

Jointly Learned Waveforms for Residual CFO Robustness and PAPR Reduction

1st Saksham Dewan

Department of ECE, University at Albany – SUNY
Albany, NY, USA sdewan@albany.edu

3rd Asaf Tzadok

IBM T. J. Watson Research Center
New York, NY, USA atzadok@us.ibm.com

2nd Hany Elgala

Department of ECE, University at Albany – SUNY
Albany, NY, USA helgala@albany.edu

4th Ming-Ching Chang

Department of CS, University at Albany – SUNY
Albany, NY, USA mchang2@albany.edu

Abstract—Orthogonal frequency division multiplexing (OFDM) relies on a fixed sinusoidal basis that is highly sensitive to carrier-frequency offset (CFO), yet conventional mitigation strategies primarily correct CFO at the receiver while leaving the transmit waveform fixed. This paper proposes a two-stage residual CFO-aware transceiver framework that combines a jointly learned transmit/receive waveform basis with a model-guided unfolded receiver, USR-Net. Stage 1 replaces OFDM’s fixed Fourier structure with a learned waveform basis trained to reduce symbol-domain leakage under bounded residual CFO, while also incorporating explicit peak-to-average power ratio (PAPR) regularization. Stage 2 keeps the transmit waveform fixed and refines the recovered symbols using the known residual CFO operator structure. For 16-QAM at $E_b/N_0 = 10$ dB, the full two-stage system extends the CFO-robustness window eightfold over standard OFDM at bit-error rate (BER) 10^{-2} , and reduces error vector magnitude (EVM) by 60% at a normalized residual CFO of 0.10. With PAPR regularization, the learned waveform reduces the 95th-percentile PAPR by up to 2.22 dB below the OFDM baseline while maintaining lower BER at the evaluated residual CFO operating points. These results establish operator-aware design as a practical direction for waveform-level robustness.

Index Terms—Waveform design, OFDM, carrier-frequency offset, transceiver-learned waveform, inter-carrier interference, PAPR, non-orthogonal multicarrier, physical layer.

I. INTRODUCTION

Orthogonal frequency division multiplexing (OFDM) dominates modern wireless systems by enabling low-complexity frequency-domain equalization, but its fixed Fourier basis makes strict orthogonality both its primary advantage and its major vulnerability. Hardware or channel impairments that violate the orthogonal structure rapidly produce severe inter-carrier interference (ICI), degrading system performance. Non-orthogonal waveform families have been proposed in the literature, including faster-than-Nyquist signaling [1], spectrally efficient frequency division multiplexing (SEFDM) [2], filter-bank multicarrier [3], and index-modulation-based designs [4]. These works demonstrate that orthogonality is not always necessary for reliable communication, but their designs leave the interference burden entirely at the receiver, with no transmitter-side mitigation. This motivates a more

principled approach: shifting part of this burden to the transmitter by learning a waveform basis that intrinsically reduces interference.

A direct step in this direction is non-orthogonal frequency shaping (NOFS) [5], which treats the multicarrier waveform basis as a learnable object, where the basis refers to the set of subcarrier pulse shapes replacing OFDM’s fixed sinusoids. In NOFS, a neural modulation matrix generates irregular sinc-like (irSinc) shapes trained to reduce internally generated ICI under spectral compression, demonstrating that the transmit basis can be optimized so that the effective symbol-domain interference structure becomes more favorable for the receiver. The present work applies that principle to a different and practically important impairment: residual carrier-frequency offset (CFO) that persists after conventional estimation and correction.

CFO is one of the most consequential synchronization impairments in multicarrier systems, arising from oscillator mismatch, Doppler shift, or residual tracking error, and it destroys subcarrier orthogonality through a common phase rotation and ICI across the entire subcarrier set. Classical synchronization procedures for frequency and timing offset estimation are established in [6], [7]. More recently, learning-based approaches [8]–[10] have revisited CFO compensation through deep and model-driven receivers, but in each case the transmit waveform remains fixed and the receiver alone corrects the CFO damage. This is particularly relevant in contested and congested wireless environments, where oscillator instability and mobility make residual CFO a persistent impairment. More broadly, since precision scales with cost, a waveform tolerant of low-grade oscillators could benefit cost-sensitive internet-of-things (IoT) and edge deployments where high-accuracy crystal references are impractical. Such scenarios also demand strict power efficiency; high peak-to-average power ratio (PAPR) forces power amplifiers to operate with backoff to avoid the nonlinear distortion region. A waveform basis that simultaneously improves CFO robustness and reduces PAPR addresses two front-end impairments, offering advantages that purely receiver-side corrections cannot

provide.

