Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

ML-Enabled Deformable Matched Filters for Bandlimitation Compensation in Free-Space Optics

Paul Anthony Haigh

p.a.haigh@qmul.ac.uk

Queen Mary University of London

Article

Keywords: Carrier-Less Amplitude and Phase Modulation, Filters, Free-Space Optics, Machine Learning, Neural Networks

Posted Date: February 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8852763/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

ML-Enabled Deformable Matched Filters for Bandlimitation Compensation in Free-Space Optics

Paul Anthony Haigh

***Abstract*—This paper proposes a neural-network-assisted deformable matched filtering framework for carrier-less amplitude and phase (CAP) modulation operating under bandwidth-limited channel conditions. Instead of replacing the analytically derived CAP matched filter, the proposed receiver learns a residual deformation of the nominal matched filter based on a compact set of physically motivated signal features extracted from the received waveform. A total of 16 time-domain, frequency-domain, and memory-related features are used to provide a low-dimensional representation of bandwidth-induced pulse distortion. These features are mapped by a fully connected neural network to complex-valued matched filter coefficients, enabling adaptive pulse-shape compensation prior to symbol-rate sampling. The network is trained end-to-end using a differentiable loss function based on error vector magnitude (EVM). Experimental results obtained using a hardware-in-the-loop CAP transmission system demonstrate that the proposed deformable matched filter significantly outperforms conventional fixed matched filtering under severe bandwidth constraints, without requiring decision feedback or increasing receiver latency.**

***Index Terms*—Carrier-Less Amplitude and Phase Modulation, Filters, Free-Space Optics, Machine Learning, Neural Networks**

I. Introduction

CARRIER-LESS amplitude and phase (CAP) modulation has emerged as an attractive solution for bandwidth-constrained optical communication systems due to its spectral efficiency and reduced implementation complexity compared to orthogonal frequency-division multiplexing [1]. CAP encodes multi-level symbols using in-phase and quadrature pulse-shaping filters without explicit carrier modulation, enabling high-order modulation formats within limited bandwidths. However, the performance of CAP systems is highly sensitive to pulse-shape distortions induced by bandwidth limitations, atmospheric turbulence, and hardware impairments such as photodetector saturation [2]. Maintaining accurate matched filtering under time-varying and bandwidth-constrained conditions remains a central challenge for practical CAP-based optical links [3].

Traditional CAP receivers employ fixed matched filters designed under ideal channel assumptions. In practice, bandwidth limitations and channel impairments cause rapid variations in received signal statistics, resulting in mismatches between nominal filters and the optimal receiver structure. Adaptive equalisation techniques [4], [5] have been proposed to mitigate these impairments, but they often rely on training sequences that consume bandwidth and require frequent retraining, or operate at the symbol level where they cannot directly compensate for pulse-shape distortions [6], [7].

Recent advances in machine learning (ML) have demonstrated remarkable capability in learning complex non-linear mappings [8], with neural networks applied to channel equalisation often outperforming conventional methods [6]. However, most ML-based optical receivers operate at symbol or post-detection level, limiting their ability to directly address pulse-shape distortions. We propose a deformable matched filter framework where a neural network dynamically estimates optimal pulse-shaping filters from compact feature representations extracted from received waveforms, enabling compensation prior to symbol-rate processing without symbol exposure or decision feedback.

Unlike conventional approaches that replace analytically derived receiver structures with black-box models, our method learns residual corrections to theoretically optimal matched filters based on 16 physically motivated signal features. The neural network maps these features to complex-valued filter coefficients, adapting the matched filter to compensate for bandwidth-induced pulse distortions while maintaining the classical receiver structure. This hybrid strategy combines the reliability and interpretability of classical communication theory with the adaptability of data-driven optimisation.

The integration of machine learning into physical layer communications has gained significant momentum in recent years [9]. Deep learning approaches have demonstrated remarkable capabilities in end-to-end communication system optimisation, autoencoder-based transceiver design [10], and neural network-based equalisation for mitigating channel impairments [11]. However, these approaches typically operate on high-dimensional waveforms or symbol sequences, requiring substantial computational resources and often lacking interpretability. Our work addresses these limitations by employing a compact feature-based representation that enables efficient matched filter adaptation while maintaining the physical interpretability essential for practical optical systems.

The remainder of the paper is organised as follows. Section II presents the experimental system, including CAP signal generation, bandwidth-limited channel modelling, and receiver processing. Section III details the feature extraction methodology and neural network design. Section IV describes the training procedure. Section V presents performance evaluation results, and Section VI concludes with key findings.

II. Test Setup

The test setup is illustrated in Fig. 1. The experimental setup evaluates a neural network-assisted matched filter estimator for a CAP communication system operating

Manuscript received January 29, 2026.

Paul Anthony Haigh is with the Centre for Networks, Communications and Systems, School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS, UK (email: p.a.haigh@qmul.ac.uk).

