Deep-OFDM: Robust Neural Modulation for High Mobility

Sravan Kumar Ankireddy [*1] **S. Ashwin Hebbar** [*2] **Pramod Viswanath** [2] **Hyeji Kim** [1]

Abstract

Orthogonal Frequency Division Multiplexing (OFDM) is the workhorse of current 5G deployments due to its robustness in quasi-static channels and efficient spectrum use. However, in high-mobility scenarios, OFDM suffers from inter-carrier interference (ICI), and its reliance on dense pilot patterns and cyclic prefixes reduces spectral efficiency significantly. In this work, we propose Deep-OFDM : a learnable modulation framework that augments traditional OFDM by incorporating neural parameterization. Instead of mapping each symbol to a fixed resource element, Deep-OFDM spreads information across the OFDM grid using a convolutional neural network modulator. This modulator is jointly optimized with a neural receiver through end-to-end training, enabling the system to adapt to time-varying channels without relying on explicit channel estimation. Deep-OFDM outperforms conventional OFDM when paired with neural receiver baselines, particularly in pilot-sparse and pilotless regimes, achieving substantial gains in BLER and goodput, particularly at high Doppler. In the pilotless setting, the neural modulator learns a low-rank structure that resembles a superimposed pilot, effectively enabling reliable communication without explicit overhead. These results highlight the potential of transmitter–receiver co-design for robust, resource-efficient communication in challenging channel conditions, paving the way for AI-native PHY designs in next-generation wireless systems.

1. Introduction

Orthogonal frequency-division multiplexing (OFDM) is the dominant waveform in 4G and 5G networks. It transforms a multipath fading channel into a set of parallel flat-fading subchannels, enabling simple equalization and efficient use of the spectrum. As a result, OFDM delivers strong average-case performance in quasi-static channels, making it well-suited for current 5G deployments.

However, this performance degrades significantly in high-mobility environments, where the wireless channel varies rapidly in both time and frequency. In such settings (eg - vehicular communications and high-speed rail scenarios), Doppler spread destroys subcarrier orthogonality, causing inter-carrier interference and rendering the standard one-tap equalizer ineffective (Ai et al., 2014). Moreover, pilot-aided channel estimation becomes unreliable as coherence time shrinks, and pilot symbols must be injected more frequently to track the channel.

To compensate for these effects, conventional systems resort to increasing pilot density and extending the cyclic prefix (CP). While these measures help maintain performance, they reduce spectral efficiency: pilots consume time-frequency resources that could otherwise carry data, and the CP occupies a growing fraction of the OFDM symbol; particularly at high carrier frequencies or with fewer subcarriers, where the absolute delay spread remains fixed but symbol durations shrink. These issues are particularly severe in uplink scenarios, where transmit SNR is low and efficient resource usage is crucial.

These limitations are expected to intensify with the demands of next-generation systems. Proposed 6G standards call for ultra-reliable, low-latency communication at data rates up to 1 Tbps, and support for mobility up to 1000 km/h (Tong & Zhu, 2022; Chafii et al., 2023). Meeting these goals will require methods that are fundamentally more robust to rapid channel variation, and that can deliver high performance even under sparse pilot conditions. While recent efforts have proposed new waveform designs, such as OTFS (Hadani et al., 2017), to better handle channel dynamics, we take a complementary approach in this work.

Designing new waveforms is challenging, as it requires balancing spectral efficiency, robustness, and implementation feasibility. OFDM represents a carefully engineered trade-off across these axes, and entirely replacing it risks losing critical system-level advantages. Instead, we adopt a hy-

[*]Equal contribution [1]University of Texas at Austin, USA [2]Princeton University, USA. Correspondence to: Sravan Kumar Ankireddy <sravan.ankireddy@utexas.edu>, S. Ashwin Hebbar <hebbar@princeton.edu>.

brid approach known as neural augmentation—building on the OFDM framework while addressing its limitations using learning-based techniques. This paradigm has shown promising results in areas like source and channel coding (Kim et al., 2018; Hebbar et al., 2022; Ankireddy & Kim, 2023; Li et al., 2023; Hebbar et al., 2024).

Following this principle, we retain the OFDM structure but learn a neural modulation scheme tailored to time-varying channels. Recent work has shown that neural receivers can replace traditional OFDM receivers, achieving competitive performance without relying on explicit channel estimation (Ye et al., 2017; Honkala et al., 2021; Wiesmayr et al., 2024). These data-driven receivers are not restricted to conventional constellations and can adapt to arbitrary modulation patterns, often with improved robustness in time-varying conditions.