Recent studies on end-to-end learning for physical-layer communications have shown that a transmitter-receiver pair can be jointly trained offline over a known channel model, converging from a standard protocol to a shared learned waveform basis that replaces conventional fixed modulation [11], and recent work has extended this to non-orthogonal multicarrier settings targeting spectral efficiency and PAPR reduction [12], [13]. In much of this literature, however, the learned transmit mapping serves as a constellation or coding layer rather than a waveform basis. Whether the basis itself can be shaped to reduce the symbol-domain damage from a specific front-end impairment like residual CFO has not been directly addressed. This paper directly addresses that gap.

Inspired by the NOFS view of waveform synthesis as a learnable basis operation, a two-stage framework is proposed that jointly learns a transmit synthesis matrix and a corresponding linear receive matrix to reduce symbol-domain leakage over a bounded residual CFO interval. Explicit PAPR regularization is further incorporated at the transmitter. Evaluations in the present work are restricted to additive white Gaussian noise (AWGN) and residual CFO to isolate the waveform basis contribution before introducing additional impairments. The contributions of this paper are as follows:

- A residual CFO-aware jointly learned transmit/receive basis is proposed as an alternative to OFDM's fixed Fourier structure, using a transmit matrix and a linear receive matrix that target the bounded residual CFO persisting after conventional coarse correction. An explicit PAPR regularization term is incorporated into the training objective, addressing two front-end impairments through a jointly learned transceiver design (discussed in Section III).
- A model-guided unfolded receiver (USR-Net) is proposed as a receiver-side enhancement that keeps the transmit waveform fixed. USR-Net exploits the known residual CFO operator structure to iteratively refine the received symbol estimates and recover remaining symbol-domain distortion (discussed in Section III).
- Simulations under AWGN and residual CFO demonstrate an eightfold extension of the CFO-robustness window over standard OFDM at BER 10^{-2} , a 60% EVM reduction at normalized residual CFO of 0.10, and a 2.22 dB reduction in 95th-percentile PAPR below the OFDM baseline.

II. RESIDUAL CFO-RESILIENT LEARNED WAVEFORM BASIS

The proposed system comprises two stages, illustrated in Fig. 1. Stage 1 jointly learns a transmit synthesis matrix and a linear receive analysis matrix whose effective symbol-domain operator remains diagonally dominant under bounded residual CFO, while simultaneously reducing PAPR through explicit waveform regularization. Stage 2 refines the resulting symbol estimates without modifying the transmit waveform, and is described in Section III.

Consider an M -sample block whose frequency-domain resources are partitioned into K payload, P

pilot/synchronization, and G guard/null dimensions, with $M = K + P + G$. The number of information-bearing symbols is $N \leq K$, with $R = K - N$ in-band redundancy dimensions available for waveform shaping. Let $\mathbf{s} \in \mathbb{C}^N$ denote the transmit data-symbol vector.

A. Waveform Model

OFDM baseline: Let $\mathbf{F}_M \in \mathbb{C}^{M \times M}$ denote the unitary M -point DFT matrix and $\mathbf{P}_D \in \mathbb{C}^{M \times N}$ select the N data tones from the payload region. The OFDM synthesis and analysis matrices are $\mathbf{W}_{\text{OFDM}} = \mathbf{F}_M^H \mathbf{P}_D$ and $\mathbf{V}_{\text{OFDM}} = \mathbf{P}_D^H \mathbf{F}_M$, yielding $\hat{\mathbf{s}} = \mathbf{s}$ under ideal conditions.

Learned basis: The proposed waveform replaces these fixed Fourier matrices with learned matrices $\mathbf{W} \in \mathbb{C}^{K \times N}$ and $\mathbf{V} \in \mathbb{C}^{N \times M}$. Let $\mathbf{P}_K \in \mathbb{C}^{M \times K}$ embed the K payload dimensions into the M -point frequency grid while leaving the pilot and guard dimensions inactive. The effective time-domain synthesis matrix is $\mathbf{W}_{\text{tx}} = \mathbf{F}_M^H \mathbf{P}_K \mathbf{W} \in \mathbb{C}^{M \times N}$, constrained to the K -dimensional payload. The transmitted block is $\mathbf{x} = \mathbf{W}_{\text{tx}} \mathbf{s}$, and the clean-channel reconstruction condition is $\mathbf{V} \mathbf{W}_{\text{tx}} \approx \mathbf{I}_N$.