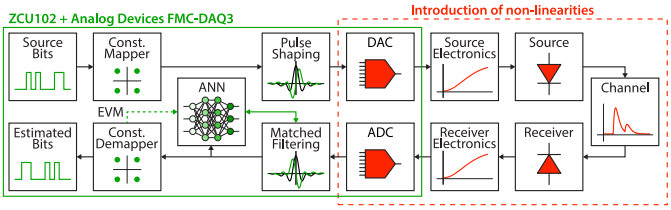


Fig. 1. Block diagram showing the high level architecture of the system.

under bandwidth-limited channel conditions. The proposed approach is benchmarked against a conventional analytically derived CAP matched filter in terms of error vector magnitude (EVM) performance. Random bits are generated in Python and mapped onto the M-ary quadrature amplitude modulation (M-QAM) alphabet. The k^{th} transmitted symbol $s[k]$ is decomposed into in-phase and quadrature components following:

$$s[k] = s_I[k] + j s_Q[k] \tag{1}$$

before upsampling and shaping the respective components with in-phase $p[n]$ and quadrature $p_b[n]$ transmit filters, which are described by [12]:

$$p[n] = b[n]\cos(2\pi f_c n t_s) \tag{2}$$

$$p_b[n] = b[n]\sin(2\pi f_c n t_s) \tag{3}$$

where n is the discrete sample instance, $t_s = f_s^{-1}$ is the sampling period, f_c is the carrier frequency and $b[n]$ is the basis function, which is a square-root raised cosine shape given by [12]:

$$b[n] = \frac{1}{\sqrt{T_s}} \frac{\sin\left(\frac{(1-\beta)\pi n t_s}{T_s}\right) + \frac{4\beta n t_s}{T_s}\cos\left(\frac{(1+\beta)\pi n t_s}{T_s}\right)}{\frac{\pi n t_s}{T_s}\left[1 - \left(\frac{4\beta n t_s}{T_s}\right)^2\right]} \tag{4}$$

where T_s is the symbol period and β is an excess bandwidth factor, set to $\beta = 0.25$ in this work. After pulse shaping, the final real-valued signal for transmission $x[n]$ is given by [12]:

$$x[n] = \sqrt{2}\sum_k s_I[k]p\left[n - \frac{kT_s}{t_s}\right] - s_Q[k]p_b\left[n - \frac{kT_s}{t_s}\right] \tag{5}$$

The system sampling rate is set to $f_s = 1.233$ GS/s, the pulse shaping filters span 8 symbol intervals, and the system bandwidth is set to $B = 100$ MHz. The carrier frequency is $f_c = 0.5B = 50$ MHz. The transmitted signal is normalised to unit power before channel impairment:

$$x[n] \longleftarrow \frac{x[n]}{\sqrt{\frac{1}{N}\sum_{n=0}^{N-1}|x[n]|^2}} \tag{6}$$

The signal is then transferred via the 'libIIO' and 'PyADI-IIO' packages to the memory of a Xilinx Zynq ZCU102 which consists of both a multi-core ARM processor subsystem and programmable logic hosting an Analog Devices FMC-DAQ3 as the digital-to-analogue and analogue-to-digital converter (DAC/ADC).

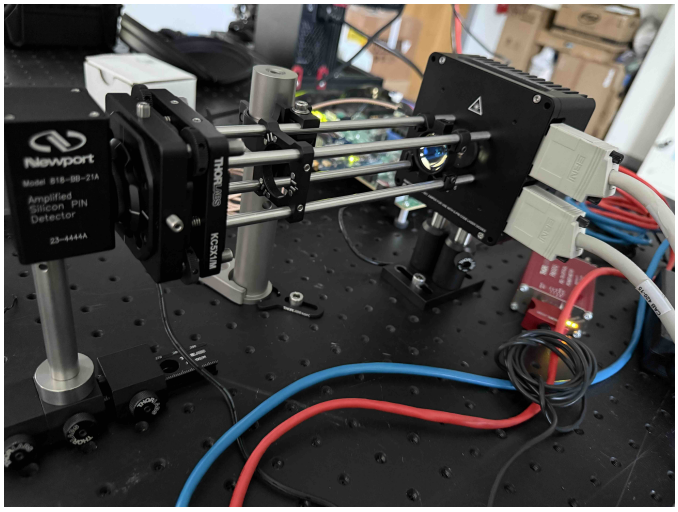


Fig. 2. Photograph of the test setup.

The transmitter electronics consists of a Thorlabs LDM56 laser mount, TED200C temperature controller (set to 25°C) and LDC205C laser driver. The RF bandwidth of these elements is 600 MHz. The laser used was a TT Electronics OPV310 which is an 850 nm VCSEL offering ∼1 GHz bandwidth. The measured L-I-V curve of the laser demonstrates high linearity in the safe operating region, confirming [13]. A photograph of the test setup is shown in Fig. 2. A 25.4 mm aspheric lens (Thorlabs ACL25416U-B) is used to collimate the beam, followed by transmission over a 200 mm free-space channel. The optical power impinging on the receiver is controlled using a series of neutral density filters (Thorlabs NDK01 series) before being focused onto the receiver via a second aspheric lens (Thorlabs C230TMD-B). The receiver is a 1.2 GHz photodiode and integrated transimpedance amplifier (Newport 818-BB-21A).