Previous work has indicated gains in transmitter co-design with neural receivers, including learning better constellations, and combining this with novel piloting schemes like Superimposed pilots (SIP) (Stark et al., 2019; Aoudia & Hoydis, 2021; Madadi et al., 2023). In our work, we propose Deep-OFDM : a two-dimensional time-frequency spreading scheme implemented via a learnable neural modulator. Unlike full joint coded modulation approaches, which learn both coding and modulation end-to-end (Jiang et al., 2019; Makkuva et al., 2021; Jamali et al., 2022; ?) but face scalability issues, Deep-OFDM retains the practical modularity of standard systems by fixing LDPC coding and decoding. This hybrid strategy captures many of the robustness benefits of joint optimization, while remaining scalable and easy to integrate into existing OFDM pipelines. Our focus is on high-mobility channels where conventional designs degrade, and we find that performance gains from neural modulation are most pronounced under sparse pilot configurations, where traditional receivers exhibit diminishing returns from further pilot insertion.

Our main contributions are summarized as follows:

- We propose Deep-OFDM, a neural modulation scheme that performs learnable time-frequency spreading, explicitly optimized to align with the inductive biases of CNN-based neural receivers. This design yields substantial performance gains in doubly-selective channels.
- We characterize the interplay between pilot density and neural modulation. Our results show that Deep-OFDM achieves its largest improvements under *pilot-sparse* settings.
- We demonstrate that Deep-OFDM maintains high performance even in the *pilotless* setting, with only modest degradation relative to pilot-assisted configurations; substantially improving spectral efficiency.

2. System Model and Problem Formulation

2.1. System Model

We consider a single-input single-output (SISO) communication system. At the transmitter, the message bit stream $\mathbf{b} \in \{0,1\}^K$, is encoded using a rate $\frac{K}{N}$ LDPC code, producing a codeword $\mathbf{c} \in \{0,1\}^N$. The encoded sequence is mapped to complex symbols using either a standard 2^m-QAM constellation or a learned constellation, resulting in the modulated symbols $\mathbf{s} \in \mathbb{C}^N$.

These symbols are transmitted over a 5G NR-compliant OFDM system with n_S subcarriers and $n_T = 14$ OFDM symbols per slot. The frequency-domain symbols are reshaped into a time-frequency grid $\mathbf{X} \in \mathbb{C}^{n_S \times n_T}$ and transformed to the time domain using an inverse FFT, followed by cyclic prefix (CP) insertion.

We assume a time-varying multipath channel modeled using the clustered delay line (CDL) model (3GPP, 2022). The UE moves at speed u m/s, inducing a maximum Doppler frequency given by $f_d = \frac{uf_c}{c}$, where f_c is the carrier frequency, and c is the speed of light. At the receiver, after CP removal and FFT, the received time-frequency grid $\mathbf{Y} \in \mathbb{C}^{n_S \times n_T}$ can be modeled as:

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{W} \tag{1}$$

where $\mathbf{X} \in \mathbb{C}^{n_S \times n_T}$ is the transmitted symbols in the time-frequency grid, $\mathbf{X} \in \mathbb{C}^{n_S \times n_S}$ is the efective channel matrix, and $\mathbf{W} \in \mathbb{C}^{n_S \times n_T}$ is AWGN noise.

In low-mobility settings, $\mathbf{H}$ is approximately diagonal, but in high-mobility environments, the channel varies within an OFDM symbol duration, breaking the orthogonality between subcarriers. This results in inter-carrier interference (ICI), where $\mathbf{H}$ becomes a full matrix with non-zero off-diagonal elements. The magnitude of these off-diagonal elements increases with mobility, significantly degrading the performance of conventional equalization techniques.

2.2. Baselines

With perfect channel state information (CSI), the Linear Minimum Mean Square Error (LMMSE) equalizer is given by:

$$\mathbf{W}_{\text{LMMSE},t} = \mathbf{H}_t^H \left(\mathbf{H}_t\mathbf{H}_t^H + \sigma_w^2\mathbf{I}\right)^{-1} \tag{2}$$

In practice, the channel $\mathbf{H}_t$ is unknown and must be esti-

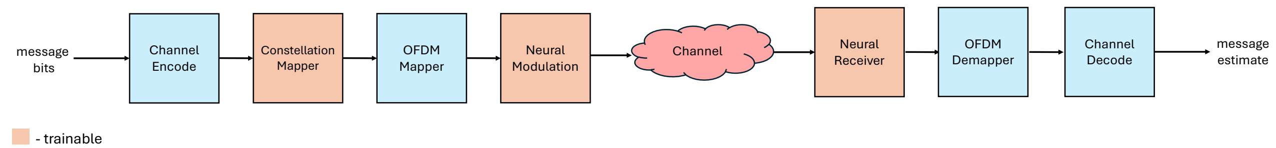

Figure 1: End-to-end optimization of transmitter and receiver.