Residual CFO channel: Let ϵ denote the CFO normalized to the subcarrier spacing $\Delta f_{\text{sc}} = f_s/M$. After conventional coarse correction using estimate $\hat{\epsilon}$, the normalized residual CFO is $\delta = \epsilon - \hat{\epsilon}$, so $|\delta| = 1$ corresponds to one full subcarrier spacing of residual offset. This induces a time-domain phase ramp

$$\Phi_\delta = \text{diag} \left(e^{j2\pi\delta n/M} \right)_{n=0}^{M-1}. \quad (1)$$

In practice, oscillator instability, Doppler shift, and imperfect coarse synchronization can push $|\delta|$ well above the threshold tolerable by conventional receivers, motivating the $|\delta| \leq 0.20$ robustness interval targeted in this study. The received block is: $\mathbf{y} = \Phi_\delta \mathbf{W}_{\text{tx}} \mathbf{s} + \mathbf{n}$, where $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_M)$ is AWGN.

For BER evaluation, the Stage 1 hard-decision estimate $\hat{\mathbf{s}}_{\text{lin}}$ is obtained by nearest-constellation-point decision on $\mathbf{z}_0$, while the Stage 2 output $\hat{\mathbf{s}}_{\text{usr}} = \hat{\mathbf{s}}^{(3)}$ is the final estimate after three unfolded iterations.

B. Effective Symbol-Domain Operator

Applying the learned analysis matrix $\mathbf{V}$ yields the Stage 1 linear receive output $\mathbf{z}_0 = \mathbf{V} \mathbf{y} = \mathbf{A}(\delta) \mathbf{s} + \mathbf{V} \mathbf{n}$, where the residual CFO-induced symbol-domain operator is defined as

$$\mathbf{A}(\delta) = \mathbf{V} \Phi_\delta \mathbf{W}_{\text{tx}}. \quad (2)$$

Decomposing $\mathbf{A}(\delta) = \mathbf{D}(\delta) + \mathbf{O}(\delta)$, where $\mathbf{D}(\delta) = \text{diag}(\mathbf{A}(\delta))$ captures per-symbol gain and phase distortion and $\mathbf{O}(\delta) = \mathbf{A}(\delta) - \mathbf{D}(\delta)$ captures off-diagonal symbol mixing, reveals the structural damage caused by residual CFO. When $\mathbf{O}(\delta)$ is small, each recovered symbol undergoes only a scalar distortion that Stage 2 can correct efficiently.

The design goal of Stage 1 is therefore to suppress $\mathbf{O}(\delta)$ and keep $\mathbf{A}(\delta)$ diagonally dominant over a bounded interval $\delta \in [-\delta_{\text{max}}, \delta_{\text{max}}]$. This formulation also explains why broad CFO immunity is not a suitable target for a fixed basis. In the square invertible case, clean reconstruction implies




Fig. 1. Proposed two-stage transceiver. The transmitter maps data symbols $\mathbf{s}$ through the learned synthesis matrix $\mathbf{W}_{\text{tx}}$. The channel applies a residual CFO phase ramp Φ_δ and AWGN. Stage 1 applies the learned analysis matrix $\mathbf{V}$ to produce the linear receive output $\mathbf{z}_0$, from which symbol estimates $\hat{\mathbf{s}}_{\text{lin}}$ are obtained. Stage 2 computes the diagonal correction vector $\hat{\mathbf{d}}$ from the residual-CFO operator model and uses it with $\mathbf{z}_0$ to form the model-guided reference $\mathbf{r}$, which initializes and conditions the USR-Net unfolded receiver. Three learned neural update iterations of the form $\hat{\mathbf{s}}^{(t+1)} = \hat{\mathbf{s}}^{(t)} + \eta_t \Delta_t$ yield the final estimate $\hat{\mathbf{s}}_{\text{usr}} = \hat{\mathbf{s}}^{(3)}$. The transmit waveform is fixed throughout Stage 2.

