
End-to-End Waveform Design for Nonlinear Satellite Links with a Convolutional Autoencoder

Yara Huleihel
Ben-Gurion University
Beer-Sheva, Israel
halihal@post.bgu.ac.il

Rom Hirsch
Ben-Gurion University
Beer-Sheva, Israel
romhi@post.bgu.ac.il

Haim H. Permuter
Ben-Gurion University
Beer-Sheva, Israel
haimp@bgu.ac.il

Abstract

We present a convolutional autoencoder (CAE) for single-carrier (SC) waveform design in satellite links with nonlinear high-power amplifiers (HPAs). The CAE jointly optimizes bit-error rate (BER), peak-to-average power ratio (PAPR), and adjacent-channel power ratio (ACPR) via an augmented-Lagrangian objective. Using measured block-upconverter (BUC) data with AM/AM and AM/PM distortions across temperature and frequency variations, we show consistent improvements over conventional PAPR-reduction techniques and baselines, with strong generalization across frequency and temperature variations. Experiments with 64-QAM demonstrate BER and spectral advantages at practical back-off levels without explicit digital predistortion (DPD).

1 Introduction

High-throughput satellite systems employ high-order QAM to maximize spectral efficiency. In practice, single-carrier (SC) waveforms are attractive for satellite links due to their reduced sensitivity to frequency offsets and phase noise, but they remain vulnerable to the nonlinearities of high-power amplifiers (HPAs). High peak-to-average power ratio (PAPR) increases the impact of nonlinear distortion, causing constellation warping, adjacent-channel interference, and degraded bit-error rate (BER). These challenges are especially pronounced in Ka-band block-upconverters (BUCs), where amplifier responses vary significantly with temperature, frequency, and hardware aging.

Classical mitigation approaches include clipping-and-filtering (CF) [3, 10], selective mapping (SLM) [10], and probabilistic shaping [7]. While effective in reducing PAPR, these techniques often increase BER, require side information, or incur substantial complexity. Digital predistortion (DPD) [9, 12, 4] remains the industry standard, but it is difficult to sustain accuracy near saturation and across varying amplifier conditions.

Building on these limitations, recent work has explored deep learning (DL) for physical layer design. Early applications demonstrated DL-based receivers, precoders, and OFDM waveform shaping [2, 11, 8]. More recently, convolutional autoencoders (CAEs) proved effective for PAPR reduction in terrestrial OFDM [6], motivating their adaptation to SC satellite links. Unlike [12, 4], which rely on explicit predistorter identification, the proposed framework learns waveform shaping directly from measured HPA data, capturing realistic AM/AM and AM/PM effects without requiring closed-form amplifier models.

This work proposes a CAE that jointly optimizes BER, spectral containment, and PAPR by embedding these requirements into an augmented-Lagrangian loss. Trained directly on measured amplifier responses, the CAE achieves robustness to hardware variability. The architecture is inspired by the waveform learning framework in [1], and after exploring several alternatives, this design offered the best balance of accuracy, complexity, and runtime.

Main contributions: The proposed study advances end-to-end learning for satellite waveform design by addressing both spectral and nonlinear challenges inherent in SC systems. The key contributions are:

- An end-to-end CAE for SC waveforms that learns shaping directly under realistic AM/AM and AM/PM distortions.
- A gradual loss training strategy to enforce spectral and power constraints during optimization.
- Evaluation on measured BUC data across diverse operating points, showing robust generalization.
- Demonstrated consistent gains in BER, PAPR, and PSD compared with CF, SLM, DPD, and a no-mitigation baseline on 64-QAM.

2 System Model

We consider an SC communication system transmitting blocks of N symbols with oversampling factor L and root-raised-cosine (RRC) pulse shaping. Let $\mathbf{B} \in \{0, 1\}^{N \times K}$ denote the input bits per block and $\mathbf{s} \in \mathbb{C}^N$ the mapped 64-QAM symbols. After upsampling, the shaped waveform is

$$x[t] = \sum_{m=0}^{L_f-1} s_{\uparrow L}[t-m] g_{\text{tx}}[m], \quad t = 0, \dots, NL + L_f - 2, \quad (1)$$

where $g_{\text{tx}}[m]$ is the RRC filter of length L_f .

