Photonic Neural Network Frontend for Automatic Modulation Classification

Sergei Kudriavtcev

Advisor: Professor Paul Prucnal

Graduate Mentor: Hyuma Umeda

Bachelor of Science in Engineering

Department of Electrical and Computer Engineering

Princeton University

May 2026

I hereby declare that I am the sole author of this thesis.

I authorize Princeton University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

_____

Sergei Kudriavtcev

I further authorize Princeton University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

_____

Sergei Kudriavtcev

Abstract

Automatic Modulation Classification (AMC) is the problem of identifying the modulation scheme of an incoming radio signal without prior coordination with the transmitter. It is central to a wide range of applications including cognitive radio systems (5G/6G, Wi-Fi), electronic warfare, and wideband spectrum sensing, yet the gigahertz bandwidths of modern signals push conventional digital receiver chains against practical limits in sampling rate, power, and latency. Silicon photonic neural networks offer a fundamentally different substrate: they perform analog computation directly on the optical signal before any digitization, at near-zero marginal energy per operation and with no sampling bottleneck. This work validates in software simulation a three-stage photonic AMC pipeline motivated by Lederman (2025): a photonic neural network frontend modeled under realistic hardware constraints, an integrator, and a lightweight digital classifier, all trained jointly end-to-end.

We show that a 2-3-2 topology containing only 15 micro-ring resonators (MRRs) achieves 97% classification accuracy across four modulation classes (OOK, BPSK, QPSK, 8PSK) at 30 dB SNR, a topology realizable on silicon photonic chips fabricated today. We further show that the optical nonlinearity is essential: a purely linear photonic frontend collapses to 60% accuracy. Sweeping over network size and channel conditions, we find that a modest scaling to a 2-8-2 topology (40 MRRs) recovers near-perfect accuracy at 20 dB SNR and achieves 95.3% at 10 dB SNR, a significant improvement over the 76.6% of the minimal topology at the same noise level. Scaling to a 2-8-4 topology (56 MRRs) improves accuracy under wireless fading from 46.4% to 68%, with robustness to channel impairments identified as the central open problem motivating future work.

Acknowledgements

I am deeply grateful to Professor Paul Prucnal for the opportunity to work in the Lightwave Communications Lab and for his direction on this project. I thank Hyuma Umeda for many hours of patient mentorship and for being the steady hand that made this work possible. I am also grateful to the broader Lightwave Lab group and the Princeton ECE department for the resources and environment that supported this independent work.

for posterity

Contents

Chapter 1

Introduction

Wireless communication systems encode information onto a carrier signal through a *modulation scheme*, which determines how bits are mapped onto the amplitude, phase, or frequency of the transmitted waveform. *Automatic Modulation Classification* (AMC) is the problem of identifying which modulation scheme is in use from the received signal alone, without prior coordination with the transmitter, in real time and under unknown channel conditions.

This problem arises constantly in practice. In 5G and emerging 6G systems, cognitive radio nodes scan a wide band of frequencies, classify the modulations occupying each channel, and identify spectral gaps on which to transmit, maximizing spectral efficiency in dense wireless environments. In electronic warfare, a receiver must identify and classify enemy transmissions in real time, as adversaries actively vary their modulation scheme to evade detection. Satellite and airborne platforms face the same classification problem across wide stretches of spectrum simultaneously, operating in contested or unknown electromagnetic environments while additionally placed under the constraint of power.

Significant progress has been made over recent decades in digital receiver chains, with modern ADCs achieving ever-higher sampling rates. However, these advances

are approaching a fundamental physical ceiling: the power, bandwidth, and latency demands of gigahertz-scale signal classification cannot be resolved by incremental improvements to digital hardware alone, motivating a fundamentally different computational substrate: photonic.

1.1 The Limits of Digital Processing

The dominant approach to AMC today is digital: sample the received analog signal with an analog-to-digital converter (ADC), then run the digitized samples through a software algorithm or neural network that performs the classification. Recent deep-learning approaches have achieved impressive accuracy [1], but the approach faces a fundamental physical bottleneck when the signals of interest occupy gigahertz of bandwidth.

While modern ADCs have achieved impressive sampling rates in the multi-gigahertz range, a single ADC covers only a limited instantaneous bandwidth, meaning that monitoring wide stretches of spectrum requires either multiple parallel ADCs or additional down-conversion hardware to subdivide the band. High-precision high-speed ADCs also routinely dissipate hundreds of milliwatts to close to a watt at the conversion stage alone, and in systems that must monitor many channels simultaneously this power burden multiplies accordingly.

Beyond digitization, running a neural network classifier on every sample in a high-bandwidth stream compounds the power and latency burden further. In latency-sensitive applications such as electronic warfare, the delay introduced by buffering, inference, and decision output can render the classification result too late to be actionable.

1.2 Photonics as a Fundamental Improvement

Photonic signal processing offers a qualitatively different approach. Rather than converting the analog RF signal to digital and then computing on it, a photonic processor operates directly on the analog signal avoiding the aforementioned issues.

The key component is the *micro-ring resonator* (MRR), a small ring of silicon waveguide whose resonance can be tuned by an integrated heating element. When an optical carrier is modulated with an RF signal and passes through a bank of such rings, each ring performs a weighted transformation of the signal in the optical domain, at the speed of light and without the thermal noise floor or quantization error of a digital processor. A bank of MRRs therefore implements a neural network layer entirely in the optical domain, performing weighted summation and nonlinear activation as the signal propagates through the chip. [3, 4].

Photonic chips thus offer a compelling alternative. Operating directly on the analog optical signal, a single photonic chip can process bandwidths spanning tens to hundreds of gigahertz without the need for parallel digitization chains or down-conversion hardware. Computation occurs as the signal propagates through the device, introducing no sampling clock, buffering, or pipeline delay. Critically, photonic multiply-accumulate operations carry effectively zero marginal energy cost per operation, orders of magnitude below the hundreds of milliwatts dissipated by high-speed ADCs. Photonic processing is therefore not an incremental improvement over digital receiver chains but a fundamentally different substrate, one that is naturally suited to the wideband, low-power, low-latency regime where digital approaches are most constrained.

Chapter 2

Background

2.1 Prior Work in Automatic Modulation Classification

Classical AMC research predating the deep learning era is built on two approaches: *likelihood-based* methods, which test candidate modulation hypotheses using a likelihood ratio, and *feature-based* methods, which extract handcrafted statistics from the signal and feed them to a classifier. Likelihood-based methods suffer from high computational complexity and sensitivity to model mismatch such as carrier phase offsets and residual channel effects [2]. Swami and Sadler proposed a practical alternative: a hierarchical classifier built on fourth-order cumulants, which characterize the shape of the received signal constellation and are naturally robust to carrier phase and frequency offsets [2]. This cumulant-based approach is the classical precursor to the moment-integral architecture this work simulates, as both exploit the fact that higher-order statistics of the received signal are characteristic of its modulation type.

