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.

Explainable deformable matched filtering reveals measurable departures from classical receiver theory in optical wireless communications

Paul Anthony Haigh

p.a.haigh@qmul.ac.uk

Queen Mary University of London https://orcid.org/0000-0002-7249-5116

Article

Keywords:

Posted Date: September 7th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10475047/v1

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

Additional Declarations: There is no conflict of interest

Explainable deformable matched filtering reveals measurable departures from classical receiver theory in optical wireless communications

Paul Anthony Haigh

Abstract

Matched filtering is a central result of communication theory, providing the optimal linear receiver when the received waveform satisfies specific assumptions. Practical communication systems rarely satisfy these assumptions, yet learned receivers that outperform the classical matched filter provide little insight into what those improvements reveal about the limitations of the underlying theory.

Here we introduce an explainable deformable matched-filter framework in which machine learning is constrained to learn a low-dimensional deformation of the classical matched filter rather than replacing it. Because every learned correction is defined relative to the theoretical matched-filter solution, the deformation becomes a measurable representation of receiver mismatch rather than an unconstrained optimisation. The communication waveform remains processed entirely by the matched filter, while a Kolmogorov-Arnold Network predicts only the deformation from physically interpretable receiver-state descriptors.

Using an optical wireless communication testbed spanning ten signalling formats, four impairment classes and 1,600 conditions, we show that learned deformations improve receiver performance, yielding a median relative error-vector-magnitude reduction of 18.1%, while revealing departures from classical matched-filter optimality. Different signalling families occupy distinct deformation regimes, spectral analysis identifies the physical mechanisms underlying receiver mismatch, and latent receiver-state organisation demonstrates that these departures are structured rather than arbitrary.

These results establish deformable matched filtering as a practical receiver design strategy and a framework for experimentally investigating the gap between analytical receiver theory and practical communication systems. More broadly, they illustrate how machine learning can interrogate analytical models by expressing learned corrections relative to theoretical optima rather than replacing them.

Matched filtering is one of the fundamental results of communication theory, providing the optimal linear receiver when the received waveform is a noisy version of a known transmit pulse in additive white Gaussian noise [1]. This result underpins modern digital communication systems and remains a standard component of practical receivers. Optical communications links rarely satisfy the assumptions under which matched-filter optimality is derived [2-4]. Transmitter bandwidth limitations, receiver filtering, AC coupling, device nonlinearities and implementation-dependent hardware distort the received waveform, creating systematic mismatch between the theoretical matched filter and the filter that is optimal in practice [4]. Although these departures are widely recognised, they are generally treated as engineering imperfections rather than as an opportunity to investigate the limits of classical receiver theory.

Modern communication receivers compensate such mismatch using adaptive equalisers, decision-feedback equalisation, Volterra methods and increasingly, machine-learning architectures [5-10]. Recent years have seen rapid growth in neural receivers for optical and wireless communications [11-15], with models trained directly on the received waveform to optimise symbol recovery [16-21]. These approaches often achieve impressive communication performance, but their internal processing is typically difficult to interpret because the learned representation replaces rather than modifies the classical receiver. Consequently, receiver optimisation has become increasingly detached from the theoretical matched-filter framework, making it difficult to determine how and why practical receivers depart from classical optimality.

This leaves an important unanswered question. If a learned receiver consistently outperforms the classical matched filter (CMF), what does that improvement reveal about the assumptions underlying matched-filter theory? Existing adaptive and neural receivers are designed to minimise detection error, not to explain the origin or structure of receiver mismatch [5, 16, 22]. As a result, they provide limited insight into whether departures from the matched-filter solution arise from bandwidth limitation, signalling structure, nonlinear distortion or other physical mechanisms. The receive filter itself is no longer an explicit object of study.

Here we introduce an explainable deformable matched-filter (DMF) framework that addresses this problem by preserving the matched filter as the central receiver element while allowing it to deform in response to the observed receiver state. Rather than processing the communication waveform directly, a compact set of physically interpretable receiver-state descriptors is extracted from the received signal and supplied to a Kolmogorov-Arnold Network (KAN), which predicts a low-dimensional parameterisation of matched-filter deformation [23, 24]. The communication waveform itself remains entirely within a CMF receiver. Machine learning

therefore configures the receiver rather than acting as the receiver, preserving a direct correspondence between learned parameters and filter shape.

This architectural distinction changes the role of machine learning within the receiver. Rather than replacing the matched filter with a learned signal-processing block, the proposed framework preserves the CMF as the reference solution and learns only the deformation required to account for practical receiver conditions. Every learned parameter therefore has meaning only through its relationship to the nominal matched filter, allowing receiver adaptation to be expressed as a measurable departure from the theoretical optimum rather than as an unconstrained optimisation. Persistent non-zero deformation therefore provides a measurable signature of departures from the assumptions under which classical matched-filter optimality is derived. The learned deformation is therefore more than a mechanism for improving receiver performance: it provides an experimentally observable representation of receiver mismatch, allowing the receiver itself to become a tool for investigating the limits of classical receiver theory.

Using an experimental optical wireless communication testbed spanning multiple signalling families, optical power levels and controlled impairment conditions, we demonstrate that learned matched-filter deformation consistently improves receiver performance while revealing systematic, signalling-dependent departures from CMF optimality. Analysis of the learned deformation coefficients, spectral responses, latent receiver-state organisation and deformation complexity shows that different signalling families occupy distinct deformation regimes rather than following a universal adaptation strategy. The principal contribution is therefore not a new neural receiver, but an interpretable framework for measuring where, how and why practical communication systems depart from one of the central results of classical receiver theory. This establishes a general framework in which learned corrections are expressed relative to analytical receiver theory, allowing machine learning to be interpreted as measurable departures from theoretical optimality rather than as unconstrained optimisation.

Relationship to previous deformable matched-filter work

Previous work demonstrated that neural-network-assisted deformable matched filters can improve receiver performance in bandwidth-limited optical communication systems [25]. That study focused primarily on communication performance and treated the learned deformation as a mechanism for compensating pulse distortion.

The present work addresses a different question. Rather than asking whether deformable matched filtering improves receiver performance, we investigate what the learned deformation reveals about departures from classical matched-filter optimality.

To address this question, the present study introduces an interpretable KAN parameterisation, extends the analysis across multiple signalling families and impairment classes, and systematically analyses the resulting deformation coefficients, filter responses and adaptation trajectories. The learned deformation is therefore treated as scientific information rather than solely as a performance enhancement mechanism.

Results

Deformable matched filtering improves performance across signalling formats and receiver conditions

Throughout this work, deformable matched filtering (DMF) refers to a receiver in which the communication waveform is processed exclusively by a matched filter whose coefficients are configured by a learned deformation. The neural network neither performs symbol detection nor processes the communication waveform directly; it predicts only the deformation parameters from a compact receiver-state description. The deformation is therefore defined relative to the theoretical matched-filter solution and can be analysed as a measurable representation of receiver mismatch. This architectural distinction underpins the interpretation developed throughout the remainder of the paper.

DMF reduced EVM across all investigated signalling formats, attenuation levels and impairment conditions (Figure 1). Across 1,600 evaluated format-impairment-attenuation-session conditions, the median relative EVM reduction was 18.1% (interquartile range, 10.9-29.7%), with improvements ranging from 0.9% to 82.6%. The median reductions were 10.5% for carrier-less amplitude and phase (CAP) modulation, 21.5% for pulse amplitude modulation (PAM), 23.1% for duobinary (DB) and 14.2% for modified duobinary (MDB), demonstrating that the deformable matched-filter framework is not restricted to a single modulation family.

Although performance gains were observed throughout the dataset, their magnitude varied substantially between receiver conditions. Bandwidth-limited conditions produced the largest median relative EVM reduction (30.8%), compared with 16.3% under baseline operation, 11.2% under isolated nonlinearity and 20.3% under combined impairment. CAP formats exhibited relatively modest improvements under baseline conditions but substantially larger gains under bandwidth-limited operation [26, 27]. DB formats displayed large improvements even without an additional imposed impairment, indicating that the mismatch revealed by DMF is not solely attributable to externally introduced distortion.