Each of the optical and electronic components operates within its nominal linear region and does not introduce significant distortion under normal operation. To systematically evaluate receiver performance under constrained spectral conditions, bandwidth limitation is emulated in the digital domain prior to DAC conversion. The channel impairment is modelled as a bandwidth-limited linear system, implemented using a 101-tap FIR filter, where the normalised cut-off frequency is defined as [14]:

$$\omega_n = \frac{2f_n}{f_s} \tag{7}$$

where f_n is the low-pass cut-off frequency. This parameter is swept experimentally to evaluate receiver robustness under progressively constrained spectral conditions. In this work, $\omega_n \in \{0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.9\}$ is swept to evaluate receiver robustness under progressively constrained spectral conditions. Smaller values of ω_n correspond to more severe bandwidth limitation.

After reception, the signal is digitised via the ADC then processed in Python. A low-pass filter with cut-off at $1.2B$ is applied to remove out-of-band noise. Timing synchronisation

is achieved via cross-correlation between the transmitted and received signals [15]:

$$\tau_{\text{sync}} = \arg\max_{\tau} \left| \sum_{n} x[n] y^*[n+\tau] \right| \tag{8}$$

The synchronised signal is then passed through conventional matched filters $\hat{p}[n] = p[-n]$ and $\hat{p_b}[n] = p_b[-n]$ to establish the baseline performance. The neural network receiver is trained to predict deformations to these baseline filters based solely on extracted signal features, enabling adaptive compensation without explicit symbol knowledge or decision feedback.

III. Feature Extraction and Neural Network Design

A. Feature Selection Rationale

The performance of the proposed deformable matched filter approach critically depends on the ability to extract robust features that capture the essential characteristics of channel impairments. Rather than directly exposing the neural network to raw signal samples or detected symbols, we extract a compact set of 16 features that efficiently encode the channel state information. This approach offers several advantages: (i) reduced computational complexity compared to processing raw waveforms, (ii) improved generalisation across diverse channel conditions, and (iii) the ability to operate without explicit symbol knowledge, thereby avoiding error propagation effects.

To improve robustness against localised signal variations, features are extracted from $K = 4$ temporal segments of the received signal and averaged:

$$\boldsymbol{f} = \frac{1}{K} \sum_{k=1}^{K} \boldsymbol{f}_k \tag{9}$$

where $\boldsymbol{f}_k$ represents the feature vector extracted from the k^{th} segment. This segmentation strategy reduces sensitivity to burst impairments while maintaining computational efficiency.

The feature set is organised into three complementary groups: (i) time-domain statistics (5 features), (ii) frequency-domain characteristics (6 features), and (iii) signal quality metrics (5 features). Each group captures distinct aspects of the received signal that are affected by bandwidth limitations and hardware impairments.

1) Time-Domain Statistical Features: Time-domain features capture the amplitude distribution and higher-order statistics of the received signal. The rms amplitude (standard deviation σ) and variance σ^2 provide basic distributional information about signal power. These features are particularly sensitive to amplitude distortions induced by system non-linearities.

Higher-order moments such as skewness and excess kurtosis quantify signal asymmetry and tail heaviness in the signal distribution. In optical communications, kurtosis is known to predict modulation format-dependent non-linear interference [16], [17], with the enhanced Gaussian noise model using kurtosis to correct for overestimation of non-linear interference [18]. Hardware compression causes amplitude distortions and kurtosis increases under non-linear distortion [17], [19], making it a valuable indicator of channel severity.

The signal envelope, obtained via the Hilbert transform $\mathcal{H}\{\cdot\}$, provides additional information about amplitude modulation characteristics. The mean envelope value captures the average signal magnitude after accounting for fast oscillations:

$$\mu_E = \frac{1}{N} \sum_{n=0}^{N-1} |\mathcal{H}\{x[n]\}| \tag{10}$$

These statistical descriptors enable the neural network to distinguish between different impairment mechanisms without requiring explicit channel models.

2) Frequency-Domain Statistical Features: Frequency-domain analysis reveals spectral distortions that are not readily apparent in time-domain representations. We compute the magnitude spectrum via windowed FFT using a Hanning window and extract features that characterise the spectral shape and energy distribution.

The spectral spread quantifies the distribution of energy around the centroid, capturing bandwidth expansion or compression effects caused by filtering. Spectral roll-off identifies the frequency below which 85% of the signal energy is concentrated, offering direct insight into high-frequency attenuation due to bandwidth limitation.

Spectral flatness measures the ratio of geometric to arithmetic mean of the power spectral density:

$$\text{SF} = \frac{\exp\left(\langle \log\left(P[m]\right) \rangle\right)}{\langle P[m] \rangle} \tag{11}$$

distinguishing between tone-like and noise-like spectra. Low flatness indicates strong spectral structure, while high flatness suggests broad spectral spreading due to distortion.

Spectral entropy quantifies the uniformity of energy distribution across frequency bins:

$$H = -\sum_{m=0}^{\frac{N}{2}-1} \tilde{P}[m] \log_2(\tilde{P}[m]) \tag{12}$$

where $\tilde{P}[m] = \frac{P[m]}{\sum P[m]}$ is the normalised power spectral density. Higher entropy indicates greater spectral dispersion, characteristic of severe bandwidth limitation, and finally the 3 dB bandwidth is computed. Together, these features enable the neural network to infer pulse-shape distortions from spectral signatures.