$\mathbf{V} \approx \mathbf{W}_{\text{tx}}^{-1}$. Then $\mathbf{A}(\delta) \approx \mathbf{W}_{\text{tx}}^{-1} \Phi_\delta \mathbf{W}_{\text{tx}}$, which is a similarity transformation of the CFO phase-ramp operator. Changing the basis redistributes CFO distortion across the symbol domain but cannot remove the underlying phase ramp. Therefore, the objective is to find a basis for which this redistribution leaves $\mathbf{A}(\delta)$ as close as possible to a scaled identity $\beta(\delta) \mathbf{I}_N$, where $\beta(\delta) = \frac{1}{N} \text{tr}(\mathbf{A}(\delta))$.

C. Training Objective

The Stage 1 training objective jointly optimizes $\mathbf{W}$ and $\mathbf{V}$ to suppress residual CFO-induced symbol mixing, preserve clean-channel reconstruction, condition the receiver against noise amplification, and directly penalize PAPR. The core objective penalizes deviation of $\mathbf{A}(\delta)$ from a scaled identity over the training interval:

$$\mathcal{L}_{\text{cfo}} = \mathbb{E}_\delta \left[N^{-1} \|\mathbf{A}(\delta) - \beta(\delta) \mathbf{I}_N\|_F^2 \right] \quad (3)$$

A common phase rotation or gain change is not penalized; off-diagonal leakage and stream-to-stream gain disparity are. To ensure the learned pair forms an approximately inverse basis at zero offset:

$$\mathcal{L}_{\text{clean}} = N^{-1} \|\mathbf{V} \mathbf{W}_{\text{tx}} - \mathbf{I}_N\|_F^2 \quad (4)$$

To connect the operator objective directly to communication performance, a noisy symbol recovery loss is included:

$$\mathcal{L}_{\text{sym}} = \mathbb{E}_{\mathbf{s}, \delta, \mathbf{n}} \left[\frac{\|\hat{\mathbf{s}}_{\text{lin}} - \mathbf{s}\|_2^2}{\|\mathbf{s}\|_2^2} \right] \quad (5)$$

Large $\mathbf{V}$ coefficients can amplify noise; a Frobenius penalty prevents this:

$$\mathcal{L}_V = N^{-1} \|\mathbf{V}\|_F^2 \quad (6)$$

Finally, front-end power variation of the transmitted block $\mathbf{x} = \mathbf{W}_{\text{tx}} \mathbf{s}$ is directly penalized via a PAPR regularization term:

$$\mathcal{L}_{\text{papr}} = \mathbb{E}_{\mathbf{s}} \left[\frac{M \cdot \max_n \|\mathbf{W}_{\text{tx}} \mathbf{s}\|_n^2}{\|\mathbf{W}_{\text{tx}} \mathbf{s}\|_2^2} \right] \quad (7)$$

The complete Stage 1 training loss is therefore:

$$\mathcal{L}_{\text{total}} = \lambda_0 \mathcal{L}_{\text{cfo}} + \lambda_1 \mathcal{L}_{\text{clean}} + \lambda_2 \mathcal{L}_{\text{sym}} + \lambda_3 \mathcal{L}_V + \lambda_4 \mathcal{L}_{\text{papr}}. \quad (8)$$

III. UNFOLDED SYMBOL REFINEMENT NETWORK (USR-NET)

After Stage 1, the received representation is $\mathbf{z}_0 = \mathbf{A}(\delta) \mathbf{s} + \mathbf{Vn}$, where the learned basis has suppressed the off-diagonal component $\mathbf{O}(\delta)$, leaving a dominant diagonal distortion $\mathbf{D}(\delta)$. Stage 2 corrects this remaining distortion without modifying $\mathbf{W}_{\text{tx}}$ or $\mathbf{V}$. Drawing on the deep unfolding framework [14], in which iterative algorithms are unfolded into trainable neural network layers, USR-Net operates as a model-guided unfolded receiver with three learned neural update iterations that produce the final symbol estimate.