The dependence of performance gain on attenuation was similarly continuous rather than abrupt (Figure 2). Mean relative EVM reduction decreased from 32.9% to 6.2% for CAP, 31.2% to 9.6% for PAM, 29.5% to 15.6% for DB and 40.8% to 13.3% for MDB between 1.3 and

2.8 dB neutral-density attenuation. DMF therefore remained beneficial throughout the investigated range, but the recoverable component of mismatch declined progressively as attenuation increased.

These observations establish that DMF provides a robust performance advantage across a wide range of operating conditions. More importantly, the substantial variation in benefit between signalling formats and impairment types suggests that the learned deformation is responding to specific departures from matched-filter optimality rather than acting as a generic equalisation mechanism. Figure 1 therefore primarily serves as a map of the mismatch landscape explored throughout the remainder of the study, as well as a performance comparison.

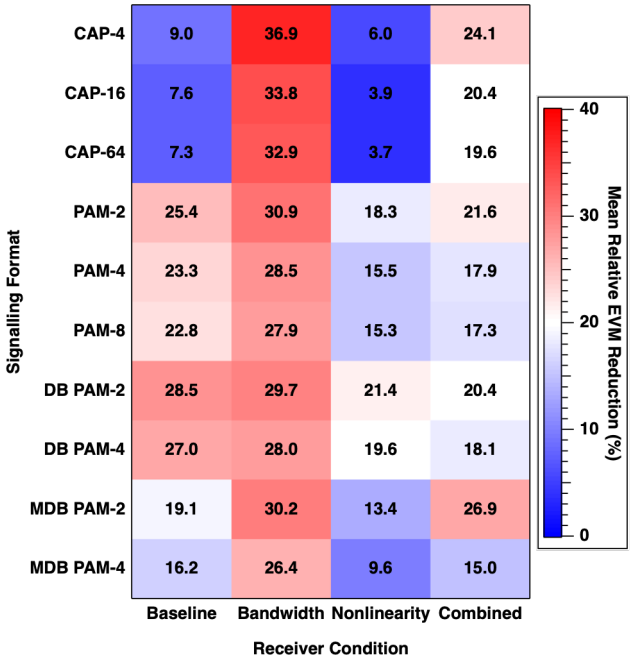


Figure 1 Performance landscape of deformable matched filtering across signalling formats and impairment conditions. This figure establishes the operating conditions under which DMF provides measurable improvement. Mean relative error vector magnitude (EVM) reduction achieved by the deformable matched filter (DMF) relative to the classical matched filter (CMF) for ten signalling formats evaluated under four receiver conditions: baseline, bandwidth-limited, non-linear and combined impairments. Values are averaged across all evaluated parameter settings for each signalling family and condition, with the numerical value shown within each cell. Colour indicates the magnitude of the mean relative EVM reduction. CAP denotes carrier-less amplitude and phase modulation, PAM pulse amplitude modulation, DB duobinary signalling and MDB modified duobinary signalling. Relative EVM reduction is defined with respect to the corresponding CMF performance under identical experimental conditions.

Signalling formats expose distinct departures from matched-filter optimality

To investigate the origin of the observed performance differences and to understand if they arose from a common deformation mechanism, we examined the relationship between learned deformation and receiver improvement (Figure 3). If the deformable matched filter simply applied a generic correction, all signalling formats would be expected to follow a common deformation-to-performance relationship. Instead, the modulation families occupied distinct deformation regimes rather than following a universal deformation-to-gain relationship.

To characterise departures from the classical matched-filter solution, we quantified the learned deformation using three complementary metrics: deformation energy, relative deformation energy and coefficient energy. Deformation energy measures the absolute energy of the learned correction relative to the nominal matched filter. Relative deformation energy normalises this quantity by the energy of the underlying matched filter, providing a scale-independent measure of mismatch. Coefficient energy characterises the magnitude of the learned KAN coefficients responsible for generating the deformation.

These quantities provide complementary descriptions of receiver mismatch. Deformation energy measures the extent to which the practical receive filter departs from the theoretical matched-filter solution, whereas coefficient energy quantifies the internal deformation required to

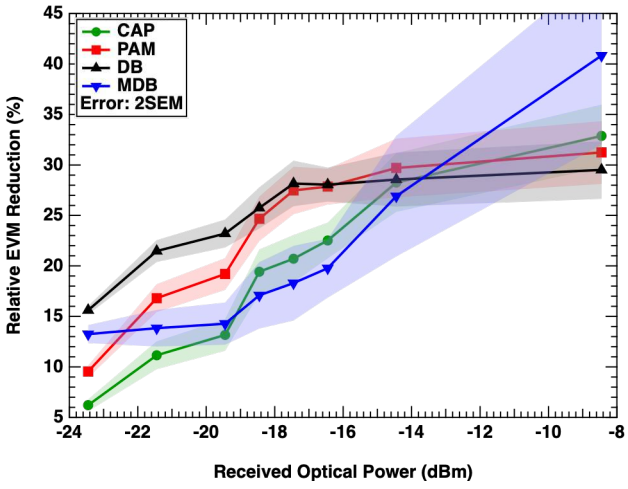


Figure 2 Sensitivity of DMF performance to received optical power. This figure examines the dependence of that improvement on received signal strength. Mean relative EVM reduction as a function of received optical power for CAP, PAM, DB and MDB signalling families. Curves represent arithmetic means over all impairment conditions evaluated at each received power, while shaded regions indicate ±2 standard errors of the mean (SEM). Relative EVM reduction is calculated with respect to the corresponding CMF operating under identical experimental conditions.

generate this departure.

The DMF can be expressed as the sum of a nominal matched filter and a learned deformation. Under ideal matched-filter assumptions, the optimal deformation would be expected to approach zero. Persistent non-zero deformations therefore provide direct evidence that practical receiver operation departs from the assumptions underlying classical matched-filter theory.

If DMF simply compensated a common impairment mechanism, all signalling formats would be expected to follow a single deformation-performance relationship. Instead, distinct modulation families occupied different deformation regimes. CAP exhibited the clearest monotonic relationship between deformation magnitude and performance improvement, with bandwidth-restricted conditions occupying the highest-gain and highest-deformation region. PAM displayed a related but weaker trend. Across all 1,600 conditions, the rank correlation between relative EVM gain and coefficient energy was ρ=0.425 ($P < 10^{-70}$), showing that greater deformation was generally associated with larger benefit while also confirming substantial family-dependent scatter.

The DB format family behaved differently. Large deformations were observed even when the associated performance improvement was comparatively modest. Consequently, the pooled deformation-performance relationship within the DB family differed substantially from those observed for CAP and PAM. Meanwhile, MDB formats occupied an intermediate position. Deformation magnitude and performance improvement remained positively related, but the overall level of deformation was substantially reduced relative to DB.

These results indicate that deformation magnitude alone cannot be interpreted independently of signalling format. Instead, each modulation family introduces a characteristic form of receiver mismatch, producing distinct deformation regimes. Importantly, the dominant organisation is by signalling family rather than modulation order. CAP-4, CAP-16 and CAP-64 occupy a common regime distinct from PAM, DB and MDB. This observation suggests that the primary determinant of receiver mismatch is the spectral and structural character of the signalling family rather than the number of transmitted symbol levels.

If the learned deformation were merely a flexible correction mechanism, no systematic relationship would be expected between deformation magnitude and receiver performance. Instead, deformation energy increased consistently in operating regimes where the conventional matched filter became increasingly suboptimal. This behaviour indicates that the learned deformation captures physical receiver mismatch rather than arbitrary filter variation.