3) Signal Quality Metrics: The third feature group comprises metrics relating to signal quality and temporal structure. Peak-to-average power ratio (PAPR) measures the dynamic range of the signal:

$$\text{PAPR} = \frac{\max\left(x[n]^2\right)}{\frac{1}{N}\sum_{n=0}^{N-1} x[n]^2} \tag{13}$$

High PAPR signals are more susceptible to non-linear distortion, making this feature critical for adaptive filter design. Similarly, the peak power $\max(x[n]^2)$ provides absolute magnitude information complementary to the normalised PAPR.

Autocorrelation features describe the temporal structure and periodicity. We compute normalised autocorrelation at two critical lags:

$$R_{xx}[\ell] = \frac{\sum_{n=0}^{N-\ell-1} x[n]x[n+\ell]}{x[n]^2} \tag{14}$$

The autocorrelation at lag $\ell = \frac{T_s}{t_s}$ (one symbol period) is particularly informative for CAP signals, as pulse-shape distortions directly affect symbol-to-symbol correlation. The autocorrelation at lag $\ell = 1$ captures rapid signal variations and provides information about the effective bandwidth.

Finally, envelope-based metrics derived via the Hilbert transform characterise amplitude modulation effects. The envelope crest factor, defined as the ratio of maximum to mean envelope, captures peak-to-average characteristics in the amplitude domain:

$$\text{CF}_E = \frac{\max\left(|\mathcal{H}\{x[n]\}|\right)}{\mu_E} \tag{15}$$

which is particularly sensitive to signal clipping and saturation effects that may occur in the optical or electronic hardware.

Features are computed after signal normalisation ($\mu = 0$). $\tilde{P}[m] = \frac{P[m]}{\sum P[m]}$ is normalised PSD, $f[m] = mf_sN^{-1}$ is frequency, and $\mathcal{H}\{\cdot\}$ denotes Hilbert transform.

B. Neural Network Architecture Design

The neural network architecture is designed to predict residual corrections to the conventional matched filters rather than learning filters from scratch. This residual learning approach is motivated by the observation that the analytically derived filters are near-optimal for undistorted channels, and only small deformations are needed to compensate for bandwidth limitations.

The network maps the 16-dimensional feature vector $\boldsymbol{f} \in \mathbb{R}^{16}$ to a set of correction terms:

$$\Delta\boldsymbol{h} = f_\theta(\boldsymbol{f}) \in \mathbb{C}^L \tag{16}$$

where L is the length of each matched filter. The final deformable filters are obtained by adding these corrections to the baseline matched filters:

$$\hat{p}[\ell] = p[-\ell] + \Re\{\Delta\mathbf{h}[\ell]\} \tag{17}$$

$$\hat{p_b}[\ell] = p_b[-\ell] + \Im\{\Delta\mathbf{h}[\ell]\} \tag{18}$$

This residual formulation offers several advantages: (i) faster convergence during training since the network only needs to learn small corrections, (ii) better generalisation as the baseline filters provide a strong prior, and (iii) graceful degradation where the network defaults to conventional matched filtering if the learned corrections are zero.

The neural network employs a compact two-layer architecture designed for computational efficiency while maintaining sufficient capacity for the regression task. The hidden layer dimension is dynamically computed as $N_{\text{hidden}} = 2^{\lceil \log_2(L) \rceil}$, which scales the network capacity with filter length while ensuring power-of-two dimensions that are efficient for hardware implementation. For the filter length $L = 192$ used in this work, this yields $N_{\text{hidden}} = 256$ neurons. The output layer consists of $2L = 382$ coefficients (real and imaginary parts concatenated).

The forward pass is computed as follows. The first layer expands the feature space:

$$\boldsymbol{h} = \text{ReLU}(\boldsymbol{W}^{(1)}\boldsymbol{f} + \boldsymbol{b}^{(1)}) \tag{19}$$

where $\boldsymbol{W}^{(1)} \in \mathbb{R}^{256\times16}$ and $\boldsymbol{b}^{(1)} \in \mathbb{R}^{256}$ are trainable parameters. The ReLU activation function $\text{ReLU}(x) = \max(0, x)$ introduces non-linearity while mitigating vanishing gradient problems. The output layer produces the concatenated filter coefficients:

$$\boldsymbol{y} = \boldsymbol{W}^{(2)}\boldsymbol{h} + \boldsymbol{b}^{(2)} \tag{20}$$

where $\boldsymbol{W}^{(2)} \in \mathbb{R}^{382\times256}$ and $\boldsymbol{b}^{(2)} \in \mathbb{R}^{382}$. The output vector is then split and recombined into complex form:

$$\Re\{\Delta\boldsymbol{h}\} = \boldsymbol{y}[0:L] \tag{21}$$

$$\Im\{\Delta\boldsymbol{h}\} = \boldsymbol{y}[L:2L] \tag{22}$$

$$\Delta\boldsymbol{h}[\ell] = \Re\Delta\boldsymbol{h}[\ell] + j\Im\{\Delta\boldsymbol{h}\}[\ell] \tag{23}$$

To ensure the predicted filters maintain consistent energy regardless of the magnitude of corrections, we apply per-sample energy normalisation to ensure learned coefficients do not arbitrarily amplify or attenuate the overall deformed matched filter response.