More recently, deep learning methods applied directly to raw IQ samples have demonstrated competitive AMC performance [1], though at the cost of substantially greater computational overhead and a dependence on large labeled training datasets.

2.2 Silicon Photonic Neural Networks at the Lightwave Lab

Silicon photonics provides an alternative foundation on which neural network can be built. The fundamental building block exploited in this work is the micro-ring resonator weight bank, in which a parallel array of silicon micro-rings, each tuned to a slightly different resonance wavelength, simultaneously imposes independent weights on a set of wavelength-multiplexed optical signals. The weighted signals are then combined and detected to yield an analog electrical sum, allowing for a fully optical multiply-and-accumulate operation. Tait et al. [3] introduced this neuromorphic photonic primitive, and the broader theoretical framework for photonic neural networks is laid out in Prucnal and Shastri's textbook [4]. The Lightwave Communications Lab at Princeton has developed silicon photonic chips implementing these ideas for a decade, and this work directly extends that line of research.

2.3 The Moment-Based Photonic AMC Architecture

The specific architecture this work simulates is discussed in Lederman's (2025) [5] work, stating that the structure of a moment-based modulation classifier maps cleanly onto a three-stage photonic system. A moment of a signal $s(t)$ is, by definition, an integral of some function of $s(t)$ over time. Different modulation schemes have characteristic higher-order cumulants and moments, which is why classical feature-based AMC methods work in the first place [2]. Thus this entire classification pipeline can be expressed simply as

$$y = f\left(\int g\big(s(t)\big)\, dt \right), \tag{2.1}$$

where $g(\cdot)$ is a nonlinear function applied to each instantaneous sample of the signal, the integral accumulates those samples over a time window, and $f(\cdot)$ maps the accumulated value to a class decision.

Crucially, the three pieces of equation (2.1) have direct physical realizations:

- $g(\cdot)$ is the bandwidth-demanding stage, applied at the full sample rate. It maps onto a small photonic neural network built from an MRR weight bank, which can run at the optical line rate.
- $\int \cdot \, dt$ is a simple accumulation, realizable with an analog integrator.
- $f(\cdot)$ is a low-rate decision function or classifier, operating on the accumulated statistic, runnings once per frame, not once per sample, and is therefore inexpensive to implement digitally.

What makes this mapping compelling is that the computationally demanding stage, $g(\cdot)$, is precisely what photonic hardware performs naturally, leaving only a low-rate digital decision at the output. The central question is then one of scale: how small a photonic network is sufficient, and how does the system behave under the hardware constraints and channel conditions of a real deployment?

2.4 Thesis and Contributions

The argument of this paper is twofold. First, it validates, in a controlled software simulation, that the three-stage photonic AMC pipeline discussed by Lederman (2025) is sufficient for accurate classification across four modulation classes. Second, by showing that a topology containing only 15 MRRs is enough to reach 97% accuracy in clean conditions, it makes the stronger claim that this design is not merely theoretically motivated but practically achievable with the photonic chip technology fabricated in research labs today. Thus shifting the question away from if photonic AMC system

could be built, but rather what its accuracy and robustness look like under realistic constraints, and where the engineering work should be focused next.

The contributions of this work are:

- A complete software simulation of a three-stage photonic AMC pipeline that enforces realistic hardware constraints, including DAC quantization and weight clamping.
- An end-to-end joint training procedure that propagates gradients through the quantized photonic weights using a Straight-Through Estimator (STE), producing trained weight values that can be utilized to tune the components of the physical chips.
- An architecture sweep over five photonic NN topologies (4 to 56 MRRs), establishing that nonlinearity is essential and that 15 MRRs is the minimum size that achieves $\geq 95\%$ accuracy in clean conditions.
- A robustness study across three SNR levels (10, 20, 30 dB) and three channel impairment levels (clean, cabled, wireless), characterizing where the system performs well and where an increased arcitecture benefits.

The remainder of this paper is organized as follows. The system architecture is described in detail in Chapter 3, covering the photonic NN model, hardware constraints, and joint training procedure. Chapter 4 documents the data generation setup, the modulation classes under study, and the channel impairment conditions evaluated. Chapter 5 presents the results of the architecture sweep and robustness studies. Chapter 6 discusses the implications of these results and the limitations of the simulation. Chapter 7 outlines future work and Chapter 8 concludes.

Chapter 3

System Architecture

3.1 Pipeline Overview

The full classification pipeline consists of three computational stages, shown in Figure 3.1. The photonic neural network processes each incoming IQ sample at the full analog sample rate, producing a small feature vector whose dimension equals the output size of the photonic NN topology (2 for the 2-3-2 and 2-8-2 models, 4 for the 2-8-4 model). The integrator accumulates these per-sample feature vectors across an entire frame, producing a single frame-level feature vector of the same dimension by summing each feature independently. Finally, a small digital classifier maps this integrated feature vector to one of four modulation classes, evaluating once per frame. The three stages correspond directly the moment-integral expression in equation (2.1), with the data generation process described separately in Chapter 4.

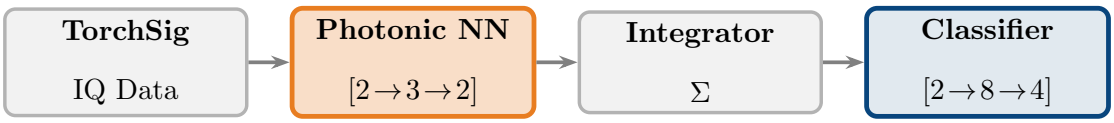

Figure 3.1: The three-stage photonic AMC pipeline. The Photonic NN produces a 2D feature vector per IQ sample, the Integrator sums these across 1024 samples per frame, and the Classifier evaluates once per frame.

3.2 Micro-Ring Resonator Weight Banks

A micro-ring resonator (MRR) is a small ring of silicon waveguide, typically a few microns in diameter, that supports optical resonances at a discrete set of wavelengths determined by the ring's circumference and refractive index. When light at a resonant wavelength enters the ring, it builds up coherently and is preferentially absorbed or redirected; light at off-resonance wavelengths passes through largely unaffected. The transmission through the ring is therefore a strong function of the wavelength relative to the ring's resonance.