The learned deformation therefore provides information not only about impairment severity but also about the

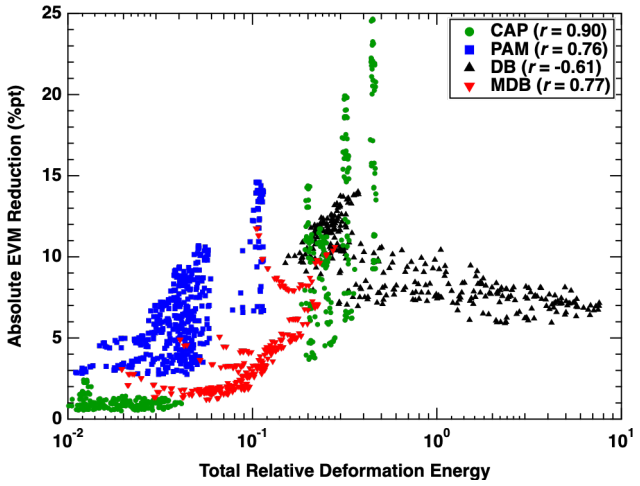


Figure 3 Relationship between matched-filter deformation and performance improvement. This figure relates the magnitude of the learned deformation to the resulting performance gain. Absolute EVM reduction is plotted against the total relative deformation energy for all evaluated receiver conditions. Each point represents an individual experimental condition. Colours and marker shapes distinguish signalling families. Pearson correlation coefficients are reported for each signalling family to quantify the relationship between deformation magnitude and absolute performance improvement. Total relative deformation energy is computed from the learned DCT deformation coefficients relative to the corresponding classical matched filter.

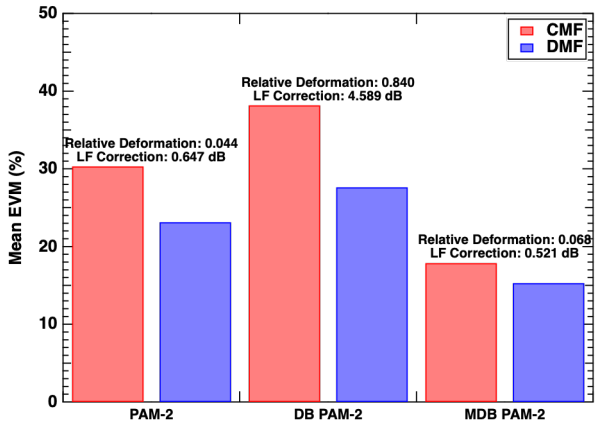


Figure 4 Comparison of deformation magnitude and receiver performance across representative signalling formats. This figure compares representative signalling families with markedly different deformation behaviour. Mean EVM achieved by the classical matched filter (CMF) and deformable matched filter (DMF) for representative PAM-2, DB PAM-2 and MDB PAM-2 signalling formats. The corresponding total relative deformation and equivalent low-frequency correction are shown above each pair of bars. Bar heights represent arithmetic mean EVM across the evaluated impairment conditions for each signalling format.

physical mechanism responsible for the mismatch. Rather than acting as a generic performance-enhancement mechanism, DMF reveals how different signalling formats depart from matched-filter optimality in distinct and measurable ways.

Partial-response signalling reveals intrinsic receiver mismatch

The contrasting behaviour of DB and MDB signalling provides a controlled experiment for understanding the origin of receiver mismatch (Figure 4).

Under baseline conditions, DB PAM-2 exhibited substantially higher EVM than conventional PAM-2 and required markedly larger deformations. In particular, the DMF applied a strong low-frequency spectral deformation, indicating a significant departure between the classical matched-filter response and the response required for optimal practical operation. MDB PAM-2 exhibited dramatically different behaviour. Baseline EVM was reduced by more than a factor of two relative to DB PAM-2, while the learned low-frequency deformation decreased from approximately 4.6 dB to approximately 0.5 dB. Deformation magnitude similarly decreased to levels comparable with conventional PAM.

The reduction in learned deformation is particularly significant because MDB differs from DB primarily through its spectral shaping properties. MDB suppresses low-frequency content that is poorly supported by the optical wireless channel, thereby reducing the mismatch between the received waveform and the assumptions underlying the CMF.

Consequently, DMF does not merely indicate that MDB achieves better performance. Instead, it identifies the specific component of receiver mismatch removed by the modified signalling strategy. Because DMF is defined relative to the CMF solution, the reduction in learned low-frequency deformation can be interpreted as a reduction in receiver mismatch rather than merely a reduction in EVM.

The DB-MDB comparison therefore transforms DMF from a performance-enhancement method into a diagnostic tool. The learned deformation provides a direct measurement of the mismatch removed by the modified signalling design.

Learned spectral deformations reveal departures from classical receiver assumptions

To investigate the physical origin of the learned deformation, we examined the frequency-domain response of the DMF under progressively increasing bandwidth limitation (Figure 5).

Using a CMF assumes that the received waveform remains matched to the expected pulse shape. Bandwidth limitation violates this assumption by selectively attenuating spectral components of the transmitted signal and altering the effective pulse shape at the receiver.

The learned deformations evolved systematically as bandwidth restriction increased. Under mild bandwidth limitation, only modest spectral modifications were required. As the available bandwidth decreased, the DMF increasingly attenuated excess low-frequency energy, with the dominant change occurring below approximately 0.1 normalised frequency. The response did not attempt an unstable inversion of the strongly attenuated high-frequency region; instead, it reshaped the portion of the spectrum that remained reliably supported by the channel.

Importantly, the growth of these spectral deformations paralleled both deformation energy and EVM improvement. The progression in Figure 5 was smooth rather than threshold-like. Increasing low-pass-filter restriction produced progressively stronger low-frequency attenuation, and conditions requiring the largest spectral deformation were also those exhibiting the largest EVM improvement. This agreement links the measured performance gain to a specific, physically interpretable change in the receiver response.

These observations provide direct evidence that the learned deformation responds to departures from the assumptions underlying classical receiver theory. Rather than applying an arbitrary non-linear transformation, DMF modifies the nominal matched-filter response in a manner consistent with the imposed bandwidth limitation while avoiding ill-conditioned high-frequency equalisation and

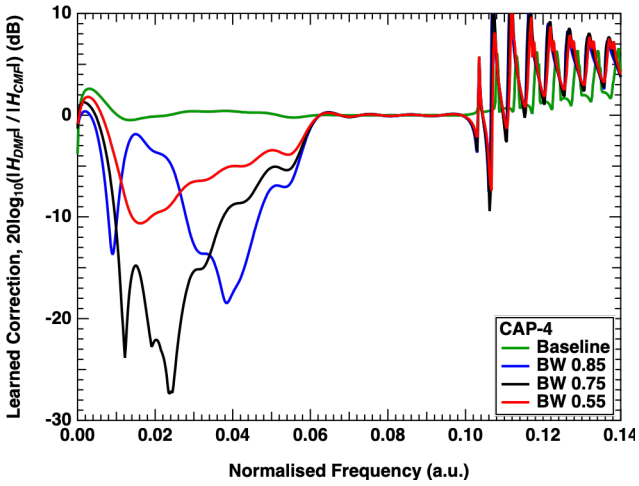


Figure 5 Frequency-domain representation of the learned matched-filter deformation. This figure illustrates the spectral mechanism through which the learned deformation is realised. Arithmetic mean frequency response correction learned by the DMF relative to the corresponding CMF for CAP-4 signalling under progressively increasing bandwidth limitation. The plotted quantity is $20\log_{10}(|H_{\mathrm{DMF}}|/|H_{\mathrm{CMF}}|)$, where H_{DMF} and H_{CMF} denote the frequency responses of the deformable and classical matched filters, respectively. Curves represent arithmetic means over 20 independently trained models. Normalised frequency is expressed relative to the discrete-time sampling frequency.

the associated performance degradations.

Among all investigated results, the frequency-domain analysis provides the clearest mechanistic link between theoretical assumption, physical impairment and measured receiver response. The deformation therefore acts as a direct observable of matched-filter mismatch.

Learned receiver-state representations exhibit family-dependent organisation

The deformable matched filter exhibited systematic organisation within its learned latent representation (Figure 6).