The nonlinear HPA applies a memoryless transformation

$$x_t^{\text{P}} = G(x_t), \quad (2)$$

capturing both AM/AM and AM/PM distortion. The channel then adds additive white Gaussian noise (AWGN):

$$y_t = x_t^{\text{P}} + w_t, \quad w_t \sim \mathcal{CN}(0, \sigma_w^2). \quad (3)$$

At the receiver, matched filtering and downsampling yield an N -symbol sequence for detection. Fig. 1 shows the system architecture, including modulation, pulse shaping, nonlinear amplification, channel noise, and receiver processing. We evaluate performance using several metrics:

BER. BER is measured after symbol detection and plotted versus the peak signal-to-noise ratio (PSNR), defined for a normalized channel as

$$\text{PSNR} = \frac{P_T}{\sigma_w^2}, \quad (4)$$

where P_T is the maximum transmit power.

PAPR. For a length- N sequence $x[n]$,

$$\text{PAPR} = \frac{\max_n |x[n]|^2}{\mathbb{E}[|x[n]|^2]}. \quad (5)$$

ACPR. ACPR quantifies spectral regrowth due to nonlinearities. From the output PSD $P_{ss}(f)$,

$$\text{ACPR} = \frac{\int_{\mathcal{F}_{\text{adj}}} P_{ss}(f) df}{\int_{\mathcal{F}_{\text{in}}} P_{ss}(f) df}, \quad (6)$$

with $\mathcal{F}_{\text{in}}$ the in-band region and $\mathcal{F}_{\text{adj}}$ the adjacent-band region.

Output Back-Off (OBO). The amplifier operating point is expressed as

$$\text{OBO [dB]} = 10 \log_{10} \left(\frac{P_{\text{sat}}}{P_{\text{out}}} \right), \quad (7)$$

where P_{sat} is the saturated output power and P_{out} is the average output power. Lower OBO means operation closer to saturation, improving efficiency but exacerbating nonlinear distortion. Together, these metrics capture the trade-off between detection reliability (BER), power efficiency (PAPR, OBO), and spectral compliance (ACPR) under nonlinear HPA conditions.




Figure 1: General SC system model with nonlinear HPA and receiver detection.

3 CAE Architecture and Training

The CAE includes a convolutional encoder $f(\cdot)$ and decoder $g(\cdot)$, modeling the transmitter and receiver, respectively. Both are built with dilated ResNet blocks. Complex I/Q samples are represented as real-valued channels (C2R at input, R2C where needed). A power-normalization layer enforces unit average transmit power.

Design rationale. Both encoder and decoder employ ResNet blocks with separable convolutions, batch normalization, ReLU activation, dilation, and skip connections [5], which mitigate vanishing gradients and enable efficient identity mapping. Dilated convolutions expand the receptive field of each layer, providing efficient representation of nonlinear distortion patterns across the waveform block without requiring very deep stacks, while residual connections stabilize training under strong nonlinearities. Alternative designs based on fully connected or attention layers consistently underperformed in BER, runtime, and scalability. Figure 2b details the encoder architecture, while the decoder mirrors this design but omits the power normalization and R2C layers.

- **Encoder (Tx):** The encoder processes symbol sequences through stacked dilated ResNet blocks, producing complex-valued waveforms using C2R split channels, with a final power-normalization stage enforcing unit transmit power.
- **Decoder (Rx):** A symmetric ResNet-based convolutional structure reconstructs LLRs from received waveforms, without the power normalization or R2C stages.

Training strategy. The HPA response, measured across frequencies and temperatures on Ka-band BUC hardware, is spline-interpolated into a high-resolution LUT model. Training begins with a warm-up phase minimizing only $\mathcal{L}_{\text{BCE}}$, which prevents premature domination by spectral or PAPR penalties. Constraints are then gradually enabled, allowing the model to balance detection accuracy, PAPR, and ACPR objectives. Early stopping on a validation set avoids overfitting and ensures generalization across amplifier realizations.