mated from pilot symbols. The standard least-squares (LS) estimate of the diagonal of $\mathbf{H}_t$ at pilot subcarriers is:

$$\hat{\mathbf{H}}_{t,\mathcal{P}} = \text{diag}\left(\frac{\mathbf{Y}_{t,\mathcal{P}}}{\mathbf{X}_{t,\mathcal{P}}}\right) \tag{3}$$

where $\mathcal{P}_t \subset \{1, \ldots, n_S\}$ denotes the set of pilot subcarrier indices in the t-th OFDM symbol, and $\mathbf{X}_{t,\mathcal{P}}, \mathbf{Y}_{t,\mathcal{P}} \in \mathbb{C}^{|\mathcal{P}_t|}$ are the transmitted and received pilot symbols, respectively.

Channel estimates at non-pilot subcarriers are obtained through interpolation over $\hat{\mathbf{H}}_{t,\mathcal{P}}$.

This approach performs adequately in low-mobility scenarios but suffers from several limitations in high-mobility environments: 1) Violation of the time-invariance assumption within an OFDM symbol, 2) Inability to capture ICI (off-diagonal terms of $\mathbf{H}_t$), 3) Standard interpolation methods cannot accurately capture rapid channel variations.

3. Deep-OFDM and end-to-end optimization

3.1. Neural modulation and receivers

Neural network-based receivers (NRx) have emerged as promising alternatives to traditional signal processing methods (Honkala et al., 2021; Balevi & Andrews, 2019; Aoudia & Hoydis, 2021; Pihlajasalo et al., 2021). These approaches unify channel estimation, interpolation, and demapping into a single trainable model. The NRx takes the received OFDM frame $\mathbf{Y}$, possibly along with pilots, and directly predicts the transmitted bits $\hat{\mathbf{b}}$. Such receivers have demonstrated strong robustness in dynamic channels, particularly in pilot-sparse scenarios. Moreover, recent work has demonstrated the practical feasibility of such systems through the real-time deployment of a 5G standard-compliant NRx, incorporating both hardware and model optimizations (Wiesmayr et al., 2024).

The availability of such trainable receiver modules opens the door for end-to-end system optimization. Specifically, in this work, we focus on mitigating the effects of mobility on OFDM systems by joint optimization of the transmitter and neural receiver, as depicted in Figure 1. Prior work has explored Tx-Rx co-design in this setting, including constellation shaping and super-imposed pilot design (Stark et al., 2019; Aoudia & Hoydis, 2021; Madadi et al., 2023). These methods maintain a one-to-one mapping between transmit symbols and resource elements (REs). In contrast, our method leverages the inductive bias of CNN-based neural receivers (NRx), which process local neighborhoods of subcarriers and time slots. To align with this structure, we introduce Deep-OFDM - a modulation scheme that relaxes the symbol-to-RE mapping by spreading information across the time-frequency grid using a CNN-based modulator. This learned spreading introduces controlled redundancy and enhances time-frequency diversity, significantly improving robustness in high-mobility channels.

3.2. Deep-OFDM Architecture

In conventional OFDM, the $n_S \times n_T$ resource grid is populated with the modulated symbols with one symbol per RE. In contrast, Deep-OFDM uses a CNN encoder to spread the modulated symbols across time and frequency, enabling learned redundancy. For simplicity, we assume all REs are available for data/pilot (i.e., no DC/guard carriers), to isolate the effect of neural modulation. We impose an average power constraint per OFDM frame to ensure parity with conventional systems and restrict our study to the SISO setting.

Neural Receiver. We adopt a CNN-based architecture for the neural receiver (NRx), closely following the design introduced in (Aoudia & Hoydis, 2021). The core building block is a `ResidualBlock`, which contains two convolutional layers. To enhance parameter efficiency without compromising performance, we replace the standard `Conv2D` layers with `SeparableConv2D`, as illustrated in Fig. 2b. The overall architecture comprises an input convolution layer, followed by five residual blocks, and an output layer that directly predicts the log-likelihood ratios (LLRs) of the transmitted bits, as shown in Fig. 3. The input to the NRx consists of the received OFDM frame (i.e., the complex-valued time-frequency grid after FFT and CP removal), along with the positions and values of any known pilot symbols if present. This setup enables the receiver to flexibly operate in pilot-rich, sparse-pilot, or fully pilotless scenarios.