An MRR *weight bank* exploits this wavelength selectivity to implement a vector of programmable analog weights. Multiple input signals are encoded onto distinct optical wavelengths and sent through a parallel array of rings; each ring is tuned to a different reference wavelength, so that ring i imposes a programmable transmission coefficient on signal i. The transmitted signals are then summed photonically (combined into a single waveguide and detected at a photodiode), yielding an analog optical multiply-and-accumulate operation. A bank of n MRRs therefore implements a single weighted-sum neuron with n inputs; layering multiple banks produces a photonic neural network.

3.3 The Photonic Neural Network

The photonic NN simulated in this work is a fully connected feedforward network with input dimension 2 (the I and Q components of a single sample) and output dimension of 2 reflecting the two extracted features (increased to 4 in the 2-8-4 model). The intermediate hidden layer dimension is varied across experiments to study the effect of network capacity. The smallest nontrivial nonlinear topology, denoted 2-3-2, has three hidden units; larger topologies 2-8-2 and 2-8-4 are also evaluated. A purely linear baseline 2-2 (no hidden layer) is included as a control to isolate the contribution

of the nonlinearity.

Hardware-constrained weights. Each weight in the photonic NN corresponds to a single MRR. The weight itself is set by passing a current through a thin-film resistive heater integrated above the ring, which thermo-optically shifts the ring's resonance and therefore the optical transmission. The current is produced by a digital-to-analog converter (DAC), which means the achievable weights live on a finite, quantized grid. To reflect this in the simulation, each weight is clamped to the interval $[-1, 1]$ and quantized to 4096 levels (12-bit DAC resolution). The clamp models the finite tuning range of the heater, and the quantization models the resolution of the DAC controlling it.

MRR Count. The MRR count for each topology is determined by the total number of weights across all layers plus the nonlinear activation units implementing the ReLU response. For the 2-3-2 topology this gives $2 \times 3 + 3 + 3 \times 2 = 15$ MRRs; the full set of topologies is summarized in Table 5.1 in Chapter 5.

Nonlinear activation. The nonlinear response of an MRR arises from its Lorentzian transmission profile combined with the square-law optical-to-electrical conversion at the photodetector, producing an inherent nonlinearity that has been shown to approximate a ReLU function in silicon photonic neural network implementations [6]. In the simulation, this is modeled explicitly as a ReLU applied after the first layer. The second layer is linear, with no activation, so that the integrator downstream sees unsaturated feature values.

Gradient flow through quantization. Quantization is mathematically a step function with derivative zero almost everywhere, which would make the photonic weights untrainable under standard backpropagation. We use the Straight-Through

Estimator (STE) [8], implemented as a custom `torch.autograd.Function`: in the forward pass the network sees the quantized weights, so its behavior matches the constrained hardware, but in the backward pass the gradient is passed through the quantization operation as if it were the identity. This allows the optimizer to move the underlying continuous weights toward better quantized neighbors.

3.4 The Integrator

A photonic chip computes instantaneously: the optical signal that emerges from the MRR weight bank reflects only the input present at that moment thus the chip has no memory and cannot directly compute any temporal statistic. The integrator is the architectural element that supplies this missing capability. The integrator accumulates the per-sample output vectors of the photonic NN across an entire frame, summing each feature dimension independently to produce a single feature vector per frame. In the moment-integral expression of equation (2.1), this corresponds to the $\int \cdot\, dt$ operation. In hardware, this stage is realized as an analog integrating photodetector, implemented in simulation as a simple cumulative sum using a custom `torch.nn.Module`.

The integrator establishes a $1024 : 1$ ratio between the rate at which the photonic NN must operate and the rate at which the classifier is called: the photonic NN sees every IQ sample, while the classifier sees only one feature vector per frame. This rate gap is what makes the overall system efficient. The expensive, bandwidth-demanding work is done in the analog optical domain, and the comparatively expensive digital classifier only ever runs at the much lower frame rate.

3.5 The Digital Classifier

The classifier is intentionally lightweight: a fully connected feedforward network of shape 2-8-4 (4-8-4 for larger 2-8-4 PNN model), trained with cross-entropy loss.

The classifier is permitted to be so simple due to the photonic NN and integrator together having already done the heavy lifting: by the time the classifier is invoked, the raw IQ frame has been projected into a compact feature space explicitly optimized for class separation through joint training. The classifier's sole task is to draw decision boundaries between the resulting clusters, which requires only a small MLP. It also runs at the frame rate rather than the sample rate, making its computational cost negligible relative to the analog frontend.

3.6 End-to-End Joint Training

The three trainable stages, photonic NN, integrator, and classifier, are trained jointly end-to-end using cross-entropy loss on the classifier output. The integrator is a fixed sum and contributes no parameters, but gradients flow through it back to the photonic NN weights. We use the Adam optimizer [9] with learning rate 1×10^{-3} for 200 epochs (300 epochs for certain high impairment tests).

Joint training is essential. If the photonic NN were trained independently to produce some hand-chosen statistic and the classifier were trained afterward on those frozen features, the photonic NN would have no way to know which projections of signal space the downstream classifier actually needs in order to discriminate between modulation types. Joint training closes this loop as the cross-entropy gradient tells the photonic NN exactly which features make the classes most separable, and the photonic NN shifts its weights accordingly.

Chapter 4

Data Generation and Experimental Setup

4.1 Modulation Classes

The four modulation classes evaluated in this work are OOK, BPSK, QPSK, and 8PSK. Each is a standard digital modulation scheme used in real communications systems, and they were chosen to span a useful range of classification difficulty while remaining simple enough to validate the smallest photonic topologies.

OOK (On-Off Keying). The simplest digital modulation: a bit value of 1 transmits the carrier at full amplitude, and a bit value of 0 transmits nothing. The constellation consists of two points on the real (in-phase) axis at amplitudes 0 and 1.

BPSK (Binary Phase-Shift Keying). Bits are encoded into the phase of the carrier: a 0-bit corresponds to phase 0 and a 1-bit to phase π. The constellation is two points on the real axis at ± 1. BPSK has the same number of constellation points as OOK but uses phase rather than amplitude, and so has very different higher-order moments.

QPSK (Quadrature Phase-Shift Keying). Two bits at a time are encoded into one of four phases ($\pi/4$, $3\pi/4$, $5\pi/4$, $7\pi/4$). The constellation is four equally spaced points on the unit circle.

8PSK (8-Phase Shift Keying). Three bits at a time are encoded into one of eight phases evenly spaced around the unit circle. 8PSK and QPSK share the same average power and the same support (the unit circle) and differ only in the density of points; this similarity is what makes them difficult to distinguish at low SNR.

The set was chosen to provide a graded difficulty curve: OOK and BPSK are easy to separate from each other and from the PSK family because their amplitude statistics are distinctive; QPSK and 8PSK live on the same circle and require the network to learn the difference between four-fold and eight-fold rotational symmetry. This set was chosen for its simplicity, making it well suited for validating a minimal photonic topology, while still presenting non-trivial classification challenges due to the overlap between QPSK and 8PSK.