The total number of trainable parameters is 102,526 (layer 1: $16\times256+256 = 4,352$, layer 2: $256\times382+382 = 98,174$). This compact architecture is significantly smaller than typical deep learning models, enabling rapid training (convergence within 500-700 epochs) and real-time inference. The single hidden layer is sufficient because the feature extraction stage (Section III-A) already provides a rich, physics-informed representation of the channel state. The network's role is to map these pre-processed features to filter corrections rather than learning representations from raw data.

The choice of $N_{\text{hidden}} = 2^{\lceil \log_2(L) \rceil}$ ensures the hidden layer has sufficient capacity to model the L-dimensional filter space while avoiding over-parameterisation. Experiments with deeper architectures (3-4 layers with 256-1,024 neurons) showed minimal performance improvement at the cost of increased training time and susceptibility to over-fitting. An illustrative example of the network structure is shown in Fig. 3, demonstrating how individual features contribute to each filter coefficient through weighted combinations in the hidden layer.

IV. Training Procedure and Implementation

A. Dataset Construction

Training employs online data generation with systematic bandwidth variation ($\omega_n \in \{0.5, 0.55, ..., 0.9\}$). For each epoch, 10,000 random M-QAM symbols are transmitted through the CAP modulator, bandwidth-limited channel, and SDR loopback. After low-pass filtering and correlation-based synchronisation, 16 features averaged over K=4 segments form the input vector. Online training ensures continuous exposure to diverse channel realisations and hardware noise, naturally coupling the network to physical hardware imperfections.

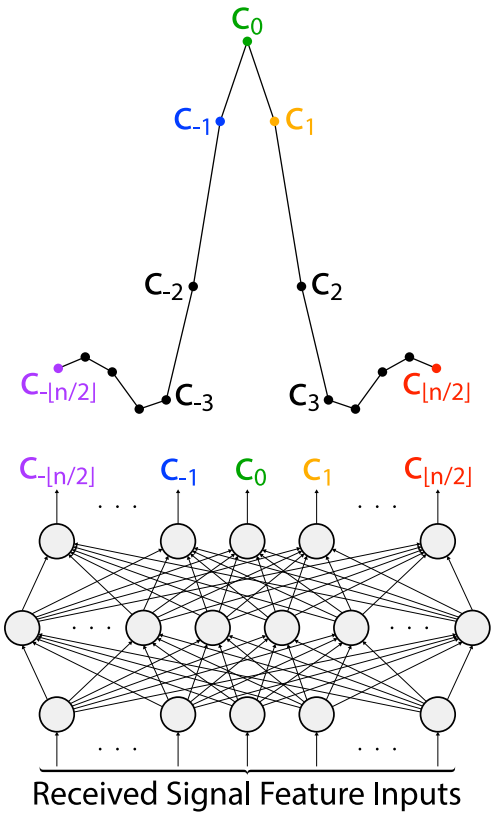


Fig. 3. An illustrative example of how the neural network connects the features to the filter coefficients.

B. Loss Function and Optimisation

The neural network is trained to minimise a composite loss function that directly optimises communication system performance. The primary loss component is error vector magnitude (EVM) computed between received symbols after applying the deformable matched filters and the transmitted reference symbols [20]:

$$\mathcal{L}_{\text{EVM}} = \frac{1}{K}\sum_{k=1}^{K} |\hat{s}[k] - s[k]|^2 \tag{24}$$

recalling that $s[k]$ are the transmitted symbols, $hat(s)[k]$ are the received symbols after matched filtering and downsampling, and $K = 10,000$ is the number of symbols in each training batch. This loss directly measures the quality of the demodulated constellation and provides a clear gradient signal for optimising filter performance.

To ensure the learned filters are physically realisable and do not contain spurious high-frequency components, we incorporate smoothness regularisation penalties. The first-order smoothness loss penalises rapid changes between adjacent filter coefficients:

$$\mathcal{L}_{s1} = \sum_{\ell=1}^{L-1} \left(|\hat{p}[\ell] - \hat{p}[\ell-1]|^2 + |\hat{p_b}[\ell] - \hat{p_b}[\ell-1]|^2 \right) \tag{25}$$

The second-order smoothness loss penalises curvature (second derivative) for even smoother filter responses:

$$\mathcal{L}_{s2} = \sum_{\ell=2}^{L-1} (|\hat{p}[\ell] - 2\hat{p}[\ell-1] + \hat{p}[\ell-2]|^2 + |\hat{p_b}[\ell] - 2\hat{p_b}[\ell-1] + \hat{p_b}[\ell-2]|^2) \tag{26}$$

These regularisation terms prevent the network from learning oscillatory or noisy filter coefficients that would be sensitive to minor variations in channel conditions. The total loss function is a weighted combination:

$$\mathcal{L}_{\text{total}} = +\mathcal{L}_{\text{EVM}} + \lambda_{s1}\mathcal{L}_{s1} + \lambda_{s2}\mathcal{L}_{s2} \tag{27}$$

where $\lambda_{s1} = 10^{-3}$ and $\lambda_{s2} = 10^{-4}$ balance the contributions of the smoothness terms. These hyper-parameters were selected through preliminary experiments to provide sufficient regularisation without excessively constraining the filter adaptation capability.