A. USR-Net Model

Model-guided reference: After the Stage 1 front-end, the received symbols satisfy $\mathbf{z}_0 = \mathbf{A}(\delta) \mathbf{s} + \mathbf{n}'$ and $\mathbf{A}(\delta) = \mathbf{V} \Phi_\delta \mathbf{W}_{\text{tx}}$. The diagonal entries of $\mathbf{A}(\delta)$ describe the dominant per-symbol gain and phase distortion, while the off-diagonal entries describe residual symbol leakage. A per-symbol diagonal correction vector is computed using the known operator form as a structural prior. Specifically, let $\hat{\mathbf{A}} = \mathbf{V} \Phi_{\hat{\delta}} \mathbf{W}_{\text{tx}}$, where $\hat{\delta}$ denotes the residual-CFO value available to the receiver, and the per-symbol diagonal correction vector is $\hat{\mathbf{d}} = \text{diag}(\hat{\mathbf{A}})$. In simulation, the true residual

offset is used and pilot-based estimation is reserved for future work. The model-guided reference is then computed as

$$\mathbf{r} = \mathbf{z}_0 \oslash (\mathbf{1} + \gamma(\hat{\mathbf{d}} - \mathbf{1})), \quad (9)$$

where $\oslash$ denotes element-wise division and $\gamma \in [0, 1]$ is a learned correction-strength parameter that interpolates between no correction ($\gamma = 0$) and full diagonal inversion ($\gamma = 1$). This reference corrects the dominant per-symbol gain and phase distortion $\mathbf{D}(\delta)$. Both $\mathbf{r}$ and $\hat{\mathbf{d}}$ are computed once and supplied as fixed conditioning inputs to every unfolded iteration.

Iterative update: The network is initialized at the model-guided reference, $\hat{\mathbf{s}}^{(0)} = \mathbf{r}$, and each unfolded layer refines the current estimate via a learned residual step:

$$\hat{\mathbf{s}}^{(t+1)} = \hat{\mathbf{s}}^{(t)} + \eta_t \Delta_t, \quad (10)$$

unfolded for three iterations to yield the final estimate $\hat{\mathbf{s}}_{\text{usr}} = \hat{\mathbf{s}}^{(3)}$. The learned scalar η_t controls the update step size at iteration t .

Neural residual branch: At each iteration, $\Delta_t = \mathcal{N}_{\theta_t}(\hat{\mathbf{s}}^{(t)}, \mathbf{z}_0, \mathbf{r}, \hat{\mathbf{d}})$ is a symbol-estimate correction produced by independent residual dilated Conv1D blocks. For each iteration, the input is a 10-dimensional real-valued feature map per symbol, comprising the real and imaginary parts of $\hat{\mathbf{s}}^{(t)}$, $\mathbf{z}_0$, $\mathbf{r}$, and $\hat{\mathbf{s}}^{(t)} - \mathbf{r}$, together with the magnitude and phase of $\hat{\mathbf{d}}$. Within each block, three Conv1D layers with dilation factors $\{1, 2, 4\}$ and GELU process the input. After offline training, Stage 1 uses fixed dense transforms with $O(MN)$ complexity per block, compared with $O(M \log M)$ FFT-based OFDM; USR-Net adds three Conv1D refinement iterations with 80,463 parameters, which is orders of magnitude smaller than comparable Transformer- or LSTM-based receivers [15].

B. Training Objective

Stage 2 is trained with all Stage 1 parameters frozen. The primary loss applies progressive supervision across all three iterations:

$$\mathcal{L}_{\text{sym}} = \sum_{t=1}^3 w_t \frac{\|\hat{\mathbf{s}}^{(t)} - \mathbf{s}\|_2^2}{\|\mathbf{s}\|_2^2}, \quad (11)$$

so that every unfolded layer is encouraged to improve the estimate, not only the final output. A constellation alignment loss steers the final estimate into correct decision regions:

$$\mathcal{L}_{\text{cls}} = \text{CE}\left(-\frac{1}{T_c} |\hat{\mathbf{s}}_{\text{usr}} - \mathbf{c}|^2, m_{\text{true}}\right), \quad (12)$$

where $\text{CE}(\cdot, \cdot)$ denotes cross-entropy loss, m_{true} is the true constellation index, $\mathbf{c}$ denotes the 16-QAM constellation points and T_c is a temperature parameter. An identity loss $\mathcal{L}_{\text{id}}$ prevents Stage 2 from degrading symbol recovery under clean or near-clean channel conditions, and a correction penalty $\mathcal{L}_{\text{corr}}$ discourages large deviations from $\mathbf{z}_0$. The complete Stage 2 objective is

$$\mathcal{L}_{\text{USR}} = \lambda_5 \mathcal{L}_{\text{sym}} + \lambda_6 \mathcal{L}_{\text{cls}} + \lambda_7 \mathcal{L}_{\text{id}} + \lambda_8 \mathcal{L}_{\text{corr}}. \quad (13)$$