4.2 TorchSig and the Choice of Synthetic Data

This work uses TorchSig [7] rather than the more widely used RadioML dataset [1]. RadioML was generated by transmitting modulated signals through a real software-defined radio (SDR) and re-capturing them, which embedded hardware-specific fingerprints of the recording device into every example. Classifiers trained on this data risk exploiting these device-specific artifacts rather than true modulation features, which raises questions about the generalizability of published accuracy numbers to other hardware and recording conditions. Additionally, certain modulation classes in RadioML have been found to contain effectively only noise rather than the labeled signal, a generation artifact that undermines the reliability of the dataset for

benchmarking.[1] TorchSig addresses these concerns by generating signals from explicit mathematical models with no recording hardware involved, providing clean train/test separation and parameterizable channel conditions.

TorchSig generates signals from explicit mathematical models with full parametric control over SNR, sample rate, bandwidth, center frequency, and channel impairments. This allows the experimental conditions to be precisely defined and reproduced, and critically, training and test sets can be drawn from fully independent generations with no shared source data. For this work, these controls allowed us to isolate specific impairment conditions, sweep SNR levels cleanly, and ensure that observed accuracy differences reflect genuine classifier behavior rather than dataset construction artifacts.

It must be acknowledged honestly that TorchSig is a much newer tool than RadioML and has not yet seen wide adoption in the published AMC literature. Using it does mean that the absolute accuracy numbers reported here are not directly comparable to most published deep-learning AMC benchmarks. We accept this tradeoff because our primary goal is to validate the photonic architecture itself under controlled and reproducible conditions, rather than to benchmark against the broader AMC literature.

4.3 Signal Parameters

Each frame consists of $N = 1024$ complex IQ samples. The simulated sample rate is 10 MHz and signals are generated with a bandwidth of 4 MHz, yielding approximately 2.5 samples per symbol across all modulation classes. The center frequency is fixed at 0 Hz (DC), meaning no carrier frequency offset is applied, so all frames share the same constellation orientation. This is a deliberate simplification for clean

[1]These issues have been documented in technical analyses of the dataset. DeepSig, the creator of RadioML, now labels its own hosted versions of the data as "erratic" on their official dataset page [11]. A detailed signal-level audit is available at the Cyclostationary Signal Processing Blog [12].

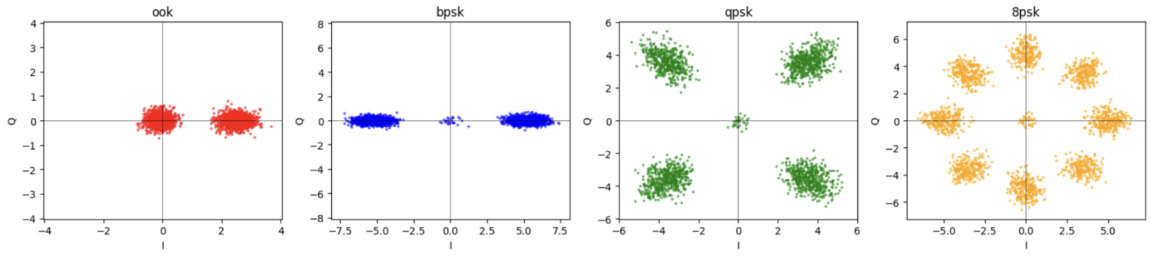

Figure 4.1: Representative clean-signal constellations for each of the four modulation classes at SNR = 30 dB with all impairments disabled. The vertical (Q) spread visible for OOK and BPSK is purely the AWGN contribution to the imaginary axis: even though their ideal constellation points are real-valued, the noise process is complex.

validation; carrier frequency effects are partially reintroduced at higher impairment levels through the applied channel transforms. Signal duration is set to exactly 1024 samples per frame, SNR is fixed at the target evaluation level, and all remaining TorchSig parameters are held at their defaults unless otherwise noted.

For each modulation class we generate 1000 training frames and 200 test frames, drawn from fully independent generations. Across the four classes this gives 4000 training frames and 800 test frames per experiment.

4.4 Impairment Levels

TorchSig parameterizes channel realism through an integer impairment level. Reading the source of `torchsig.transforms.impairments.Impairments`, the three documented levels apply progressively richer transform sets to the transmitted signal and the received frame.

Level 0 (clean). No transmit-side hardware impairments are applied to the signal. The received frame is processed by a `RandAugment` augmentation layer that randomly selects two of: `RandomDropSamples`, `ChannelSwap`, `TimeReversal`, and `AddSlope`. The signal is degraded only by the additive white Gaussian noise corresponding to the chosen SNR level.

Level 1 (cabled). The transmit-side transforms now include a probabilistic chain of hardware impairments: amplitude quantization, clock drift and jitter, passband ripple, IQ imbalance, carrier phase noise, carrier frequency drift, carrier phase offset (always applied), intermodulation products, nonlinear amplifier compression, spurious tones, and spectral inversion. The receive-side transforms include the same augmentation layer plus an independent second pass of hardware impairments and a digital automatic gain control. This level models a high-quality cabled connection between transmitter and receiver, in which there is no propagation channel but both ends contribute their own electronic imperfections.

Level 2 (wireless). Identical to Level 1 with one critical addition: the transmit-side chain now includes `Fading`, applied with probability 0.25, which simulates multipath fading characteristic of a wireless propagation channel. The fading process introduces frequency-selective amplitude attenuation and phase distortion that destroys much of the constellation structure relied on by phase-based classifiers. This is the most realistic and the most challenging condition we evaluate.

Figures 4.2 and 4.3 show IQ constellations for all four modulation classes across the three SNR levels and three impairment conditions respectively. At 30 dB SNR under clean conditions the constellation structure of each class is clearly visible; at 10 dB the clusters smear substantially due to AWGN. Under cabled impairments the constellations are distorted by the probabilistic hardware transform chain, which applies each impairment independently at a fixed probability per sample. Under wireless conditions, multipath fading is additionally applied with probability 0.25, meaning its effect is present in a subset of frames rather than uniformly. Because transforms are selected stochastically at generation time, the specific distortion pattern will differ between independent dataset generations, which is reflected in the variability visible across the constellation plots.