Projection of more than 8×10^5 latent vectors showed that signalling family was the dominant organising factor. CAP, PAM, DB and MDB occupied characteristic regions of the latent space, consistent with the family-specific deformation regimes identified in Figure 3. By contrast, impairment classes overlapped substantially within these regions and did not form clean, globally separable clusters.

The latent representation should therefore not be interpreted as an impairment classifier. Instead, it provides a compact receiver-state description in which operating conditions redistribute the population within a family-dependent structure. Bandwidth and non-linear conditions produce systematic shifts and changes in occupancy, but these condition-dependent redistributions are superimposed on the stronger organisation imposed by signalling family.

This result reinforces the central interpretation of DMF. The learned state is structured, but its structure reflects the interaction between signalling family and receiver condition rather than a one-to-one encoding of impairment identity. The latent space therefore provides a measurable representation of how the receiver configuration changes across the experimental landscape.

Together with the deformation-gain analysis, Figure 6 shows that receiver mismatch is not governed by a single scalar severity axis. Different signalling families require different deformation states, while individual impairments move the receiver within those broader family-dependent regimes.

Deformation complexity depends on impairment mechanism

The dimensionality of the learned deformation also varied systematically between signalling formats and impairment conditions (Figure 7).

Effective coefficient rank varied from 0 to 15.1 across condition and branch combinations (mean 4.11), showing that mismatch could not generally be represented by a single gain or bandwidth adjustment. Mean rank was lowest for DB PAM-2 (1.15) and DB PAM-4 (1.48), intermediate for MDB and conventional PAM (2.86-4.25), and highest for CAP (6.53-7.42 across modulation orders). Different receiver conditions therefore recruited deformation subspaces of differing dimensionality. The dependence of coefficient rank on both signalling family and impairment class further supports the interpretation that DMF responds to distinct physical mechanisms rather than a single universal distortion process.

Taken together, the deformation magnitude, spectral response, latent organisation and coefficient complexity all indicate that the deformable matched filter learns structured deformations reflecting the specific departures from classical matched-filter assumptions present within the optical wireless channel. The convergence of these independent analyses strengthens the central interpretation of this work: the learned deformation is not simply a means of improving receiver performance, but a measurable and interpretable representation of receiver mismatch.

Discussion

Matched filtering occupies a central position in communication receiver theory because it provides the optimal linear receiver when the received waveform, channel response and noise statistics satisfy the assumptions under which the filter is derived. In practical optical wireless systems, however, these assumptions are frequently violated by bandwidth limitation, nonlinear transfer characteristics and signalling formats whose spectral properties interact unfavourably with the channel response.

The DMF should not be confused with conventional adaptive filtering. Classical adaptive filters iteratively update receiver coefficients to minimise an error criterion. In contrast, DMF preserves the matched-filter structure and learns a deformation of the theoretical optimum. The deformation therefore remains interpretable with respect to the underlying matched-filter solution and can be analysed as a measurable representation of receiver mismatch.

The present results demonstrate that DMF can be used to quantify departures from theoretical matched-filter optimality. The performance landscape in Figure 1 establishes that the benefit is widespread but highly condition dependent, while Figure 2 shows that the gain changes continuously with optical attenuation rather than appearing only beyond a discrete operating threshold. These observations motivate interpretation of the deformation itself: the receiver is beneficial across the dataset, but the amount and form of the required correction reveal where the nominal matched filter is most mismatched.

Bandwidth limitation provides the clearest example. CAP shows modest DMF benefit under baseline conditions but substantially larger improvements under bandwidth-constrained operation (Figure 1 and Figure 3). Figure 5 identifies the corresponding mechanism: stronger low-pass restriction produces progressively