Neural Modulator. To introduce time-frequency diversity at the transmitter, we design the neural modulator as a CNN. While we use a similar architecture as the NRx (Figure 2), we choose a smaller number of convolution

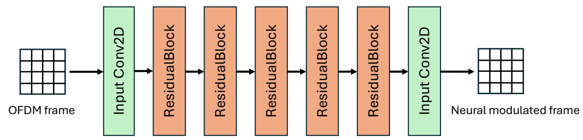

(a) Neural modulation of OFDM frame

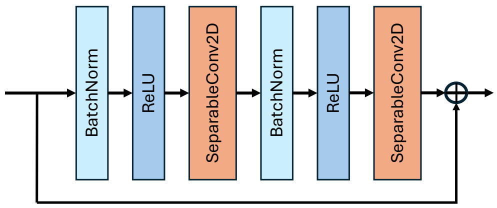

(b) Parameter efficient `ResidualBlock`.

Figure 2: Neural Modulation architecture

channels to minimize the computational burden on the UE during transmission. The modulator takes as input a time-frequency grid populated with complex-valued symbols, where each symbol is generated from a learnable constellation mapping applied to encoded bits. This constellation is optimized jointly with the modulator and receiver as part of the end-to-end training process. The modulator outputs a full OFDM frame that spans all subcarriers and OFDM symbols. The final architecture choices for the neural modulator and neural receiver are shown in Tab. 1. Empirically, we observed that reducing the number of parameters in the neural modulator had a negligible impact on performance, suggesting that the encoder's role in introducing time-frequency diversity can be fulfilled with a lightweight architecture. We leave a more detailed analysis to future work.

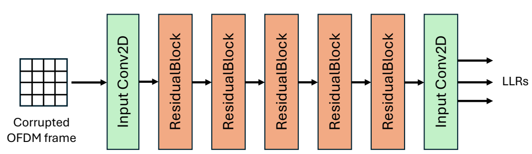

Figure 3: Neural receiver to predict LLRs.

Layer	Channels (Modulator)	Channels (Receiver)	Kernel size	Dilation rate
Input Conv2D	36	48	(3,3)	(1,1)
ResNet block 1	72	96	(7,7)	(7,2)
ResNet block 2	72	96	(7,5)	(7,1)
ResNet block 3	72	96	(5,3)	(1,2)
ResNet block 4	72	96	(3,3)	(1,1)
ResNet block 5	72	96	(3,3)	(1,1)
Output Conv2D	2	$2m$	(1,1)	(1,1)

Table 1: Architecture details for neural modulator and receiver.

Training. We simultaneously optimize the neural modulator and neural receiver blocks by optimizing the end-to-end binary cross entropy (BCE) loss, using the rate as

$$\mathcal{L}_{\text{rate}} = -\left(1 - \frac{\mathcal{L}_{\text{BCE}}}{\ln 2}\right) = \frac{\mathcal{L}_{\text{BCE}}}{\ln 2} - 1. \tag{4}$$

Parameter	**Symbol (if any)**	**Value**
Number of OFDM symbols	n_T	14 (1 slot)
Number of subcarriers	n_S	128
Carrier frequency	–	2.6 GHz
Subcarrier spacing	–	15 kHz
Cyclic prefix duration	n_{CP}	6 symbols
Channel model	–	CDL-C
Learning rate	–	10^{-3}
Batch size for training	S	128
Modulation order	m	6
UE speed range (training)	–	0 m/s to 40 m/s

Table 2: Training parameters

As will be evident from results in Sec. 4, our approach is particularly applicable in scenarios where pilot-based channel estimation is ineffective, such as in very high mobility.

4. Results

In this section, we evaluate the reliability and robustness of the neural modulator scheme introduced in Section 3.

Evaluation setup. We conduct extensive experiments to evaluate our proposed neural modulation approach against multiple baselines. Our simulation environment follows the system model introduced in Section 3. Key training parameters are outlined in Table 2. To reduce training complexity, we omit LDPC encoding/decoding during training and reintroduce a rate-1/2 LDPC code (blocklength equal to the total number of resource elements in the OFDM grid) only at inference. All models are trained on the CDL-C channel model, with user mobility uniformly sampled from 0 to 40 m/s, covering both stationary and vehicular scenarios. Simulations are implemented using the GPU-accelerated Sionna framework (Hoydis et al., 2022). To ensure a fair comparison, we allocate an identical total parameter budget to both our end-to-end neural link and the QAM + neural receiver (NRx) baseline. In this way, any performance differences arise from how complexity is apportioned between transmitter and receiver, rather than from an overall parameter increase. We evaluate performance at three mobility regimes - low speeds (10m/s = 22mph), medium speeds (40m/s = 88mph), and high speeds (60m/s = 132mph). Notably, 60m/s lies outside the training distribution and is used to assess the generalization capability of the proposed scheme.