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Modulation	16QAM
Training E_b/N_0	15 dB
Evaluation E_b/N_0	10 dB
Channel	Residual CFO + AWGN
Residual CFO grid	$[-0.20, 0.20]$
SNR sweep	0–20 dB
$\lambda_0, \lambda_1, \lambda_2, \lambda_3, \lambda_4$	1.0, 1.0, 0.18, 5×10^{-3} , 10^{-4}
$\lambda_5, \lambda_6, \lambda_7, \lambda_8$	1.0, 0.3, 0.1, 10^{-3}

IV. RESULTS AND DISCUSSION

This section evaluates the proposed two-stage system against classical OFDM across four receiver configurations: *OFDM* (linear receiver with fixed Fourier basis), *OFDM + USR-Net* (same basis with learned unfolded refinement), *Learned Basis* (learned transmit/receive pair with linear receiver), and *Learned Basis + USR-Net* (full two-stage system). The evaluation is restricted to the AWGN and residual CFO channel to isolate the contribution of each stage before introducing additional impairments such as multipath or phase noise.

A. Simulation Setup

Each transmitted block contains $M = 64$ time-domain samples, with the frequency grid partitioned into $K = 50$ payload, $P = 6$ pilot, and $G = 8$ guard dimensions. Of the K payload tones, $N = 45$ carry data, leaving $R = 5$ redundant dimensions for waveform shaping. The pilot dimensions are reserved for a future synchronization budget and are not used for CFO estimation here. Both stages are trained at $E_b/N_0 = 15$ dB and evaluated at 10 dB over the AWGN and residual CFO channel to isolate the waveform basis contribution. Simulation parameters are summarized in Table I, where all λ_i are empirically tuned hyperparameters.

B. Residual CFO BER Performance

Figure 2 shows that all configurations achieve comparable zero-CFO performance ($\text{BER} \approx 6 \times 10^{-4}$ at $E_b/N_0 = 10$ dB), confirming the proposed methods preserve clean-channel integrity. As $|\delta|$ increases, the USR-Net configurations diverge




Fig. 2. BER versus normalized residual CFO at $E_b/N_0 = 10$ dB for all four receiver configurations, illustrating the progressive extension in CFO tolerance achieved by each stage of the proposed system.




Fig. 3. BER versus E_b/N_0 at fixed residual CFO offsets $\delta = 0.05$ (top) and $\delta = 0.10$ (bottom), showing how the noise-BER relationship evolves across all four configurations as CFO severity increases.

sharply toward lower BER. The CFO tolerance window, defined as the maximum $|\delta|$ sustaining a given BER target, reveals the clearest separation: At a 10^{-2} BER target, OFDM and the Learned Basis reach only $|\delta| \leq 0.025$. While OFDM + USR-Net extends to $|\delta| \leq 0.075$, the Learned Basis + USR-Net maintains performance across the entire evaluated range, an eightfold extension over standard OFDM.

C. SNR Generalization

Figure 3 evaluates BER versus E_b/N_0 at fixed residual CFO offsets $\delta = 0.05$ and $\delta = 0.10$. At $\delta = 0.05$, the Learned Basis maintains a clear advantage over OFDM at moderate and high SNR, while Learned Basis + USR-Net sustains a steep BER falloff where linear receivers saturate, with a 10 dB SNR gain at 10^{-1} BER over OFDM. At $\delta = 0.10$, OFDM and the Learned Basis both saturate near 10^{-1} , whereas USR-Net configurations continue to improve with SNR. Learned Basis + USR-Net maintains the steepest falloff at both CFO levels, confirming generalization beyond the training SNR.

D. Constellation-Level Behavior

Figure 4 shows decoded 16-QAM constellations for all four configurations at $\delta = 0.05$ and $\delta = 0.10$. At $\delta = 0.05$, OFDM exhibits noticeable cluster spreading and phase rotation, while the Learned Basis maintains more compact clusters under the same conditions. Both USR-Net configurations recover well-separated decision regions, with Learned Basis + USR-Net producing the tightest clusters. At $\delta = 0.10$, OFDM constellations are severely degraded with strong inter-cluster leakage; the Learned Basis reduces but does not eliminate this distortion, and only the USR-Net refinement step restores a recognizable 16-QAM grid. The corresponding EVM decreases from 0.380 for OFDM to 0.317 for the Learned Basis, 0.232 for OFDM + USR-Net, and 0.151 for Learned Basis +




Fig. 4. Decoded 16-QAM constellations at $\delta = 0.05$ and $\delta = 0.10$ for all four receiver configurations, illustrating the progressive improvement in cluster separation and phase stability as each stage of the proposed system is added.