Loss function. Optimization employs an augmented-Lagrangian composite loss:

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \lambda_2 \mathcal{L}_{\text{PAPR}} + \frac{\rho_2}{2} \|\mathcal{L}_{\text{PAPR}}\|_2^2 + \frac{1}{2\rho_3} (\max\{0, \lambda_3 + \rho_3 \mathcal{L}_{\text{ACPR}}\})^2 - \frac{\lambda_3^2}{2\rho_3}, \quad (8)$$


with components:

- $\mathcal{L}_{\text{BCE}}$: binary cross-entropy reconstruction loss between predicted and true bit labels,
- $\mathcal{L}_{\text{PAPR}} = \text{PAPR}\{x_n^E\} - \text{PAPR}_{\text{req}}$: constraint on encoder output relative to a predefined threshold,
- $\mathcal{L}_{\text{ACPR}} = \text{ACPR}\{x_n^P\} - \text{ACPR}_{\text{req}}$: constraint on spectral leakage at the HPA output with respect to adjacent-channel mask limits.


The augmented multipliers (λ_i, ρ_i) adapt dynamically, enabling stable multi-objective optimization. Training begins with $\mathcal{L}_{\text{BCE}}$ only and gradually activates the PAPR/ACPR penalties. This staged procedure proved essential to prevent collapse, ensure convergence across diverse HPA profiles, and maintain computational efficiency for real-time satellite deployment.

4 Results and Insights


Setup. Experiments were conducted on SC blocks of $N=1000$ symbols with oversampling factor $L=4$ and RRC pulse shaping. Independent training and test sets were generated with randomized



(a) Conv-AE overall scheme.



(b) Waveform design - Encoder scheme $f(x)$.



(c) ResNet block structure.

Figure 2: Structure of the proposed CAE system. (a) End-to-end scheme. (b) Encoder block. (c) ResNet block.

symbol streams, AWGN realizations, and HPA instances to capture variability across hardware, temperature, and frequency.


Baselines. We compare the proposed CAE against classical approaches: CF, SLM employing $U=64$ candidate phase sequences, DPD, a no-mitigation baseline (BL), and BL-IDEAL (a distortion-free reference). All classical receivers use maximum-likelihood or LLR demapping as appropriate.

Measured HPA model. The BUC originates from a Gilat Ka-band transceiver. AM/AM and AM/PM responses were measured across -40°C – 60°C and 29–30 GHz using a 67 GHz vector network analyzer in mixer mode. The 1 dB compression point is 2.5 W, and a spline-interpolated LUT captures fine nonlinearities. To isolate effects, we also consider a reduced AM/AM-only model (no phase distortion).


BER and Generalization. Figure 3a shows BER for 64-QAM under mixed HPA training. Solid lines denote mean performance, shaded areas indicate variability across amplifier realizations, outperforming CF and SLM. CAE matches the noise-limited floor at low PSNR, then achieves substantially lower BER as nonlinear distortion dominates. Narrow confidence bands confirm robustness across diverse HPAs without per-device tuning, an essential property for satellite deployments where amplifier responses drift with temperature, frequency, and aging.

PAPR and Spectral Containment. Figures 3b and 4a report PAPR CCDF and PSD. CAE shifts the CCDF left, lowering PAPR by over 1 dB relative to CF and SLM. Spectral regrowth is also suppressed, with PSD comfortably within mask limits. Figure 4b highlights the OBO–ACPR trade-off: CAE consistently reaches a given ACPR target at lower output back-off (OBO), confirming that it preserves spectral purity while operating closer to saturation, reflecting better efficiency–purity balance. The flatter slope stems from training at zero input back-off (IBO), unlike classical methods that rely on power back-off.

Comparison with DPD. We further benchmark against interpolation-based DPD, which constructs a spline-fitted LUT inverse from measured AM/AM and AM/PM curves. DPD effectively compensates AM–PM by adjusting input phase, yielding competitive BER in moderate regimes, but struggles near saturation where clipping and compression dominate. Figure 5 show that CAE consistently outperforms DPD. Moreover, CAE inherently reduces PAPR, whereas DPD leaves peak excursions largely unchanged. Importantly, DPD requires amplifier-specific calibration, limiting



(a) BER vs. PSNR. Shaded: HPA instances variance. (b) PAPR CCDF for 64-QAM.
Figure 3: 64-QAM performance with mixed HPA training.