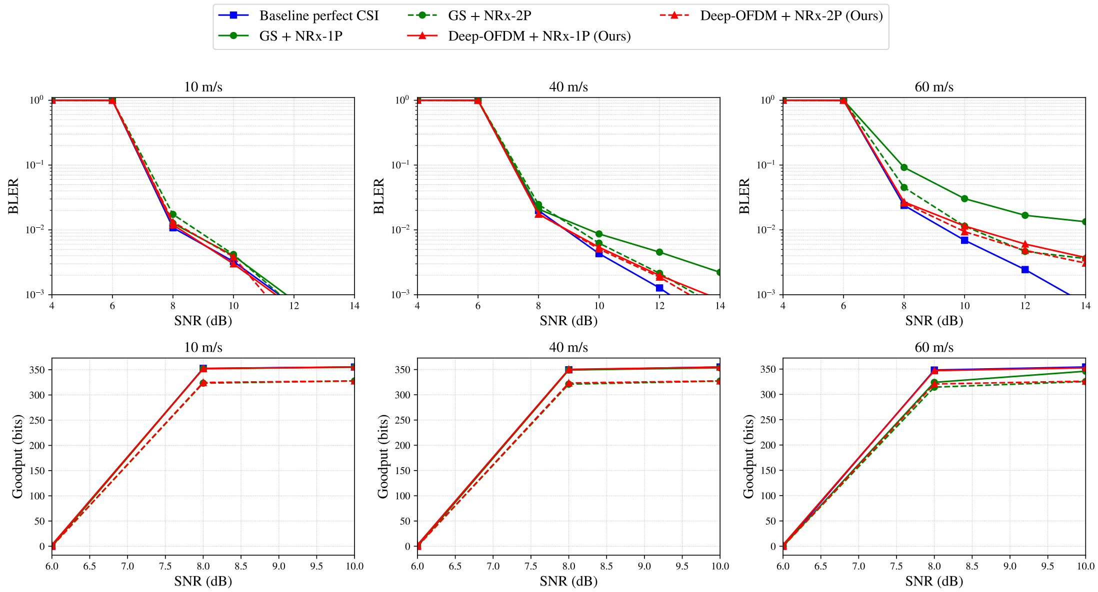

Figure 4: NTF-OTFS exhibits gains in BLER compared to NRx with trainable constellations (GS). The gains are more pronounced in the single-pilot case than with two-pilots

4.1. Pilot-sparse: Deep-OFDM excels at high mobility.

Under high mobility, the wireless channel varies significantly between consecutive symbols, necessitating a large number of pilot symbols for reliable channel estimation. However, increasing pilot density reduces spectral efficiency, presenting a fundamental tradeoff. This motivates the need for designs that perform reliably under limited pilot overhead.

To study this regime, we evaluate system performance when only a single pilot symbol is inserted at the third OFDM symbol (1P). At high mobility, the LMMSE receiver with least-squares (LS) channel estimation performs poorly due to channel aging, where the estimated channel quickly becomes outdated. In contrast, the neural receiver (NRx) substantially outperforms the LMMSE+LS baseline by better accommodating the time-varying nature of the channel. However, even the neural receiver alone begins to degrade at higher Doppler frequencies. Prior work has shown that geometric shaping (GS), i.e., learning the constellation mapping, can yield performance gains in both AWGN and fading channels (Stark et al., 2019; Aoudia & Hoydis, 2021; Madadi et al., 2023). Consistent with this, Figure 4 shows that training a learnable constellation provides gains over fixed QAM when used with NRx.

Beyond geometric shaping, Deep-OFDM introduces additional degrees of freedom across both time and frequency dimensions, enabling the system to better cope with mobility-induced distortions. While the performance gains at low Doppler shifts (e.g., 10 m/s) are modest, the advantages of neural modulation become increasingly pronounced as Doppler frequencies rise. For instance, at higher speeds (e.g., 60 m/s), neural modulation yields a substantial improvement in both BLER and spectral efficiency (Figure 4). Notably, the performance gains are more significant in the single-pilot configuration compared to the two-pilot case.

We quantify system performance using the achieved goodput, defined as

$$\text{Goodput} = R \cdot \rho \cdot (1 - \text{BLER}), \tag{5}$$

where R denotes the transmission rate in bits per channel use, and ρ is the fraction of symbols allocated to data (excluding pilots).