USR-Net, a 60% reduction relative to OFDM for the complete two-stage system.

E. Peak-to-Average Power Ratio (PAPR)

As Stage 2 does not alter the transmit waveform, this PAPR profile characterizes the Stage 1 output for both OFDM and Learned configurations. CFO-aware learning also reduces the 95th-percentile PAPR from 8.70 dB for OFDM to 8.21 dB even without explicit PAPR regularization ($\lambda_4 = 0$). This intrinsic 0.49 dB reduction is notable given the limited shaping headroom of only $R = 5$ redundant dimensions out of $K = 50$ payload dimensions. Activating $\mathcal{L}_{\text{papr}}$ provides an

TABLE II
PAPR-BER OPERATING POINTS UNDER $\mathcal{L}_{\text{papr}}$ REGULARIZATION

λ_4	p95 PAPR (dB)	BER ($\delta=0.05$)	BER ($\delta=0.10$)
OFDM	8.70	1.34×10^{-2}	1.67×10^{-1}
0	8.21	8.84×10^{-3}	7.04×10^{-2}
10^{-4}	7.64	9.37×10^{-3}	6.88×10^{-2}
3×10^{-3}	6.48	1.16×10^{-2}	1.43×10^{-2}

additional PAPR-BER tradeoff, as shown in Table II. Setting $\lambda_4 = 10^{-4}$ reduces the 95th-percentile PAPR to 7.64 dB with little BER change relative to the unregularized learned waveform. Stronger regularization ($\lambda_4 = 3 \times 10^{-3}$) further reduces PAPR to 6.48 dB, a 2.22 dB reduction below OFDM, while maintaining lower BER than OFDM at both evaluated CFO offsets. These results indicate that PAPR reduction can be incorporated into the learned waveform design without eliminating the residual-CFO robustness gains.

V. DISCUSSION AND CONCLUSION

This work demonstrates that the transmit waveform need not be treated as a fixed constraint imposed on the receiver. Rather than relying on OFDM's fixed Fourier structure, jointly learning the transmit and receive basis allows the transmit waveform itself to reduce the impact of front-end impairments before demodulation. By formulating residual CFO sensitivity as an operator-shaping problem and suppressing off-diagonal leakage in the residual CFO-impaired symbol-domain operator over a bounded CFO interval, the proposed method learns a transmit basis that produces a more diagonally dominant symbol representation and improves recoverability after CFO distortion. In scenarios characterized by mobility-induced residual CFO and stringent power-efficiency requirements, the proposed framework addresses residual-CFO robustness and PAPR reduction within a unified learned transceiver design. Stage 1 generates a representation that is intrinsically more robust to CFO-induced leakage, while Stage 2 efficiently removes the remaining distortion. The two stages are therefore complementary, and the results suggest a broader principle: operator-aware basis shaping and model-guided unfolded refinement are most effective when co-designed.

Several future directions remain open, with the most immediate being integration of the reserved P pilot dimensions into USR-Net reference construction. A structured pilot-based CFO estimate could replace the oracle residual offset currently assumed in Stage 2, improving robustness at low SNR where offset estimation error is most consequential. Another direction is to extend the current linear receive matrix to nonlinear architectures such as CNNs, LSTMs, or Transformers [15]. More fundamentally, the generalization of the operator-aware shaping principle beyond the AWGN-plus-CFO setting remains an important open question. The present study intentionally isolates residual CFO to evaluate the contribution of transmit-basis design; whether similar benefits can be achieved in jointly impaired environments involving

multipath propagation, time-varying channels, and additional hardware nonidealities warrants further investigation.

ACKNOWLEDGMENT

This work was supported by the Center for Emerging Artificial Intelligence Systems (CEAIS), a collaborative initiative between the University at Albany-SUNY, and IBM.