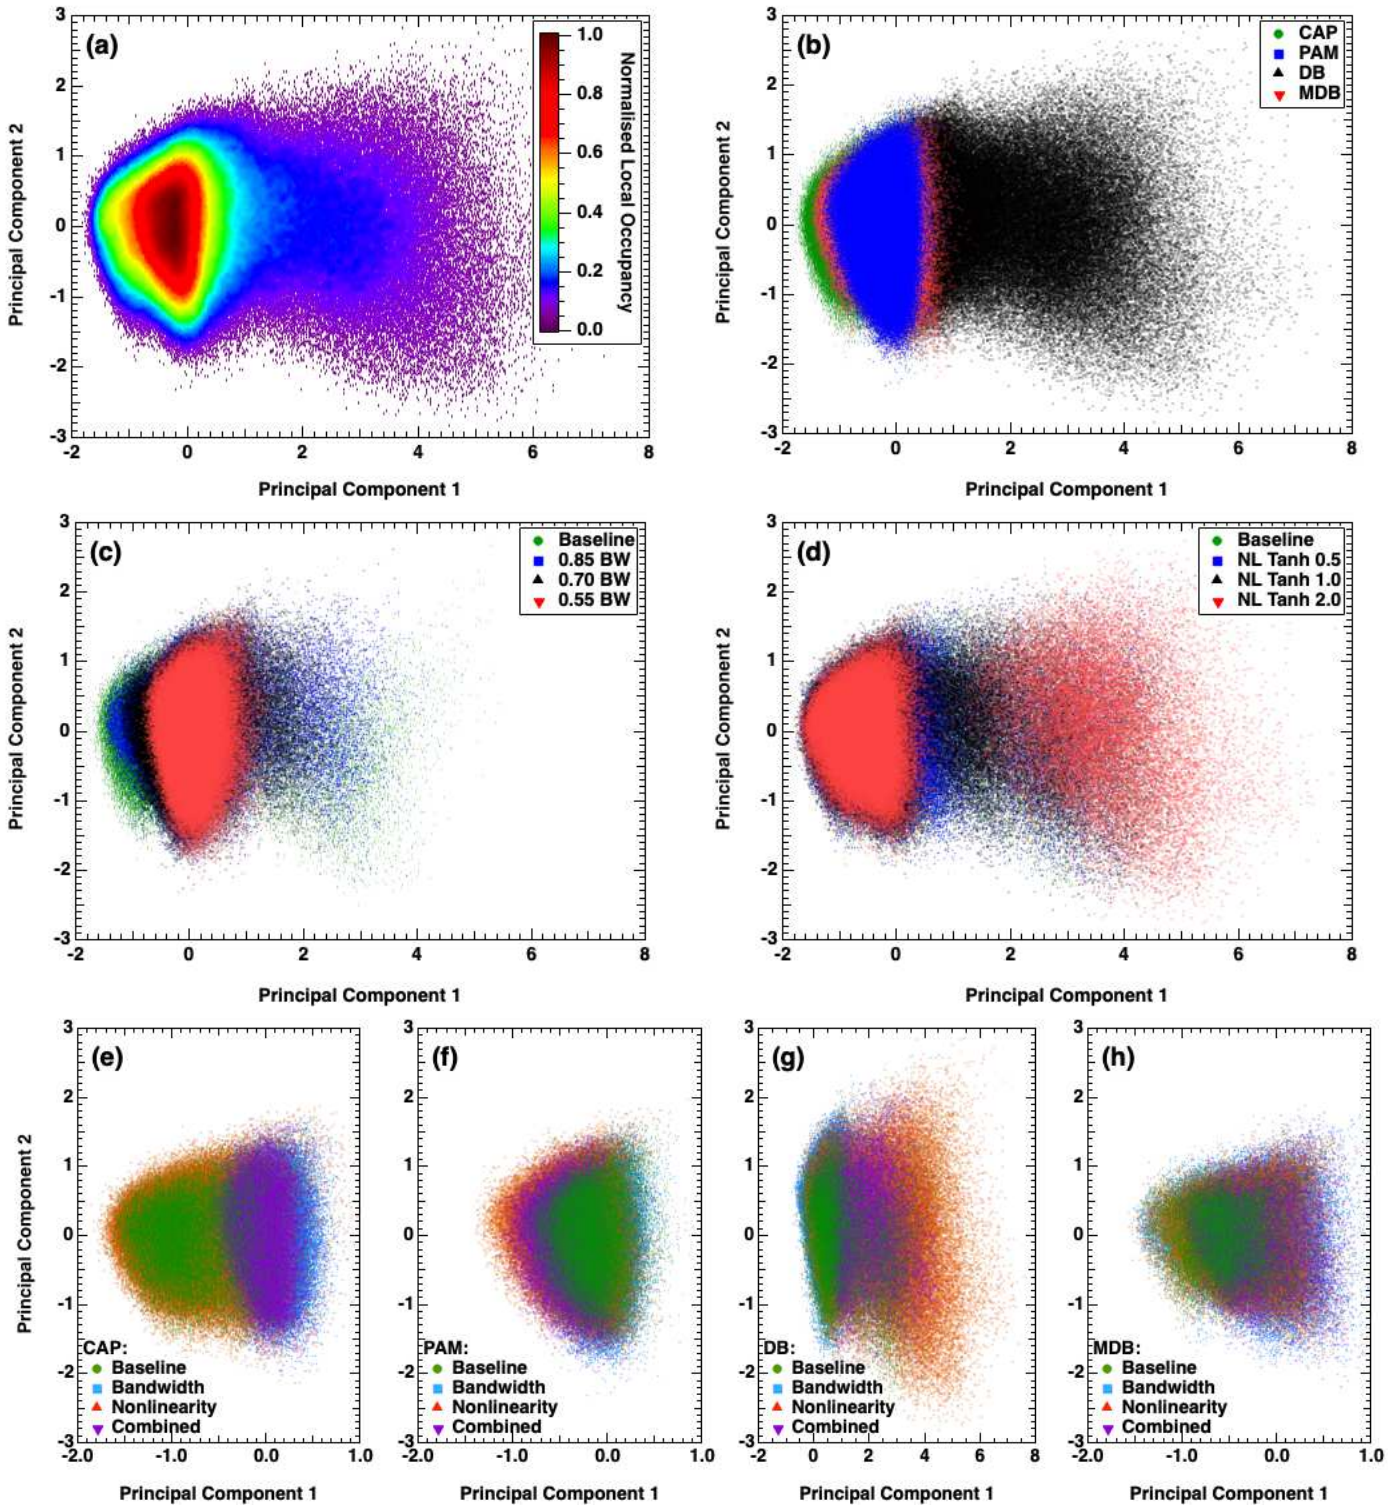


Figure 6 Organisation of the learned latent receiver state. This figure visualises the organisation of the learned receiver state. Principal component analysis (PCA) of the learned 64-dimensional latent receiver representation obtained from the deformable matched-filter network using 819,200 latent vectors pooled across all signalling formats and impairment conditions. **(a)** Global latent-space distribution coloured by the normalised local occupancy obtained from a Gaussian-smoothed two-dimensional occupancy estimate, normalised to the maximum occupancy and displayed after square-root compression. **(b)** Signalling-family organisation within the latent space. **(c,d)** Redistribution of the latent representation under progressively increasing bandwidth limitation and nonlinear distortion, respectively. **(e-h)** Family-specific latent-space organisation for CAP, PAM, DB and MDB signalling under baseline, bandwidth-limited, nonlinear and combined impairment conditions. Principal components were obtained from a single global PCA fitted to the complete latent dataset.

principally below approximately 0.1 normalised frequency, while the receiver avoids unstable inversion of the attenuated high-frequency region. The deformation therefore reports both the presence and the spectral form of matched-filter mismatch.

The comparison between PAM, DB and MDB provides a complementary test of the same idea (Figure 4). DB requires a large low-frequency deformation even under baseline conditions, showing that mismatch can be intrinsic to the interaction between signalling structure and the practical channel rather than caused solely by an externally imposed impairment. MDB suppresses the poorly supported low-frequency content at the transmitter and correspondingly reduces both EVM and learned deformation. The DMF therefore identifies the spectral component of mismatch removed by the modified signalling design. Importantly, deformation magnitude does not map universally onto performance gain: Figure **3** shows that CAP, PAM, DB and MDB occupy distinct deformation regimes despite sharing the same physical channel. Receiver mismatch must therefore be interpreted relative to signalling family and through the pattern of the deformation, not through a single pooled scalar metric.

The latent-state and coefficient-rank analyses extend this conclusion beyond deformation magnitude. In Figure **6**, signalling family provides the dominant organisation of the learned receiver state, while impairment classes overlap and produce systematic redistributions within those broader regimes. This is evidence of structured receiver configuration, not clean impairment classification. Figure 7 shows that the dimensionality of the coefficient space also varies across signalling families and operating conditions, indicating that some departures from matched-filter optimality require coordinated changes across multiple deformation components rather than a single gain or bandwidth adjustment.

The selected feature representation should be viewed as one physically motivated parameterisation of receiver state rather than a unique or optimal description. The present work deliberately favours compact, interpretable descriptors that maintain a transparent relationship between measurable waveform characteristics and the resulting matched-filter deformation. Alternative feature sets, automatically learned waveform embeddings, or hybrid representations could equally be incorporated within the proposed framework without changing the underlying deformable filtering methodology.

Taken together, Figures 1-7 form a linked argument. Figures 1 and 2 define the experimental mismatch landscape; Figure 3 shows that deformation and performance are related in family-specific regimes; Figure 4 validates the interpretation through the controlled DB-MDB comparison; Figure 5 identifies the spectral mechanism under bandwidth restriction; Figure 6 shows how the learned receiver state is organised; and Figure 7 quantifies the complexity of the required deformation. The agreement among these independent views is more important than any single performance metric because it shows that the learned correction contains reproducible information about receiver mismatch.

DMF should therefore not be viewed as an AI receiver, an adaptive equaliser, or a replacement for classical receiver theory. The waveform remains processed by a matched-filter receiver; the learned model only selects a constrained deformation of that filter. This preserves a physically interpretable reference point and makes each learned change expressible as a departure from the theoretical matched-filter solution.

In this sense, the learned deformation is a measurable representation of the gap between theoretical matched-filter optimality and practical receiver operation. Performance improvement is supporting evidence, but the principal contribution is the ability to observe where

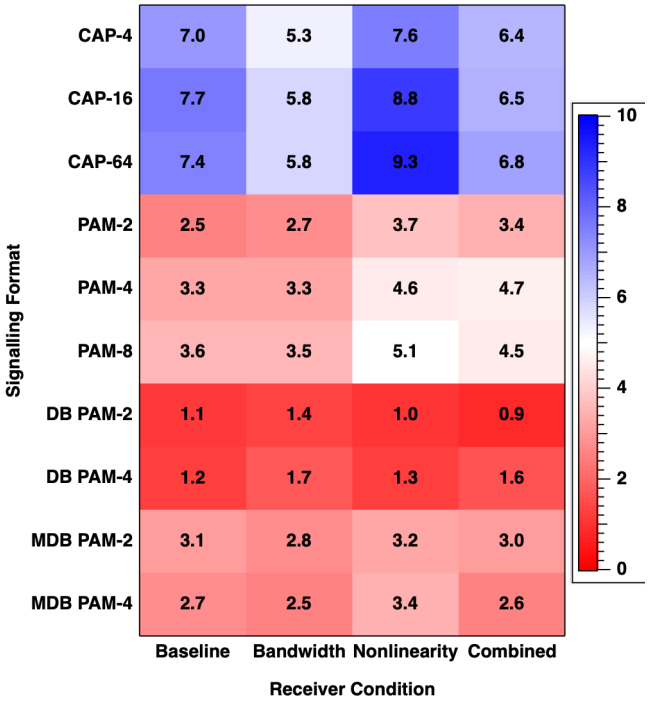


Figure 7 Complexity of the learned matched-filter deformation. This figure quantifies the dimensional complexity of the learned deformation. Mean effective coefficient rank of the learned DCT deformation coefficients for each signalling format under baseline, bandwidth-limited, nonlinear and combined impairment conditions. Effective coefficient rank was computed using the entropy effective-rank metric applied to the singular values of the coefficient matrix. Numerical values within each cell indicate arithmetic means across all evaluated parameter settings for the corresponding signalling format and receiver condition, while colour denotes the mean effective coefficient rank. Lower effective ranks indicate that the learned deformation is concentrated within fewer dominant deformation modes.

matched-filter assumptions fail, identify the spectral or structural form of the mismatch, and compare those departures across signalling families and channel conditions.