The difficulty of reliable channel estimation at high Doppler is underscored by the observation that the standalone NRx achieves nearly the same goodput at 60 m/s with one pilot (1P) as with two pilots (2P). This stagnation indicates that simply increasing pilot density does not sufficiently mitigate mobility impairments. In contrast, our end-to-end neural link achieves significantly higher goodput than both NRx-1P and NRx-2P, despite maintaining the same pilot overhead as NRx-1P. These results highlight the efficacy of transmitter-side adaptation in challenging channel conditions.

Finally, despite operating under the same parameter budget, Deep-OFDM consistently outperforms the QAM + NRx baseline across all mobility levels. This highlights a key architectural insight: jointly optimizing transmitter and receiver yields better robustness than concentrating model capacity entirely at the receiver. While prior work has

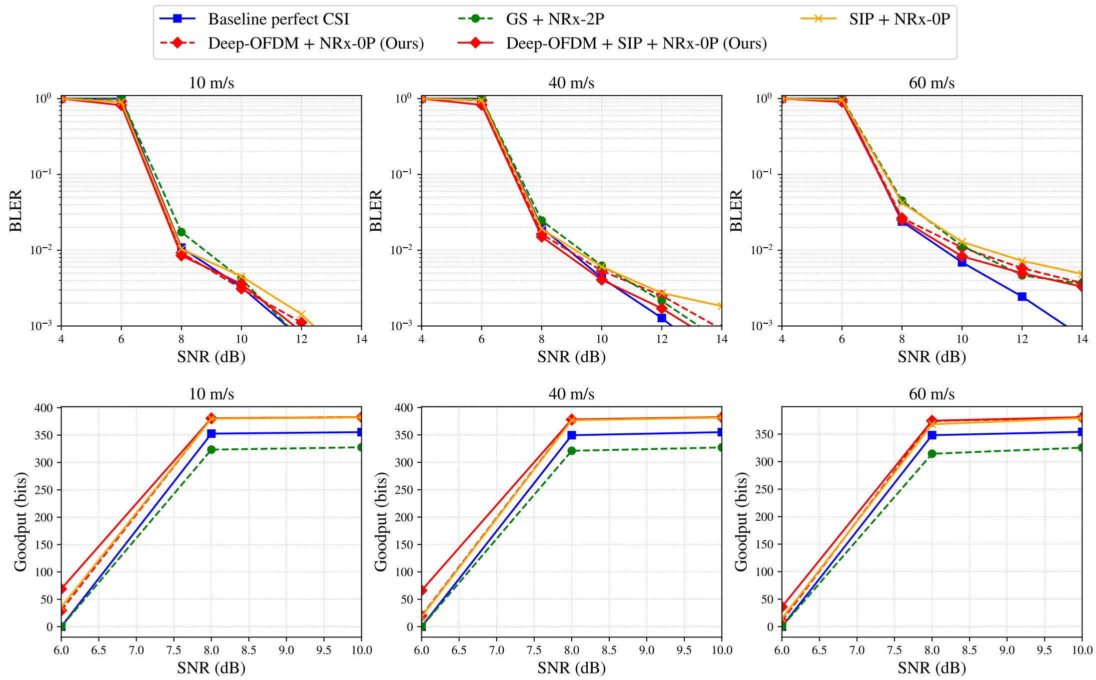

Figure 5: Pilotless Deep-OFDM outperforms NRx baseline and explicit SIP; Gains attributed to both time-frequency mixing and implicitly learned SIP

shown that large neural receivers alone can achieve robustness (Aoudia & Hoydis, 2021), our design achieves similar or better performance with far fewer parameters.

4.2. Deep-OFDM enables pilotless communication.

We now evaluate the performance of neural modulation in the absence of any pilot signalling. In this challenging setting, traditional methods for coherent communication fail, as they require reference signals for channel estimation. Similarly, conventional OFDM with neural receiver performs poorly across all mobility conditions due to the lack of explicit channel information.

A prior approach to this problem uses superimposed pilots (SIP) (Aoudia & Hoydis, 2021), which overlays pilot signals onto data-carrying symbols. Here, a BPSK-modulated pilot matrix $P \in \pm 1^{N \times M}$ is defined; and an energy allocation matrix A is learned. $\sqrt{1-A} \cdot X + \sqrt{A} \cdot P$ is transmitted. Unlike orthogonal pilots, which reserve dedicated REs for pilots, SIP overlays pilot and data within every RE, forcing the receiver to jointly separate and decode them.

We compare pilotless Deep-OFDM against SIP and LMMSE with perfect CSI. Within the training range (0–40 m/s), SIP and pilotless modulation deliver similar reliability and goodput. However, it is more robust to unseen channel conditions - when the relative speed between UE and the base station is 60m/s plot, neural modulation outperforms SIP. When SIP is combined with Deep-OFDM, we observe further performance enhancements, highlighting the complementary benefits of time-frequency diversity.