(a) PSD of considered methods. (b) OBO-ACPR trade-off.
Figure 4: 64-QAM spectral performance with mixed HPA training.




Figure 5: BER vs. $PSNR$ comparison: CAE vs. DPD, 64-QAM.

scalability in hardware-diverse systems. By contrast, CAE generalizes across amplifiers after a single training phase, offering robustness without repeated recalibration.

Summary. CAE achieves consistent BER, PAPR, and ACPR gains across nonlinear regimes, generalizes across amplifiers, and surpasses both heuristic PAPR methods and classical DPD. These results underline the feasibility of deploying NN-based waveform design in realistic satellite communication systems, offering robustness and efficiency simultaneously.

5 Conclusion

This work presented a CAE that jointly optimizes BER, PAPR, and ACPR for SC waveform design under nonlinear HPAs. Trained directly on measured BUC responses, the CAE consistently outperformed CF, SLM, and interpolation-based DPD across BER, spectral, and efficiency metrics, while generalizing across frequency, temperature, and hardware variations. Unlike DPD, which requires amplifier-specific calibration, the CAE scales across devices after a single training phase, reducing deployment cost and complexity.

Beyond demonstrating technical feasibility, these results highlight the broader potential of learning-based waveform design to complement or replace model-driven compensation in satellite systems. Future work will extend the framework to coded operation with advanced FEC, integration into non-terrestrial network (NTN) scenarios, and real-time implementation on DSP/FPGA platforms.

References

- [1] F. A. Aoudia and J. Hoydis. “Waveform learning for next-generation wireless communication systems”. In: *IEEE Transactions on Communications* (2022).
- [2] Mathieu Goutay et al. “Learning OFDM waveforms with PAPR and ACLR constraints”. In: *arXiv preprint arXiv:2110.10987* (2021).
- [3] I. Gutman, I. Iofedov, and D. Wulich. “Iterative Decoding of Iterative Clipped and Filtered OFDM Signal”. In: *IEEE Transactions on Communications* 61.10 (2013), pp. 4284–4293.
- [4] Muhammad Furqan Haider et al. “Predistortion-based linearization for 5G and beyond millimeter-wave transceiver systems: A comprehensive survey”. In: *IEEE Communications Surveys & Tutorials* 24.4 (2022), pp. 2029–2072.
- [5] Kaiming He et al. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.
- [6] Yara Huleihel and Haim H Permuter. “Low papr mimo-ofdm design based on convolutional autoencoder”. In: *IEEE Transactions on Communications* 72.5 (2024), pp. 2779–2792.
- [7] Or Levi, Dan Raphaeli, and Yonathan Tate. “A novel shaping scheme for PAPR reduction in single-carrier modulation”. In: *IEEE Transactions on Communications* 66.9 (2018), pp. 4222–4233.
- [8] Zhijun Liu et al. “A Joint PAPR Reduction and Digital Predistortion Based on Real-Valued Neural Networks for OFDM Systems”. In: *IEEE Transactions on Broadcasting* (2021).
- [9] Albert Molina, Kannan Rajamani, and Kamran Azadet. “Digital predistortion using lookup tables with linear interpolation and extrapolation: Direct least squares coefficient adaptation”. In: *IEEE Transactions on Microwave Theory and Techniques* 65.3 (2016), pp. 980–987.
- [10] A. P. More and S. B. Somani. “The reduction of PAPR in OFDM systems using clipping and SLM method”. In: *2013 International Conference on Information Communication and Embedded Systems (ICICES)*. 2013, pp. 593–597.
- [11] I. Sohn and S. C. Kim. “Neural Network Based Simplified Clipping and Filtering Technique for PAPR Reduction of OFDM Signals”. In: *IEEE Communications Letters* 19.8 (2015), pp. 1438–1441.
- [12] Ovais Bin Usman and Andreas Knopp. “Digital predistortion in high throughput satellites: Architectures and performance”. In: *IEEE Access* 9 (2021), pp. 42291–42304.