REFERENCES

- [1] J. B. Anderson, F. Rusek, and V. Owall, "Faster-Than-Nyquist Signaling," *Proceedings of the IEEE*, vol. 101, no. 8, p. 1817-1830, Aug. 2013. [Online]. Available: <http://dx.doi.org/10.1109/JPROC.2012.2233451>
- [2] I. Darwazeh, T. Xu, and R. C. Grammenos, "Bandwidth-compressed Multicarrier Communication: SEFDM," p. 90-114, Aug. 2016. [Online]. Available: <http://dx.doi.org/10.1002/9781119116493.ch5>
- [3] P. Siohan, C. Siclet, and N. Lacaille, "Analysis and design of OFDM/OQAM systems based on filterbank theory," *IEEE Transactions on Signal Processing*, vol. 50, no. 5, p. 1170-1183, May 2002. [Online]. Available: <http://dx.doi.org/10.1109/78.995073>
- [4] Y. Chen, T. Xu, and I. Darwazeh, "Index Modulation Pattern Design for Non-Orthogonal Multicarrier Signal Waveforms," *IEEE Transactions on Wireless Communications*, vol. 21, no. 10, p. 8507-8521, Oct. 2022. [Online]. Available: <http://dx.doi.org/10.1109/TWC.2022.3166850>
- [5] T. Xu and I. Darwazeh, "Identification and practical validation of spectrally efficient non-orthogonal frequency shaping waveform," *Communications Engineering*, vol. 2, no. 1, Aug. 2023. [Online]. Available: <http://dx.doi.org/10.1038/s44172-023-00108-w>
- [6] T. Schmidl and D. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Transactions on Communications*, vol. 45, no. 12, p. 1613-1621, 1997. [Online]. Available: <http://dx.doi.org/10.1109/26.650240>
- [7] H. Bolcskei, "Blind estimation of symbol timing and carrier frequency offset in wireless OFDM systems," *IEEE Transactions on Communications*, vol. 49, no. 6, p. 988-999, Jun. 2001. [Online]. Available: <http://dx.doi.org/10.1109/26.930629>
- [8] Z. Wang, S. Wei, L. Zou, F. Liao, W. Lang, and Y. Li, "Deep-Learning-Based Carrier Frequency Offset Estimation and Its Cross-Evaluation in Multiple-Channel Models," *Information*, vol. 14, no. 2, p. 98, Feb. 2023. [Online]. Available: <http://dx.doi.org/10.3390/info14020098>
- [9] X. Lin, Y. Shen, and C. Jiang, "A Model-Driven Deep Learning-Based Receiver for OFDM System With Carrier Frequency Offset," *IEEE Communications Letters*, vol. 28, no. 4, p. 813-817, Apr. 2024. [Online]. Available: <http://dx.doi.org/10.1109/LCOMM.2024.3354990>
- [10] V. Ninkovic, A. Valka, D. Dumić, and D. Vukobratovic, "Deep Learning-Based Packet Detection and Carrier Frequency Offset Estimation in IEEE 802.11ah," *IEEE Access*, vol. 9, p. 99853-99865, 2021. [Online]. Available: <http://dx.doi.org/10.1109/ACCESS.2021.3096853>
- [11] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, p. 563-575, Dec. 2017. [Online]. Available: <http://dx.doi.org/10.1109/TCCN.2017.2758370>
- [12] B. Zhu, J. Wang, L. He, and J. Song, "Joint Transceiver Optimization for Wireless Communication PHY Using Neural Network," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, p. 1364-1373, June 2019. [Online]. Available: <http://dx.doi.org/10.1109/JSAC.2019.2904361>
- [13] C. Liu, G. Liu, F. Xu, Q. Li, H. Liu, L. Zhang, and M. Ali Imran, "End-to-End Optimized Non-Orthogonal Multicarrier Waveform Design via Deep Learning," *IEEE Transactions on Communications*, vol. 73, no. 12, p. 14444-14457, Dec. 2025. [Online]. Available: <http://dx.doi.org/10.1109/TCOMM.2025.3615782>
- [14] S. Deka, K. Deka, N. Thanh Nguyen, S. Sharma, V. Bhatia, and N. Rajatheva, "Comprehensive Review of Deep Unfolding Techniques for Next-Generation Wireless Communication Systems," *IEEE Internet of Things Journal*, vol. 13, no. 6, p. 10379-10406, Mar. 2026. [Online]. Available: <http://dx.doi.org/10.1109/JIOT.2026.3652234>
- [15] M. S. Sarwar, M. Awais, and J. Hossain, "Transformer-Based Multi-Class Detection for Multi-Mode SEFDM with Global Inter-Carrier Coupling," *IEEE Transactions on Communications*, p. 1-1, 2026. [Online]. Available: <http://dx.doi.org/10.1109/TCOMM.2026.3681659>