Remarkably, despite using no pilots, pilotless Deep-OFDM matches the performance of a GS+NRx system with two pilots: gaining 14% in goodput. In addition to performance, Deep-OFDM also offers architectural flexibility. Owing to its convolutional structure, a model trained on one OFDM frame size naturally generalizes to configurations with different numbers of subcarriers. In contrast, SIP requires a separately trained pilot pattern for each OFDM frame configuration

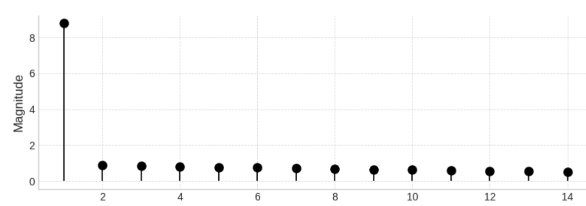

Figure 6: SVD of the neural modulator's output mean reveals a strong rank-one component, suggesting the emergence of a learned implicit pilot.

Pilotless neural modulation learns implicit *pilots*. To gain insight into the *structure* that the neural modulator learns, we analyze the statistics of its transmitted vectors $\mathbf{X}$. First, we compute the empirical mean $\boldsymbol{\mu} \triangleq \mathbb{E}[\mathbf{X}]$,

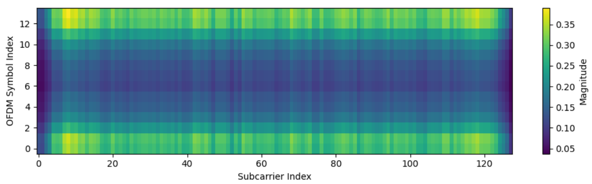

Figure 7: Implicitly learned pilot pattern

and perform a Singular Value Decomposition (SVD)

$$\boldsymbol{\mu} \;=\; \mathbf{U}\,\boldsymbol{\Sigma}\,\mathbf{V}^{\mathsf{H}}.$$

Figure 6 shows that the mean has dominant rank–one structure ($\sigma_1 \gg \sigma_2$), suggesting the presence of a determinstic reference. Reconstructing only the first singular component $\mathbf{p} \triangleq \sigma_1 \mathbf{u}_1 \mathbf{v}_1^{\mathsf{H}}$ into the time–frequency grid (Fig. 7) reveals a deterministic, constant-phase pattern that is superimposed on *every* transmit block.

Consequently we can decompose the modulator output :

$$\mathbf{X} \;=\; \underbrace{\mathbf{P}}_{\text{learned "pilot"}} \;+\; \underbrace{\widetilde{\mathbf{X}}}_{\substack{\text{data}\\ \text{(zero mean)}}}, \qquad \mathbb{E}[\widetilde{\mathbf{X}}] = \mathbf{0}. \tag{6}$$

Equation (6) is formally identical to the *super-imposed pilot* (SIP) architecture studied in (Cui & Tellambura, 2005). Unlike that work, where the pilot vector P is predetermined and only the data symbols $\tilde{X}$ vary, our model learns *both* P and $\tilde{X}$ jointly via gradient descent, without imposing any explicit pilot constraint. The implicitly learned pilots P are visualized in Figure 7.

5. Conclusion and Remarks

In this work, we introduced Deep-OFDM, a learnable modulation scheme that augments OFDM by spreading transmit symbols across time and frequency, tailored to the inductive biases of neural receivers. By jointly optimizing a low-complexity neural modulator and a CNN-based receiver, the system achieves strong robustness to high-mobility fading, particularly in sparse or even zero-pilot regimes. Beyond performance improvements, our analysis shows that Deep-OFDM learns to embed implicit pilot structures, enabling reliable pilotless communication through end-to-end training. Future work includes interpreting the learned time–frequency mixing patterns, analyzing the kernels of the neural modulator, and developing analytical approximations to better understand the learned behavior. Further reducing transmitter-side complexity, extending the approach to MIMO settings, and studying generalization across diverse channel distributions are promising directions toward practical deployment.

References

3GPP. Study on Channel Model for Frequencies from 0.5 to 100 GHz. Technical Report TR 38.901, 3rd Generation Partnership Project (3GPP), 2022. URL https://www.3gpp.org/DynaReport/38901.htm. Technical Specification Group Radio Access Network.

Ai, B., Cheng, X., Kürner, T., Zhong, Z.-D., Guan, K., He, R.-S., Xiong, L., Matolak, D. W., Michelson, D. G., and Briso-Rodriguez, C. Challenges toward wireless communications for high-speed railway. *IEEE transactions on intelligent transportation systems*, 15(5):2143–2158, 2014.