This end-to-end loss formulation offers several advantages over mean-squared error on filter coefficients. It directly optimises symbol detection quality rather than an intermediate representation, and the differentiable nature of the convolution and downsampling operations enables direct back-propagation from symbol errors to filter coefficients. This allows the network to learn task-specific filters rather than generic approximations.

C. Training Procedure

We employ the Adam optimiser [21] with an initial learning rate of $\eta = 10^{-3}$, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. No learning rate scheduling or early stopping is employed. The model typically converges within 500-700 epochs as evidenced by stable EVM loss, though training continues for the full 1,000 epochs to ensure thorough exploration of the parameter space. Training time is approximately 120 seconds.

To evaluate the performance of the proposed deformable matched filter approach, we compare against the conventional matched filter pair, which uses the analytically derived time-reversed filters $\hat{p}[n] = p[-n]$ and $\hat{p_b}[n] = p_b[-n]$, which are optimal for the undistorted channel. In the case that the EVM obtained when using the conventional matched filter is superior to that of the deformable filter (i.e. with very little ISI or with high noise/ISI), the NN will yield to the conventional matched filter approach. Hence, the NN is only utilised when it can offer an EVM advantage.

V. Results

The performance of the proposed NN-based deformable matched filter is evaluated against the conventional matched filter baseline across a range of bandwidth limitation conditions and received optical powers controlled by calibrated neutral density filters. Fig. 4 presents the comprehensive EVM performance comparison as a function of received optical power for the eight different normalised cut-off frequencies with optical power swept from -25 dBm to +5 dBm for 4-QAM.

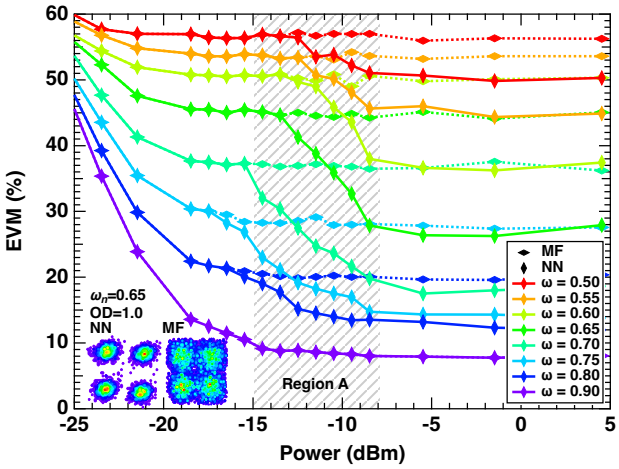


Fig. 4. EVM vs received optical power for 4-QAM across eight ω_n values. Solid: conventional MF; dashed: proposed NN. Region A shows maximum improvement.

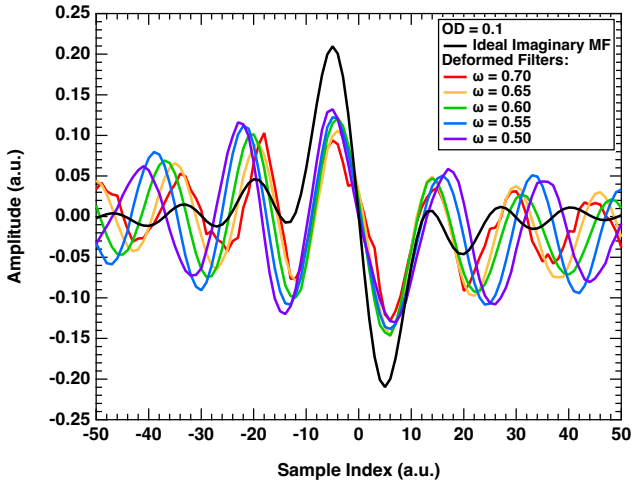


Fig. 5. Imaginary component of the deformed matched filters learned by the neural network for different normalised cut-off frequencies at OD = 0.1. The ideal matched filter (dashed black line) is shown for reference.

Under severe bandwidth limitation ($\omega_n = 0.5$), the conventional matched filter exhibits EVM values exceeding 50% at low optical powers, degrading the constellation to the point where reliable demodulation becomes infeasible. In contrast, the neural network-assisted deformable matched filter achieves EVM values below 30% under identical conditions, representing a relative performance improvement >40%. This substantial gain demonstrates the ability of the learned filter deformations to compensate for the severe ISI induced by aggressive bandwidth restriction.

As the bandwidth constraint is relaxed (increasing ω_n), the performance gap between conventional and deformable matched filters diminishes. For $\omega_n = 0.9$, which represents mild bandwidth limitation, both approaches converge to similar EVM performance, with values below 10% achieved at received powers above -10 dBm. This convergence validates the design philosophy of residual learning: when channel distortion is minimal, the neural network learns correction terms close to zero, effectively defaulting to the analytically optimal matched filter.