Limitations and future work

Several limitations define the scope of the present interpretation. First, the learned deformation depends on the engineered feature set and the 24-component DCT basis; alternative descriptors or bases may expose different, although related, representations of mismatch. Second, a non-zero deformation demonstrates departure from the nominal matched-filter solution but does not prove that the learned filter is the globally optimal receiver under every channel and noise model. Third, PCA provides a useful common projection of the latent state but is linear and does not preserve all high-dimensional relationships; overlap between impairment classes should therefore not be interpreted as absence of condition dependence, nor should visual separation be treated as a classifier-performance claim. Fourth, the experimental evidence is restricted to the investigated IM/DD optical link, signalling families and controlled impairment models, however speculatively DMF should transfer across communication mediums and domains.

Future work should test whether deformation signatures transfer across transmitters, photodetectors, link geometries and coherent or radio-frequency receivers; compare fixed DCT deformations with data-driven and physically derived bases; and examine uncertainty and non-uniqueness in the inferred deformation. Online or few-shot calibration could establish whether receiver-state trajectories can track hardware ageing or environmental drift. A further direction is to use the measured deformation to guide transmitter pulse shaping and signalling design, thereby closing the loop between diagnosis of matched-filter mismatch and mitigation at the waveform level.

The objective of this study is not to establish a universally superior receiver architecture relative to conventional adaptive equalisers or unconstrained learned filters. Such receivers optimise different objectives and do not, by construction, express their solution as a deformation of the theoretical matched filter. Comparative benchmarking of receiver architectures is therefore outside the principal scope here; the relevant contribution is an interpretable coordinate system for measuring how a practical matched-filter solution changes.

Although the engineered features were selected to capture diverse and physically interpretable characteristics of the received waveform, they do not represent the complete space of possible receiver descriptors. Future work will investigate systematic feature-selection strategies, sensitivity of the learned deformation to alternative descriptor sets, and fully learned latent waveform representations to determine the minimum information required to predict matched-filter deformation. Extending the framework to coherent optical systems, radio-frequency communication links and online receiver adaptation also represents promising directions for future research

Methods

Experimental free-space optical link

Experiments were performed using an IM/DD free-space optical communication system comprising of a 1.25 GS/s digital-to-analogue and analogue-digital converter (Analog Devices FMC-DAQ3) hosted on a Xilinx ZCU102 field programmable gate array, optical transmitter (850 nm TT Electronics OPV310 hosted on Thorlabs LDM56 laser mount and TED200C temperature controller (set to 25°C), free-space propagation path (200 mm path length) and photodetector receiver (Newport 818-BB-21A), with a system bandwidth of ~600 MHz. The signals were transmitted with a signal bandwidth of 100 MHz. All reported results were obtained from experimentally acquired waveforms Optical attenuation was controlled using calibrated neutral-density filters (Thorlabs NDK01 series). The same physical link and receiver procedures were used throughout the campaign; full apparatus details follow the methods reported in [25].

Signal generation

Four signalling families were investigated:

- PAM-2, PAM-4 and PAM-8
- CAP-4, CAP-16 and CAP-64
- Duobinary PAM-2 and PAM-4
- Modified duobinary PAM-2 and PAM-4

Pulse shaping employed root-raised-cosine matched filters with a nominal filter length of 192 taps and four samples per symbol. CAP signalling was implemented using orthogonal in-phase and quadrature pulse shapes derived from root-raised-cosine prototype filters modulated onto orthogonal carriers [28, 29].

To isolate specific sources of receiver mismatch, controlled impairments were introduced during signal generation. Bandwidth-limited conditions used low-pass-filter factors of 0.85, 0.70 and 0.55 relative to the nominal signal bandwidth. Nonlinear conditions used hyperbolic-tangent transfer functions with parameters 0.50, 1.00 and 2.00, and combined conditions applied the 0.70 bandwidth factor together with the 1.00 nonlinearity setting. Duobinary and modified-duobinary signalling were generated using the first-order partial-response polynomials $H_{DB}(z) = 1 + z^{-1}$ and $H_{MDB}(z) = 1 - z^{-1}$, respectively [30, 31].

Deformable matched filter formulation

The receiver is based on a CMF (h_{CMF}) augmented by a learned deformation:

$$h_{DMF} = h_{CMF} + \Delta h$$

where Δh denotes the learned filter deformation. Rather

than learning a completely unconstrained receive filter, the deformation is represented using a low-dimensional orthonormal basis expansion:

$$\Delta h = Bc = \sum_{i=1}^{M} c_i \, \phi_i$$

where B is a basis matrix, ϕ_i are basis functions and c_i are learned coefficients. A discrete cosine transform (DCT) basis was used throughout this work [32]. The basis comprised 24 orthonormal components spanning a 192-tap receive filter.

This formulation constrains the receiver to remain close to the nominal matched-filter solution while permitting systematic departures when required by the received signal conditions.

Feature-driven deformation estimation

The deformation coefficients were estimated using a KAN. A compact set of twenty-four physically interpretable waveform descriptors was extracted to characterise the receiver state.

Unlike conventional machine-learning receivers, the proposed network never operates directly on the communication waveform. Instead, each received waveform is transformed into a compact feature vector describing the receiver state. The KAN predicts only the deformation coefficients required to configure the matched filter, while symbol recovery continues to be performed entirely by the matched-filter receiver.

The feature vector contains 24 engineered descriptors extracted from each received waveform: root-mean-square amplitude, variance, skewness, kurtosis, peak-to-average power ratio, crest factor, zero-crossing rate, mean envelope, envelope standard deviation, maximum envelope, spectral entropy, spectral flatness, spectral centroid, spectral spread, 85% spectral roll-off, effective bandwidth, high-frequency energy ratio, autocorrelation at lags 1, 8 and 16, the 95th and 5th amplitude percentiles, mean absolute first difference and standard deviation of the first difference. Spectral features were calculated from a 4096-point real FFT after normalising the power spectrum to unit sum.

Consequently, the network acts only on receiver parameters rather than on the communication waveform itself. Changes in network output therefore correspond directly to changes in matched-filter shape, preserving a physically interpretable connection between the learned deformation and receiver behaviour [33].

Feature selection rationale

The engineered feature set was not intended to constitute either a minimal or exhaustive description of the received waveform. Instead, the descriptors were selected to span complementary physical characteristics known to influence departures from classical matched-filter optimality, including temporal statistics, amplitude distributions, transition behaviour, energy measures and spectral characteristics. The full table is provided in the supplementary information. Collectively, these descriptors provide an interpretable representation of the principal distortion mechanisms encountered under bandwidth limitation, nonlinear transfer characteristics and correlative signalling. The proposed DMF framework is independent of the specific feature definitions, and alternative or expanded feature sets could be incorporated without modification to the receiver architecture.

Kolmogorov-Arnold network architecture

The KAN encoder contained two hidden layers of 64 units. Each layer applied a channel-wise learnable piecewise-linear spline over 15 uniformly spaced control points between -3 and 3, followed by a learned linear projection and rectified-linear activation. The spline functions were identity-initialised. Layer normalisation was applied to the 24-dimensional input feature vector.

The network accepted the 24-dimensional feature vector and produced a 64-dimensional latent receiver state. Two independent linear heads mapped this state to 24 deformation coefficients for the in-phase and quadrature matched-filter branches. The coefficients were reconstructed into 192-tap filter deformations through the DCT basis expansion.