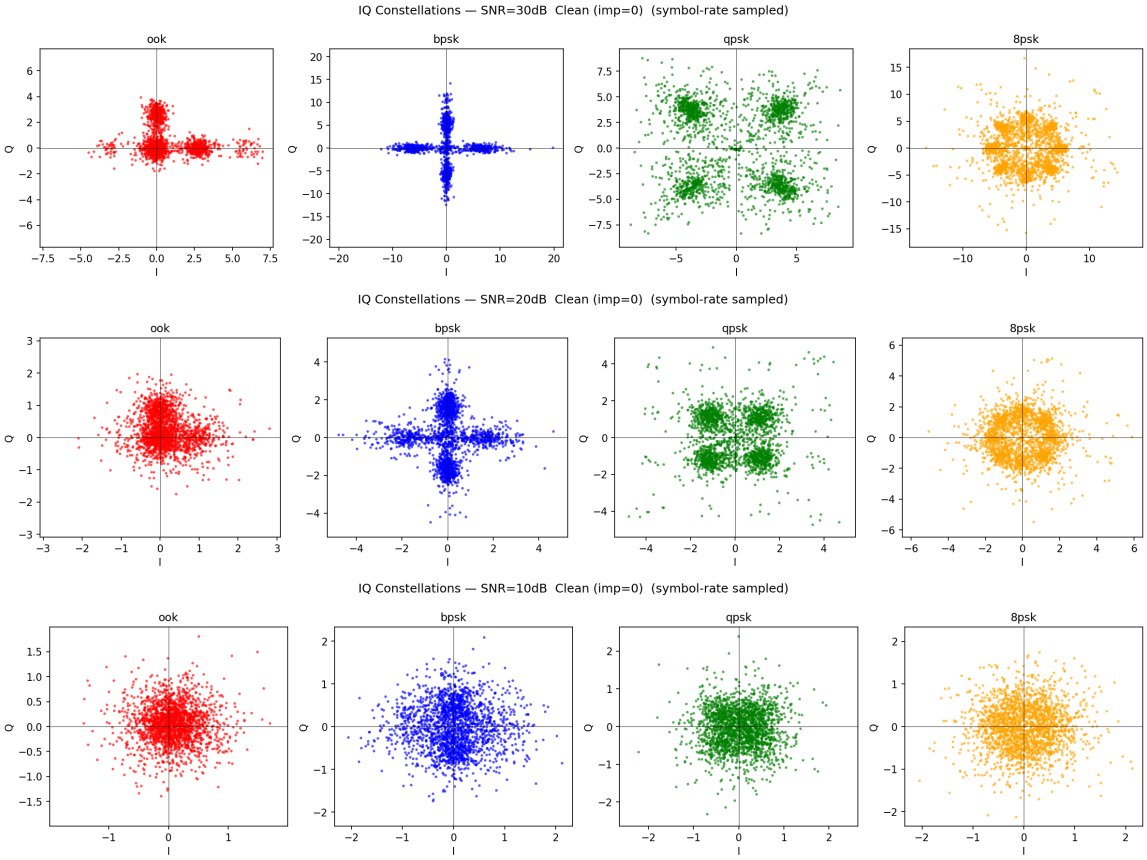

Figure 4.2: IQ constellations for all four modulation classes at SNR = 30, 20, and 10 dB (top to bottom) under clean (imp= 0) conditions. Constellation structure becomes progressively less clear and fuzzy as SNR decreases.

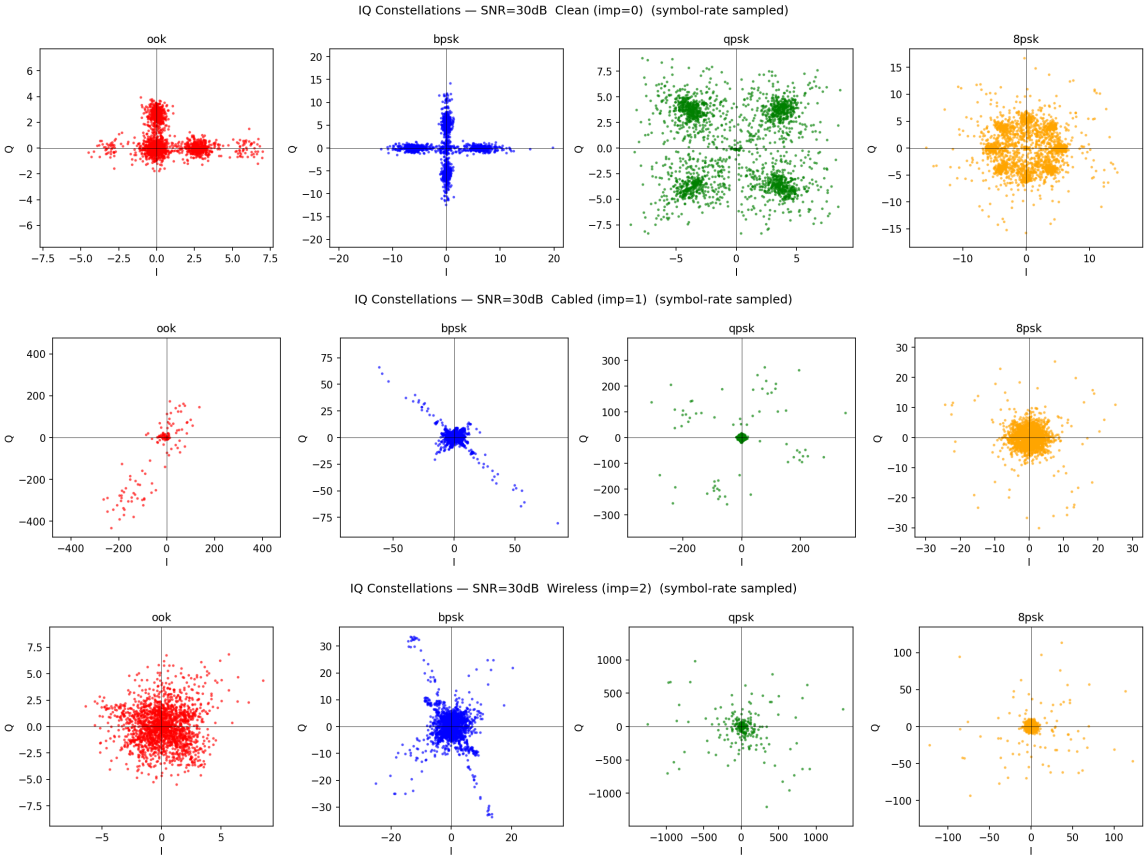

Figure 4.3: IQ constellations for all four modulation classes at SNR = 30 dB across the three impairment levels: clean (imp= 0), cabled hardware impairments (imp= 1), and wireless fading (imp= 2) (top to bottom). The distortion patterns under cabled and wireless conditions will vary between dataset generations due to the probabilistic selection of transforms applied at each level.

