolution Time-of-Flight Estimation for Ranging via Multi-band Splicing," in *GLOBECOM 2023-2023 IEEE Global Communications Conference*, pp. 4656–4661, IEEE, 2023.
- [16] S. Dimce, A. Zubow, and F. Dressler, "mmSplicer: Toward Experimental Multiband Channel Splicing at mmWave Frequencies," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1–2, IEEE, 2024.
- [17] A. Bakshi, Y. Mao, K. Srinivasan, and S. Parthasarathy, "Fast and Efficient Cross Band Channel Prediction using Machine Learning," in *The 25th Annual International Conference on Mobile Computing and Networking*, pp. 1–16, 2019.
- [18] Z. Liu, G. Singh, C. Xu, and D. Vasisht, "FIRE: Enabling Reciprocity for FDD MIMO Systems," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pp. 628–641, 2021.
- [19] M. S. Safari, V. Pourahmadi, and S. Sodagari, "Deep UL2DL: Data-driven Channel Knowledge Transfer from Uplink to Downlink," *IEEE Open Journal of Vehicular Technology*, vol. 1, pp. 29–44, 2019.
- [20] Y. Wang, Z. Tian, Z. Li, K. Liu, S. Gui, and C. Huang, "Integrating Multiband Channel State Information for Enhanced Ranging and Localization," in *GLOBECOM 2024-2024 IEEE Global Communications Conference*, pp. 2930–2935, IEEE, 2024.
- [21] H. Zhu, Y. Zhuo, Q. Liu, and S. Chang, " π -Splicer: Perceiving Accurate CSI Phases with Commodity WiFi Devices," *IEEE Transactions on Mobile Computing*, vol. 17, no. 9, pp. 2155–2165, 2018.
- [22] R. D. Bhavani, D. Sudhakar, et al., "Design and Implementation of Inverse Fast Fourier Transform for OFDM," *International Journal of Science and Engineering Applications*, vol. 2, no. 7, pp. 155–157, 2013.
- [23] J. Zheng, T. Su, W. Zhu, X. He, and Q. H. Liu, "Radar High-Speed Target Detection based on the Scaled Inverse Fourier Transform," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 3, pp. 1108–1119, 2014.
- [24] W. M. Snyder, "Continuous Parabolic Interpolation," *Journal of the Hydraulics Division*, vol. 87, no. 4, pp. 99–111, 1961.
- [25] 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 - Redline," *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007) - Redline*, pp. 1–5229, March 2012.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is All You Need," *Advances in neural information processing systems*, vol. 30, 2017.
- [27] F. Ait Aoudia and J. Hoydis, "End-to-end Learning of Communications Systems Without a Channel Model," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, pp. 298–303, IEEE, 2018.
- [28] I. Ahmad and S. Shin, "Channel Model for End-to-End Learning of Communications Systems: A Survey," *arXiv preprint arXiv:2204.03944*, 2022.
- [29] A. Gupta, R. Ramanath, J. Shi, and S. S. Keerthi, "Adam vs. SGD: Closing the Generalization Gap on Image Classification," in *OPT2021: 13th annual workshop on optimization for machine learning*, pp. 1–7, 2021.
- [30] M. W. Iruthayarajan and S. Baskar, "Covariance Matrix Adaptation Evolution Strategy based Design of Centralized PID Controller," *Expert systems with Applications*, vol. 37, no. 8, pp. 5775–5781, 2010.
- [31] F. Meneghello, F. Gringoli, M. Cominelli, M. Rossi, and F. Restuccia, "How to BREAK MU-MIMO Precoding in IEEE 802.11 Wi-Fi Networks," in *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*, IEEE, 2025.
- [32] K. Rumman, F. Meneghello, K. F. Haque, F. Gringoli, and F. Restuccia, "SHRINK: Reducing MIMO Feedback Overhead in Wi-Fi with Dynamic Data-Driven Channel Sounding," in *Proc. of the Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, 2025.
- [33] F. Gringoli, M. Schulz, J. Link, and M. Hollick, "Free Your CSI: A Channel State Information Extraction Platform For Modern Wi-Fi Chipsets," in *Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, (New York, NY, USA), p. 21–28, Association for Computing Machinery, 2019.