The reconstructed deformation was added to the nominal matched filter, and the resulting filter was normalised to unit ℓ_2 energy:

$$h_{DMF} = \frac{h_{DMF}}{\left|\left|h_{DMF}\right|\right|_2 + \varepsilon}$$

where $\varepsilon = 10^{-12}$ prevents numerical instability. This normalisation constrains the optimisation to learn only the shape of the matched-filter deformation, ensuring that performance improvements arise from changes in the filter response rather than arbitrary gain scaling

Training and evaluation

Training was performed using experimentally acquired waveform datasets. Data were partitioned into training, validation and test subsets using fractions of 70%, 15% and 15%, respectively, with random seed 1234. The network was optimised for 50 epochs using AdamW with learning rate 1×10^{-3}, weight decay 1×10^{-4} and batch size 32. The objective combined receiver EVM with smoothness, curvature, coefficient-magnitude and stability terms weighted 10^{-3}, 10^{-4}, 10^{-5} and 10^{-4}, respectively. To improve robustness to small measurement variations, zero-mean Gaussian noise ($\sigma = 0.01$) was added to the engineered feature vector during training.

Receiver performance was quantified using EVM. For

both CMF and DMF, symbol estimates were obtained by matched filtering, symbol-rate sampling and timing alignment using identical receiver procedures. Performance comparisons therefore isolate the effect of filter deformation while keeping all other receiver operations unchanged.

Latent-state analysis and visualisation

The 64-dimensional latent receiver state was extracted from the output of the second KAN hidden layer. Up to 2,000 examples per experimental condition were selected using random seed 0, yielding 819,200 latent vectors. A single global principal-component analysis was fitted after mean-centring the pooled latent matrix using singular-value decomposition [34]. The first three components explained 27.15%, 9.44% and 8.25% of the variance, respectively. Condition centroids and dispersions were calculated only after projection into this common coordinate system.

For Figure 6, relative local latent-space occupancy was estimated on a 512×384 two-dimensional histogram over PC1 and PC2. The histogram was Gaussian-smoothed with σ = 2 bins, and the smoothed values were bilinearly interpolated back to every latent vector. Colour values were normalised by the global maximum and square-root compressed for display, $D_{display} = \sqrt{\frac{D_{raw}}{\max(D_{raw})}}$. The colour scale therefore represents relative local occupancy rather than an absolute probability density.

Coefficient effective rank and statistical analysis

For each format, impairment, attenuation and receiver branch, coefficient vectors were pooled across sessions and mean-centred. Effective rank was calculated from the singular values σ_i of the coefficient matrix using the entropy definition $r_{eff} = \exp -\Sigma_i p_i \ln p_i$, where $p_i = \frac{\sigma_i^2}{\Sigma_j \sigma_j^2}$ [35]. This continuous measure equals one for a strictly one-dimensional coefficient distribution and increases as variance is distributed across additional orthogonal coefficient directions. To control computational cost, a maximum of 2,000 coefficient vectors per condition and branch was sampled using random seed 0.

Unless otherwise stated, condition-level summaries are arithmetic means over independent sessions. Ninety-five percent confidence intervals in summary tables were computed across session-level observations. Relationships between deformation and performance were assessed using Pearson correlation, Spearman rank correlation and partial correlation controlling for neutral-density attenuation. No hypothesis test was used to define whether DMF was beneficial; effect sizes and their distributions are reported because the central purpose is to characterise receiver mismatch rather than to establish a binary treatment effect.

Deformation metrics

The learned deformations were characterised using deformation energy, relative deformation energy and coefficient energy.

Deformation energy was defined as:

$$E_D = \|\Delta h\|_2^2$$

and quantifies the overall departure from the nominal matched filter.

Relative deformation energy was defined as:

$$E_R = \frac{\| \Delta h \|_2^2}{\| h_{\mathrm{CMF}} \|_2^2}$$

and provides a scale-independent measure of receiver mismatch.

Coefficient energy was defined as:

$$E_C = \sum_i c_i^2$$

and quantifies the magnitude of the latent deformation required to generate the observed filter modification.

These metrics were used throughout the study to compare receiver mismatch across signalling formats, impairment classes and optical attenuation levels. Because the deformable matched filter is defined relative to the classical matched-filter solution, these quantities provide direct measurements of departures from matched-filter optimality.

Data availability

The data used in this study is available at: https://github.com/qmul-optocomms/kan-dmf-public-data

Code availability

The code used in this study is available at: https://github.com/qmul-optocomms/kan-dmf-public-data

Acknowledgements

No funding was received for this research.

Author contributions

PAH is the sole author of the work. PAH conceptualised the paper, developed the code, performed the laboratory experiments, analysed the results, prepared the figures and wrote and reviewed the manuscript.

Competing interests

The author declares no competing interests.

References

[1] G. Turin, "An introduction to matched filters," *IRE transactions on Information theory,* vol. 6, no. 3, pp. 311–329, 1960.

[2] H. Nguyen, I. Al, and Y. M. Jang, "Survey of next-generation optical wireless communication technologies for 6G and Beyond 6G," *ICT Express,*