Chapter 5

Results

5.1 Architecture Sweep

Table 5.1 summarizes the classification accuracy of five photonic NN topologies on the four-class problem at 30 dB SNR under clean conditions. The table sweeps from a purely linear baseline through progressively larger nonlinear topologies.

Topology	**MRRs**	**Nonlinear?**	**Accuracy**
2-2 (linear)	4	No	60.2%
2-2-2	10	Yes	83.2%
2-3-2	**15**	**Yes**	**97.0**%
2-8-2	40	Yes	$\sim$100%
2-8-4	56	Yes	$\sim$100%

Table 5.1: Test accuracy across photonic NN topologies at 30 dB SNR under clean conditions. The 2-2 baseline contains no hidden layer and no nonlinearity. The 2-3-2 topology, with only 15 MRRs, is the smallest topology to clear 95%.

The central result of the architecture sweep is that remarkably few MRRs are needed to achieve high classification accuracy. A 2-3-2 topology with only 15 MRRs reaches 97% accuracy, indicating that the photonic NN is able to project the four modulation classes into a well-separated 2D feature space even at this minimal scale. The jump from 2-2-2 (10 MRRs, 83.2%) to 2-3-2 (15 MRRs, 97.0%) with the addition

of just one hidden unit suggests that the photonic frontend is operating close to a sufficient representational capacity for this four-class problem, and that the classifier has little work left to do once the features are formed. The linear baseline (60.2%) confirms that the nonlinearity is a necessary structural component, as expected, but the more practically significant finding is how quickly accuracy saturates once nonlinearity is present when using clean signals.

5.2 Decision Boundary and Confusion Matrix

Figure 5.1 visualizes the decision boundaries learned by the classifier in the 2-3-2 system's 2D feature space at 30 dB clean. The four classes occupy clearly separated regions of the feature plane, with OOK and BPSK located in distinct positions and QPSK and 8PSK forming neighboring clusters that the classifier separates along an approximately linear boundary.

The corresponding confusion matrix (Figure 5.2) makes the residual error structure explicit: misclassifications are confined to the QPSK $\leftrightarrow$ 8PSK boundary, with both OOK and BPSK classified perfectly. This is the expected pattern given the geometry of the modulations as QPSK and 8PSK occupy the same circle in IQ space and differ only in the density of constellation points, so they remain the hardest pair for the small network.

5.3 Robustness to SNR

Figure 5.3 shows accuracy at three SNR levels (30, 20, and 10 dB) for the 2-3-2, 2-8-2, and 2-8-4 topologies. The 2-3-2 system holds 97.0% at 30 dB, drops modestly to 94.1% at 20 dB, and degrades to 76.6% at 10 dB. The larger topologies are substantially more robust: 2-8-2 reaches near-perfect accuracy at 20 dB and 95.3% at 10 dB, with 2-8-4 achieving comparable performance at 92.4% at 10 dB.

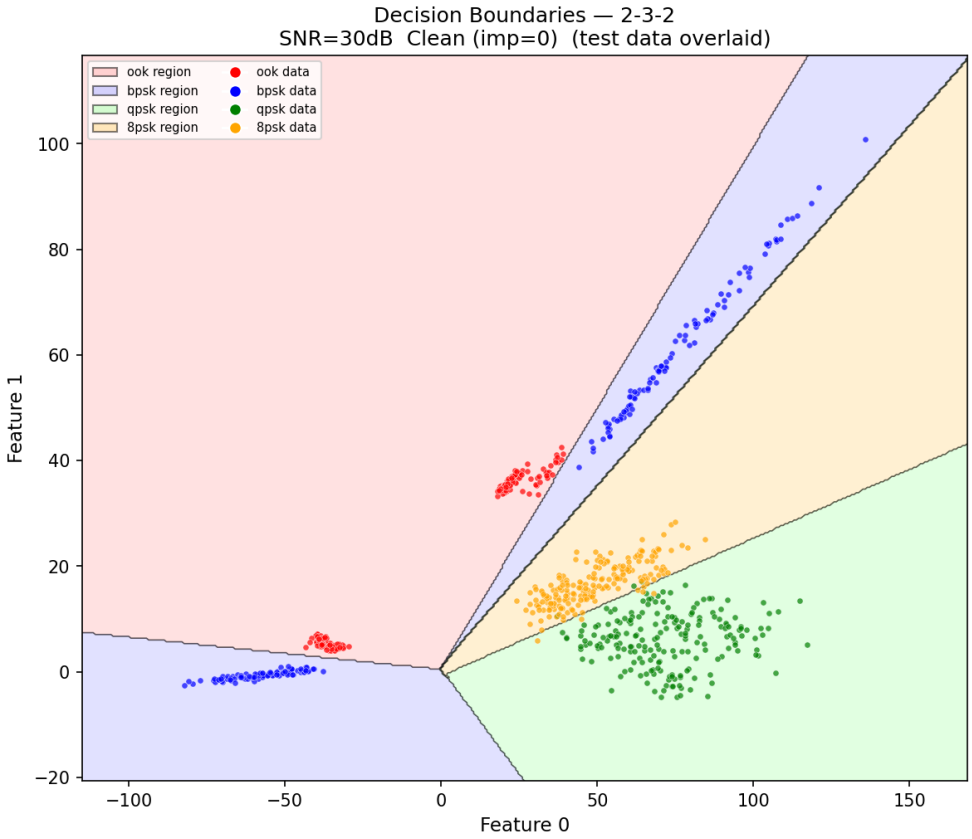

Figure 5.1: Decision boundaries learned by the joint photonic-NN / classifier system in the 2D feature space produced by the integrator, for the 2-3-2 topology at 30 dB SNR under clean conditions. Test points are overlaid on the colored class regions. OOK and BPSK occupy distinct clusters; QPSK and 8PSK are adjacent and account for the small residual error.

5.4 Robustness to Channel Impairments

Figure 5.4 shows accuracy at 30 dB SNR across the three impairment levels for all three topologies. All models degrade substantially from clean to impaired conditions: the 2-3-2 system falls from 97.0% to 58.2% under cabled impairments and further to 46.4% under wireless fading. Larger topologies show meaningfully better retention, with the 2-8-4 reaching 68% under wireless conditions compared to 46.4% for the smallest model, though all topologies remain well below their clean-condition performance.