A critical observation from Fig. 4 is the transition region between severe and mild bandwidth limitation ($\omega_n \in [0.65, 0.75]$). Within this regime, the deformable matched filter provides moderate but consistent improvements of 5-15% in EVM compared to conventional filtering. This intermediate performance gain suggests that the neural network successfully learns to balance between preserving the theoretically optimal filter structure and adapting to channel-induced distortions.

The relationship between received optical power and EVM follows the expected inverse trend for both receiver architectures. At low optical powers (below -20 dBm), photon shot noise and thermal noise in the photodetector dominate the error budget, limiting the benefit of improved matched filtering. However, as optical power increases and signal-to-noise ratio improves, the ISI introduced by bandwidth limitation becomes the dominant impairment. In this regime (approximately -15 dBm to -8 dBm, denoted region A in Fig. 4), the deformable matched filter achieves maximum relative performance gains, with EVM improvements exceeding 50% under moderate bandwidth limitations. At high received powers (above -5 dBm), the performance of both receivers is stable. For $\omega_n = 0.9$, the excess bandwidth of the conventional matched filters is sufficient to eliminate the ISI and no further performance is gained, due to the SNR ceiling of the system. Similar improvements are seen for 16-QAM but not shown due to space constraints.

The learned filter coefficients provide insight into the adaptation mechanism employed by the neural network. Fig. 5 displays the imaginary component of the deformed matched filters for a fixed optical density (OD = 0.1) across varying bandwidth constraints. The ideal imaginary matched filter $\hat{p}_b[n] = p_b[-n]$ is shown for reference as the baseline.

The deformable filter shifts energy into side-lobes and repositions peaks as needed, unlike static matched filters. As ω_n decreases, deformations become more pronounced, where central lobe broadening compensates for temporal spreading, while side-lobe reshaping counteracts bandwidth-induced distortion. The antisymmetric Hilbert-pair structure is largely preserved, suggesting implicit orthogonality maintenance. Under severe limitation ($\omega_n \leq 0.55$), substantial restructuring suppresses ISI-causing oscillations while maintaining filter smoothness, validating the regularisation strategy, eq. (27).

Fig. 6 shows training convergence for different ω_n (OD = 0.1). All configurations converge rapidly within 100-200 epochs, with refinement continuing to epoch 1000. Moderate constraints ($\omega_n \geq 0.7$) stabilise around $10^{-2} - 10^{-3}$, while severe constraints ($\omega_n \leq 0.6$) converge higher ($\sim 10^{-1}$) due to fundamental ISI limits. Hardware-in-the-loop training automatically accommodates analogue imperfections and component tolerances. One limitation apparent from the results is that performance gains diminish when channel conditions are either very favourable or extremely degraded. Under mild bandwidth limitation ($\omega_n \geq 0.85$), conventional matched filtering already approaches near-optimal performance, leav-

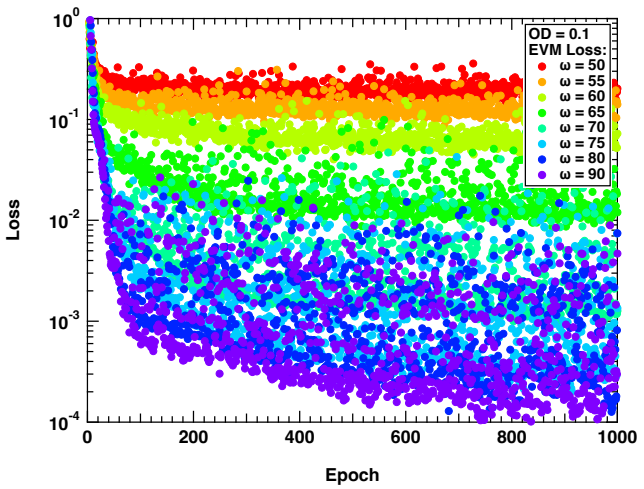


Fig. 6. Training convergence curves showing EVM loss as a function of epoch for different ω_n. All configurations exhibit rapid initial convergence within 100-200 epochs, followed by gradual refinement.

ing little room for improvement through adaptive filtering. Under severe degradation ($\omega_n < 0.5$, not tested in this work), the fundamental ISI introduced by extreme bandwidth restriction exceeds the ability of linear filtering, and more sophisticated techniques such as decision-feedback equalisation, maximum-likelihood sequence estimation or non-linear equalisation would be required.

Across the intermediate regime ($\omega_n \in [0.5, 0.8]$) where most practical systems operate, the deformable matched filter provides consistent improvements with minimal implementation complexity, particularly in region A where ISI dominates.

VI. Conclusion

In this paper we demonstrated a neural network-assisted deformable matched filter framework for CAP modulation in a bandwidth-constrained optical communication system. Unlike conventional approaches that replace analytical receiver structures with black-box machine learning models, the proposed method learns residual corrections to theoretically optimal matched filters based on a compact set of 16 physically motivated signal features. This hybrid strategy combines the reliability of classical communication theory with the adaptability of data-driven optimisation. The feature-based adaptation framework could be extended to jointly optimise transmit and receive filters in systems where both link endpoints are adaptive. By learning adaptive corrections to matched filters rather than replacing them entirely, the proposed approach achieves robust performance improvements while maintaining the ability to gracefully degrade to conventional operation when conditions are favourable. Experimental validation used hardware-in-the-loop transmission to demonstrate substantial performance improvements under bandwidth-limited channel conditions. Under moderate constraints ($\omega_n \in [0.65, 0.75]$), consistent improvements up to 50% are observed. Importantly, when channel conditions are favourable ($\omega_n \geq 0.85$), the network gracefully defaults to conventional matched filtering, eliminating the risk of performance degradation.