vol. 11, no. 3, pp. 576–589, 2025/06/01/ 2025, doi: https://doi.org/10.1016/j.icte.2025.04.006.
[3] P. H. Pathak, X. Feng, P. Hu, and P. Mohapatra, "Visible light communication, networking, and sensing: A survey, potential and challenges," *IEEE communications surveys & tutorials,* vol. 17, no. 4, pp. 2047–2077, 2015.
[4] J. M. Kahn and J. R. Barry, "Wireless infrared communications," *Proceedings of the IEEE,* vol. 85, no. 2, pp. 265–298, 1997.
[5] S. U. Qureshi, "Adaptive equalization," *Proceedings of the IEEE,* vol. 73, no. 9, pp. 1349–1387, 1985.
[6] D. Maiti and M. Brandt-Pearce, "Modified Nonlinear Decision Feedback Equalizer for Long-Haul Fiber-Optic Communications," *Journal of Lightwave Technology,* vol. 33, no. 18, pp. 3763–3772, 2015, doi: 10.1109/JLT.2015.2444273.
[7] N. Stojanovic, F. Karinou, Z. Qiang, and C. Prodaniuc, "Volterra and Wiener Equalizers for Short-Reach 100G PAM-4 Applications," *Journal of Lightwave Technology,* vol. 35, no. 21, pp. 4583–4594, 2017/11/01 2017. [Online]. Available: https://opg.optica.org/jlt/abstract.cfm?URI=jlt-35-21-4583.
[8] V. Lauinger, F. Buchali, and L. Schmalen, "Blind Equalization and Channel Estimation in Coherent Optical Communications Using Variational Autoencoders," *IEEE Journal on Selected Areas in Communications,* vol. 40, no. 9, pp. 2529–2539, 2022, doi: 10.1109/JSAC.2022.3191346.
[9] M. Sana and E. C. Strinati, "Semantic Channel Equalizer: Modelling Language Mismatch in Multi-User Semantic Communications," in *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*, 4–8 Dec. 2023 2023, pp. 2221–2226, doi: 10.1109/GLOBECOM54140.2023.10437371.
[10] P. J. Freire, A. Napoli, B. Spinnler, N. Costa, S. K. Turitsyn, and J. E. Prilepsky, "Neural Networks-Based Equalizers for Coherent Optical Transmission: Caveats and Pitfalls," *IEEE Journal of Selected Topics in Quantum Electronics,* vol. 28, no. 4: Mach. Learn. in Photon. Commun. and Meas. Syst., pp. 1–23, 2022, doi: 10.1109/JSTQE.2022.3174268.
[11] M. A. Amirabadi, S. A. Nezamalhosseini, M. H. Kahaei, and L. R. Chen, "A Comprehensive Survey on Machine and Deep Learning for Optical Communications," *IEEE Access,* vol. 13, pp. 88794–88846, 2025, doi: 10.1109/ACCESS.2025.3569559.
[12] S. R. Doha and A. Abdelhadi, "Deep Learning in Wireless Communication Receivers: A Survey," *IEEE Access,* vol. 13, pp. 113586–113605, 2025, doi: 10.1109/ACCESS.2025.3584000.
[13] Y. Osadchuk, D. Kong, D. Zibar, and F. D. Ros, "Experimental Investigation of Nonlinear Vector Autoregressor-Based Channel Equalization for IM/DD Transmission With Silicon Microring Modulators," *Journal of Lightwave Technology,* vol. 43, no. 3, pp. 1130–1139, 2025, doi: 10.1109/JLT.2024.3478743.
[14] S. F. Nielsen, F. D. Ros, M. N. Schmidt, and D. Zibar, "End-to-End Learning of Transmitter and Receiver Filters in Bandwidth Limited Fiber Optic Communication Systems," *Journal of Lightwave Technology,* vol. 43, no. 8, pp. 3749–3760, 2025, doi: 10.1109/JLT.2025.3528542.
[15] M. Honkala, D. Korpi, and J. M. J. Huttunen, "DeepRx: Fully Convolutional Deep Learning Receiver," *IEEE Transactions on Wireless Communications,* vol. 20, no. 6, pp. 3925–3940, 2021, doi: 10.1109/TWC.2021.3054520.
[16] Y. Xu, X. Guan, W. Jiang, X. Wang, W. Hu, and L. Yi, "Low-complexity end-to-end deep learning framework for 100G-PON," *J. Opt. Commun. Netw.,* vol. 16, no. 11, pp. 1093–1103, 2024/11/01 2024, doi: 10.1364/JOCN.532742.
[17] A. Sun *et al.*, "End-to-End Deep-Learning-Based Photonic-Assisted Multi-User Fiber-mmWave Integrated Communication System," *Journal of Lightwave Technology,* vol. 42, no. 1, pp. 80–94, 2024, doi: 10.1109/JLT.2023.3304006.
[18] S. Xing *et al.*, "End-to-end deep learning for a flexible coherent PON with user-specific constellation optimization," *J. Opt. Commun. Netw.,* vol. 16, no. 1, pp. 59–70, 2024/01/01 2024, doi: 10.1364/JOCN.500500.
[19] S. Cammerer *et al.*, "Sionna Research Kit: A GPU-Accelerated Research Platform for AI-RAN," in *2025 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN)*, 26–29 May 2025 2025, pp. 1–2, doi: 10.1109/ICMLCN64995.2025.11140427.
[20] Q. Liu, M. Fu, X. Liu, L. Yi, W. Hu, and Q. Zhuge, "End-to-End Joint Digital and Optical Signal Processing Enabled by Interpretable Deep Learning for Coherent Optical Communication Systems," in *2024 Optical Fiber Communications Conference and Exhibition (OFC)*, 24–28 March 2024 2024, pp. 1–3.
[21] I. Roumpos *et al.*, "Physics-inspired End-to-End Deep Learning for High-Performance Optical Fiber Transmission Links," in *CLEO 2023*, San Jose, CA, 2023/05/07 2023: Optica Publishing Group, in Technical Digest Series, p. SF1F.6, doi: 10.1364/CLEO_SI.2023.SF1F.6. [Online]. Available: https://opg.optica.org/abstract.cfm?URI=CLEO_SI-2023-SF1F.6
[22] Y. Osadchuk *et al.*, "Low-Complexity Samples Versus Symbols-Based Neural Network Receiver for Channel Equalization," *Journal of Lightwave Technology,* vol. 42, no. 15, pp. 5167–5174, 2024, doi: 10.1109/JLT.2024.3390227.
[23] Z. Liu *et al.*, "KAN: Kolmogorov–arnold networks," in *International conference on learning representations*, 2025.
[24] T. Zheng *et al.*, "A KAN-Based Decoupled Equalization Framework Towards Latency-Optimized 6G Optical Transport Systems," *Journal of Lightwave Technology,* vol. 44, no. 13, pp. 5364–5371, 2026/07/01 2026, doi: 10.1364/JLT.44.005364.

[25] P. A. Haigh, "ML-Enabled Deformable Matched Filters for Band-Limit Compensation in Free-Space Optics," *npj Wireless Technology,* 2026.
[26] M. I. Olmedo *et al.*, "Multiband carrierless amplitude phase modulation for high capacity optical data links," *Journal of Lightwave Technology,* vol. 32, no. 4, pp. 798–804, 2013.
[27] R. Rodes *et al.*, "Carrierless amplitude phase modulation of VCSEL with 4 bit/s/Hz spectral efficiency for use in WDM-PON," *Optics express,* vol. 19, no. 27, pp. 26551–26556, 2011.
[28] B. R. Saltzberg, "Carrierless Amplitude–Phase Modulation," *Wiley Encyclopedia of Telecommunications,* 2003.
[29] N. Bamiedakis, R. Penty, and I. White, "Carrierless amplitude and phase modulation in wireless visible light communication systems," *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences,* vol. 378, no. 2169, p. 20190181, 2020.
[30] A. Lender, "The duobinary technique for high-speed data transmission," *IEEE Transactions on Communication and Electronics,* vol. 82, no. 2, pp. 214–218, 1963.
[31] E. Kretzmer, "Generalization of a techinque for binary data communication," *IEEE Transactions on Communication Technology,* vol. 14, no. 1, pp. 67–68, 1966.
[32] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE transactions on Computers,* vol. 100, no. 1, pp. 90–93, 1974.
[33] Z. C. Lipton, "The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery," *Queue,* vol. 16, no. 3, pp. 31–57, 2018.
[34] H. Abdi and L. J. Williams, "Principal component analysis," *WIREs Computational Statistics,* vol. 2, no. 4, pp. 433–459, 2010/07/01 2010, doi: https://doi.org/10.1002/wics.101.
[35] O. Roy and M. Vetterli, "The effective rank: A measure of effective dimensionality," in *2007 15th European Signal Processing Conference*, 3–7 Sept. 2007 2007, pp. 606–610.

Supplementary Information

No.	Feature	Definition
1	RMS amplitude	sqrt(mean(x^2))
2	Variance	var(x)
3	Skewness	mean[$((x-\mu)/\sigma)^3$]
4	Kurtosis	mean[$((x-\mu)/\sigma)^4$]
5	Peak-to-average power ratio	max(x^2)/mean(x^2)
6	Crest factor	max|x|/RMS
7	Zero-crossing rate	Mean sign-bit transition rate
8	Mean envelope	mean|x|
9	Envelope standard deviation	std|x|
10	Maximum envelope	max|x|
11	Spectral entropy	$-\Sigma P \log_2 P$
12	Spectral flatness	geometric mean(P)/arithmetic mean(P)
13	Spectral centroid	ΣfP
14	Spectral spread	sqrt[$\Sigma(f-\text{centroid})^2 P$]
15	85% spectral roll-off	Frequency containing 85% cumulative spectral power
16	Effective bandwidth	Fraction of FFT bins above half the peak normalised power
17	High-frequency energy ratio	Power above normalised frequency 0.5 / total power
18	Autocorrelation lag 1	corr(x_t,x_(t+1))
19	Autocorrelation lag 8	corr(x_t,x_(t+8))
20	Autocorrelation lag 16	corr(x_t,x_(t+16))
21	95th percentile	P95(x)
22	5th percentile	P5(x)
23	Mean absolute first difference	mean|Δx|
24	First-difference standard deviation	std(Δx)

Category	Examples	Purpose
Temporal statistics	Mean, variance, skewness, kurtosis	Capture distributional changes
Transition descriptors	Zero crossings, transition density, run lengths	Capture ISI and correlative signalling
Energy descriptors	RMS, peak-to-average ratio, energy measures	Capture amplitude distortions
Spectral descriptors	Spectral centroid, bandwidth, spectral moments	Capture bandwidth limitation
Correlation descriptors	Autocorrelation and related measures	Capture waveform memory