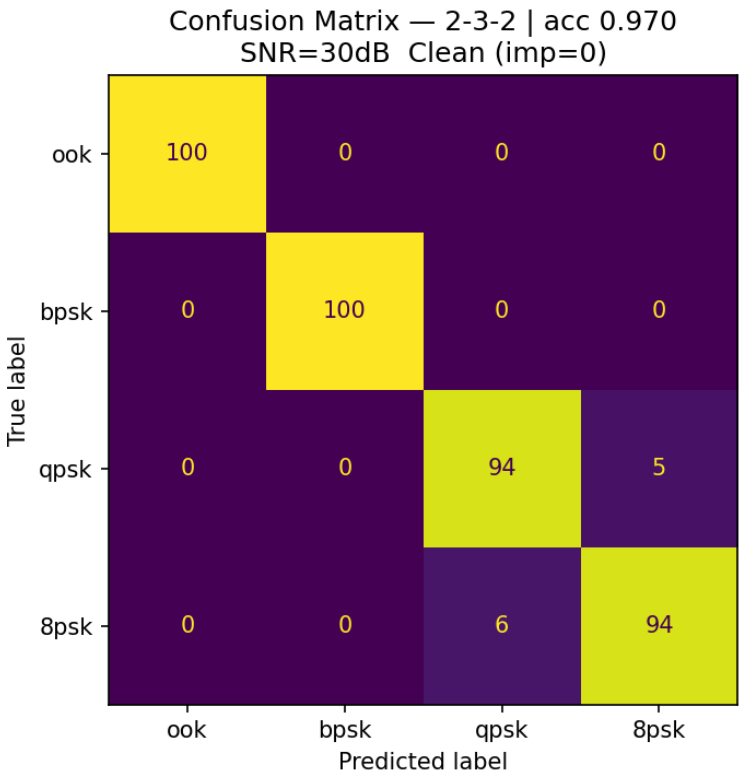

Figure 5.2: Confusion matrix for the 2-3-2 topology at 30 dB clean. OOK and BPSK are classified perfectly; the residual error consists entirely of QPSK $\leftrightarrow$ 8PSK confusions, consistent with their geometric similarity in IQ space.

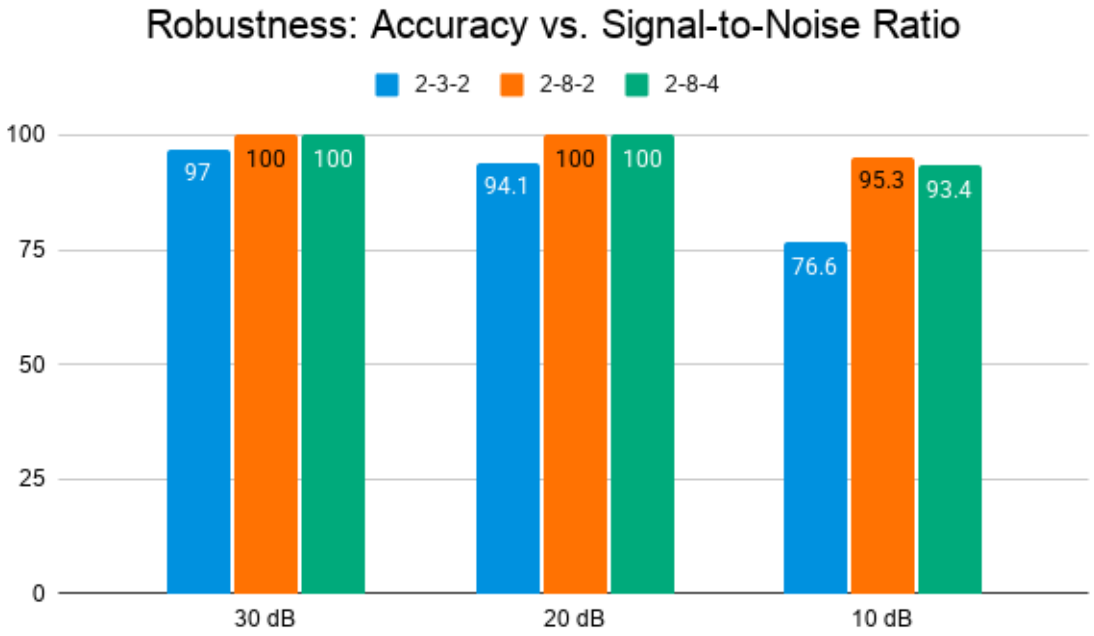

Figure 5.3: Test accuracy versus SNR for three photonic NN topologies under clean conditions. The 2-3-2 network is sufficient at high SNR but degrades visibly at 10 dB; scaling to 2-8-2 or 2-8-4 (40–56 MRRs) recovers near-perfect accuracy across the full SNR range.

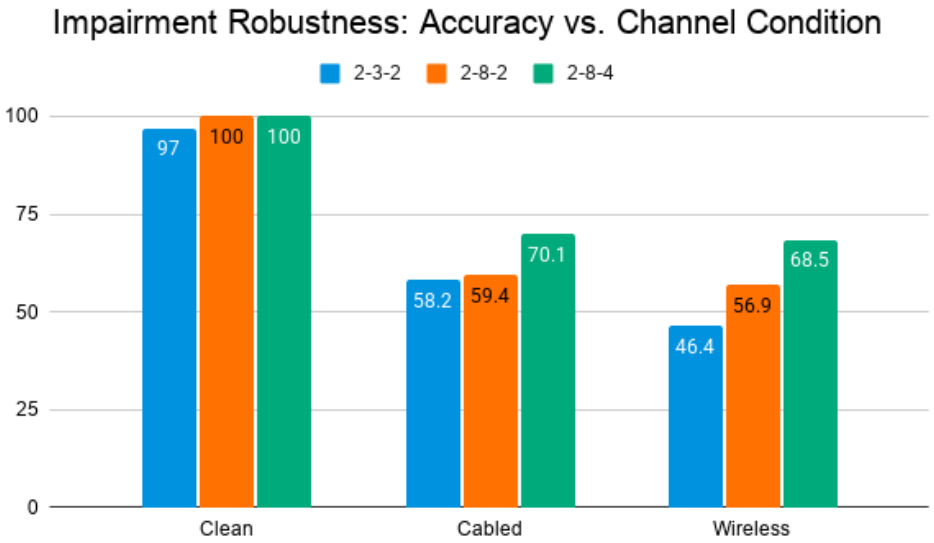

Figure 5.4: Test accuracy versus channel impairment level at 30 dB SNR for three photonic NN topologies. All topologies degrade sharply from clean to cabled to wireless conditions. The largest topology (2-8-4, 56 MRRs) retains the most accuracy under impairment but still falls substantially below its clean-condition performance.