References

[1] J. Zheng, X. Li, S. Qiang, and Y. Wang, "Decoding scheme based on cnn for differential free space optical communication system," *Optics Communications*, vol. 559, p. 130449, 2024.

[2] H. Yu, L. Li, Y. Hou, Y. Li, C. Yin, C. Gao, and S. Fu, "Deep learning-based prediction of atmospheric turbulence toward satellite-to-ground laser communication," *Optics Letters*, vol. 50, no. 2, pp. 273–276, 2025.

[3] W. Jiang, M. Cheng, L. Guo, X. Yi, J. Li, J. Wang, and A. Forbes, "Machine learning assisted speckle and oam spectrum analysis for enhanced turbulence characterization," *Photon. Res.*, vol. 13, no. 10, pp. B29–B37, Oct 2025. [Online]. Available: https://opg.optica.org/prj/abstract.cfm?URI=prj-13-10-B29

[4] D. Zhang, J. Zhang, Y. Gao, and T. Du, "Optical field-to-field translation under atmospheric turbulence: A conditional gan framework," *Photonics*, vol. 12, no. 4, p. 339, 2025.

[5] X. Li, Y. Wang, X. Liu, Y. Ma, Y. Cai, S. A. Ponomarenko, and X. Liu, "Deep learning and random light structuring ensure robust free-space communications," *Applied Physics Letters*, vol. 124, no. 21, 2024.

[6] J. Zheng, X. Li, S. Qiang, and Y. Wang, "Decoding scheme based on cnn for differential free space optical communication," *Optics Communications*, vol. 559, p. 130449, 2024.

[7] Y. Gao, B. Yang, S. Fan, L. Xu, T. Wang, B. Yang, and S. Jiang, "A hybrid deep learning-based modeling methods for atmosphere turbulence in free space optical communications," *Photonics*, vol. 12, no. 12, p. 1210, 2025.

[8] Y. Gao, Q.-W. Jing, M.-F. Liu, W.-H. Zong, and Y.-Q. Hong, "Deep learning-assisted high-pass-filter-based fixed-threshold decision for free-space optical communications," *Photonics*, vol. 11, no. 7, p. 599, 2024.

[9] 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, pp. 563–575, Dec 2017.

[10] S. Dorner, S. Cammerer, J. Hoydis, and S. ten Brink, "Deep learning based communication over the air," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 132–143, 2018.

[11] B. Karanov, M. Chagnon, F. Thouin, T. A. Eriksson, H. Bülow, D. Lavery, P. Bayvel, and L. Schmalen, "End-to-end deep learning of optical fiber communications," *Journal of Lightwave Technology*, vol. 36, no. 20, pp. 4843–4855, 2018.

[12] J. G. Proakis and M. Salehi, *Digital Communications*, 5th ed. New York: McGraw-Hill, 2008.

[13] TT Electronics, *OPV310 850 nm VCSEL Datasheet*, TT Electronics, 2015, rev. 2.0. [Online]. Available: https://www.ttelectronics.com/...

[14] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 1999.

[15] U. Mengali and A. N. D'Andrea, "Synchronization techniques for digital receivers," 1997.

[16] K. Wu, G. Liga, A. Sheikh, F. M. Willems, and A. Alvarado, "Temporal energy analysis of symbol sequences for fiber nonlinear interference modelling via energy dispersion index," *Journal of Lightwave Technology*, vol. 39, no. 18, pp. 5766–5782, 2021.

[17] J. Cho and R. W. Tkach, "On the kurtosis of modulation formats for characterizing the nonlinear fiber propagation," *Journal of Lightwave Technology*, vol. 40, no. 12, pp. 3739–3748, 2022.

[18] A. Carena, G. Bosco, V. Curri, Y. Jiang, P. Poggiolini, and F. Forghieri, "Egn model of non-linear fiber propagation," *Optics Express*, vol. 22, no. 13, pp. 16 335–16 362, 2014.

[19] Y. C. Gültekin, A. Alvarado, F. M. Willems, and G. Liga, "Kurtosis-limited sphere shaping for nonlinear interference noise reduction in optical channels," *Journal of Lightwave Technology*, vol. 40, no. 11, pp. 3356–3367, 2022.

[20] R. A. Shafik, M. S. Rahman, and A. H. M. R. Islam, "On the extended relationships among EVM, BER and SNR as performance metrics," in *International Conference on Electrical and Computer Engineering*. IEEE, 2006, pp. 408–411.

[21] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017. [Online]. Available: https://arxiv.org/abs/1412.6980

Acknowledgements

This work declares no funding.

Author contributions

P.A.H conceived the idea, wrote the paper, performed the field tests and wrote and confirmed approval of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.