Ankireddy, S. K. and Kim, H. Interpreting Neural Min-Sum Decoders. In *ICC 2023-IEEE International Conference on Communications*, pp. 6645–6651. IEEE, 2023.

Aoudia, F. A. and Hoydis, J. End-to-end learning for ofdm: From neural receivers to pilotless communication. *IEEE Transactions on Wireless Communications*, 21(2):1049–1063, 2021.

Balevi, E. and Andrews, J. G. One-bit ofdm receivers via deep learning. *IEEE Transactions on Communications*, 67(6):4326–4336, 2019.

Chafii, M., Bariah, L., Muhaidat, S., and Debbah, M. Twelve scientific challenges for 6g: Rethinking the foundations of communications theory. *IEEE Communications Surveys & Tutorials*, 25(2):868–904, 2023.

Cui, T. and Tellambura, C. Superimposed pilot symbols for channel estimation in ofdm systems. In *Proc. IEEE Globecom*, volume 5, pp. 2229–2233, 2005.

Hadani, R., Rakib, S., Tsatsanis, M., Monk, A., Goldsmith, A. J., Molisch, A. F., and Calderbank, R. Orthogonal time frequency space modulation. In *2017 IEEE wireless communications and networking conference (WCNC)*, pp. 1–6. IEEE, 2017.

Hebbar, S. A., Mishra, R. K., Ankireddy, S. K., Makkuva, A. V., Kim, H., and Viswanath, P. TinyTurbo: Efficient Turbo Decoders on Edge. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pp. 2797–2802. IEEE, 2022.

Hebbar, S. A., Ankireddy, S. K., Kim, H., Oh, S., and Viswanath, P. DeepPolar: Inventing non-linear large-kernel Polar Codes via Deep Learning. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Honkala, M., Korpi, D., and Huttunen, J. M. Deeprx: Fully convolutional deep learning receiver. *IEEE Transactions on Wireless Communications*, 20(6):3925–3940, 2021.

Hoydis, J., Cammerer, S., Aoudia, F. A., Vem, A., Binder, N., Marcus, G., and Keller, A. Sionna: An open-source library for next-generation physical layer research. *arXiv preprint arXiv:2203.11854*, 2022.

Jamali, M. V., Saber, H., Hatami, H., and Bae, J. H. Productae: Toward training larger channel codes based on neural product codes. In *ICC 2022-IEEE International Conference on Communications*, pp. 3898–3903. IEEE, 2022.

Jiang, Y., Kim, H., Asnani, H., Kannan, S., Oh, S., and Viswanath, P. Turbo autoencoder: Deep learning based channel codes for point-to-point communication channels. *Advances in neural information processing systems*, 32, 2019.

Kim, H., Jiang, Y., Kannan, S., Oh, S., and Viswanath, P. Deepcode: Feedback codes via deep learning. *Advances in neural information processing systems*, 31, 2018.

Li, J., Li, B., and Lu, Y. Neural video compression with diverse contexts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22616–22626, 2023.

Madadi, P., Cho, J., Zhang, C. J., and Burghal, D. Ai/ml optimized high-order modulations. In *ICC 2023-IEEE International Conference on Communications*, pp. 6373–6378. IEEE, 2023.

Makkuva, A. V., Liu, X., Jamali, M. V., Mahdavifar, H., Oh, S., and Viswanath, P. Ko codes: inventing nonlinear encoding and decoding for reliable wireless communication via deep-learning. In *International Conference on Machine Learning*, pp. 7368–7378. PMLR, 2021.

Pihlajasalo, J., Korpi, D., Honkala, M., Huttunen, J. M., Riihonen, T., Talvitie, J., Uusitalo, M. A., and Valkama, M. Deep learning based ofdm physical-layer receiver for extreme mobility. In *2021 55th Asilomar Conference on Signals, Systems, and Computers*, pp. 395–399. IEEE, 2021.

Stark, M., Ait Aoudia, F., and Hoydis, J. Joint learning of geometric and probabilistic constellation shaping. In *2019 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6. IEEE, 2019.

Tong, W. and Zhu, P. 6g: The next horizon. In *TITLES*, pp. 54, 2022.

Wiesmayr, R., Cammerer, S., Aoudia, F. A., Hoydis, J., Zakrzewski, J., and Keller, A. Design of a standard-compliant real-time neural receiver for 5g nr. *arXiv preprint arXiv:2409.02912*, 2024.

Ye, H., Li, G. Y., and Juang, B.-H. Power of deep learning for channel estimation and signal detection in ofdm systems. *IEEE Wireless Communications Letters*, 7(1): 114–117, 2017.