Chapter 6

Discussion

6.1 15 MRRs is Realistic Today

The central result of the architecture sweep is that 97% classification accuracy is achievable with only 15 MRRs. This is significant because it places the photonic frontend requirement well within the capabilities of silicon photonic chips fabricated today. The Tait et al. silicon photonic chip [3], which forms the direct hardware predecessor to the architecture simulated here, demonstrated 16 tunable MRR weights on a single fabricated die. The 15-MRR topology validated in this work is therefore not a hypothetical future device but one that is directly realizable with existing fabrication processes, and the larger topologies evaluated (up to 56 MRRs) remain well within the scale of what current silicon photonic platforms support. Importantly, the trained weights produced by this simulation are directly intended for deployment on the fabricated MRR chip available at Princeton's Lightwave Communications Lab, where the full pipeline will be validated in hardware using time multiplexing to realize the weight bank on the available device.

6.2 Nonlinearity is Structural

The 60% accuracy of the linear baseline confirms what the moment-integral formulation predicts: a linear photonic frontend followed by integration computes only a linear functional of the signal, which is insufficient to extract the higher-order statistics needed to distinguish between modulation classes that share the same power and phase symmetry. Higher-order statistics, of the kind exploited by classical cumulant-based classifiers [2], require a nonlinear transformation before integration to be accessible. The nonlinearity is therefore not an implementation detail but a structural requirement of the architecture. It should be noted that the photonic NN does not explicitly compute cumulants; rather, it learns its own nonlinear projection of the signal through end-to-end training, one that is optimized for class separability in the integrated feature space rather than corresponding to any analytically defined statistic. The MRR transmission profile provides the required nonlinearity naturally, and the ReLU in the simulation models this physical response. Any fabricated photonic AMC chip implementing this approach must preserve this nonlinear stage.

6.3 The Impairment Gap

Impairment robustness is the principal open problem identified by this work. The degradation under cabled and wireless conditions reflects the genuine difficulty of classifying signals under hardware impairments and fading: wireless fading introduces frequency-selective amplitude attenuation and phase distortion that partially destroys the constellation structure the photonic NN learned to project into a separable feature space. Critically, each model was trained on data from its own impairment level, so this is not a training distribution mismatch issue but an inherent challenge of the classification problem under these conditions.

Several observations suggest the results reported here are not a ceiling. Even

modest increases in model size produce significant robustness gains, with the jump from 2-3-2 to 2-8-4 recovering roughly 20 percentage points under wireless fading, indicating the architecture is not fundamentally limited in its robustness capacity. Additionally, each model was trained on only 1000 frames per class for 300 epochs, a deliberately small amount chosen for initial validation. Targeted training with larger datasets, longer schedules, and mixed impairment conditions would be expected to yield meaningful further improvement. This work establishes the baseline behavior of the hardware-constrained photonic frontend under each condition; impairment-aware training at scale is the natural and immediate next step.

6.4 Limitations of the Software Simulation

Several physical phenomena that affect a real fabricated photonic chip are not captured by the simulation reported here.

Fabrication variation. Real MRRs vary in radius and effective index across a die, shifting their resonance wavelengths unpredictably. Compensating for this requires per-ring calibration and feedback control. Our simulation assumes ideal, identical rings.

Thermal drift and crosstalk. The same heaters that program the weights also dissipate power, and the temperature of one ring can drift the resonance of its neighbors. Real systems require active thermal stabilization. Our simulation has no thermal coupling between weights.

Photodetector noise. The optical-to-electrical conversion at the integrating photodetector contributes shot noise and electronic noise that further degrade the integrator output. We model the photodetector as an ideal sum.

Optical losses and finite extinction. MRRs have finite extinction ratios, so a weight intended to be zero leaks a small amount of signal. Our quantized weight model treats every level as exactly representable; in hardware the achievable dynamic range is narrower.

These limitations mean the accuracy numbers reported here should be treated as optimistic estimates of what a fabricated chip would achieve at first measurement, and closing the gap between simulation and hardware is the primary motivation for the physical validation work described in the following chapter.

Chapter 7

Future Work

Physical chip validation. The most direct extension of this work is to replace the simulated photonic NN with a real fabricated MRR weight bank and validate the end-to-end pipeline against measured chip data. The Lightwave Communications Lab has the necessary fabricated devices and measurement infrastructure. The trained weights produced by this simulation are quantized to 12-bit DAC precision and are compatible with the tuning range of the available MRR chip, and will be loaded onto hardware using time multiplexing to realize the full weight bank on the fabricated device.

Analog integrator. The integrator in the present simulation is a digital accumulator. The natural hardware realization is an integrating photodetector that performs the accumulation directly in the analog domain, completing a fully analog frontend in which the first digital operation is the classifier itself.

FPGA classifier deployment. A path to FPGA deployment via hls4ml [10] on a Xilinx ZCU216 RFSoC has been prototyped in Vivado as part of this project, with the classifier converted to HLS C++ and packaged as an IP block. The process for loading and configuring the ZCU216 for this pipeline has been documented in detail

by Lederman [5], and this work will follow that established procedure. Full hardware-in-the-loop validation, in which the FPGA receives live integrator output and emits class decisions in real time, is the next step toward a complete deployable system.

Impairment-robust training. Each model in this work was trained on a single impairment level, with a relatively small dataset of 1000 frames per class. Training jointly across all impairment levels simultaneously, increasing the number of training frames, and exploring curriculum schedules that progressively introduce harder impairment conditions are all candidate approaches for improving robustness in a focused follow-on study.

Chapter 8

Conclusion

Automatic Modulation Classification places real demands on receiver hardware: as signal bandwidths grow, the cost of digitizing and processing every sample with conventional electronics becomes a practical constraint in power-limited and latency-sensitive deployments. Silicon photonic neural networks built from micro-ring resonator weight banks offer a fundamentally different substrate for this problem, performing analog computation directly on the optical signal at near-zero marginal energy per operation. The moment-integral formulation shows that AMC is not merely a problem to which photonic processing can be applied, but one whose mathematical structure maps naturally onto a three-stage photonic pipeline.

This work has validated that pipeline in simulation under realistic hardware constraints. A 2-3-2 photonic neural network of only 15 MRRs achieves 97% classification accuracy across four modulation classes at 30 dB SNR, a topology realizable on silicon photonic chips available today. The nonlinearity provided by the MRR transmission profile is confirmed to be structurally essential, with a linear baseline collapsing to 60%. Robustness to channel impairments remains the principal open problem, with accuracy degrading substantially under cabled and wireless conditions, though larger topologies show meaningful improvement and targeted training at scale is a clear path

forward.

The trained weights produced by this simulation are quantized to 12-bit DAC precision and are ready for direct transfer to the fabricated MRR chip at Princeton's Lightwave Communications Lab. In parallel, a deployment path for the digital classifier onto a Xilinx ZCU216 RFSoC via hls4ml has been prototyped in Vivado, bringing the full hardware-in-the-loop system within reach. The minimum viable photonic AMC design has been characterized in simulation; physical validation is the immediate next step.

Bibliography

[1] T. J. O'Shea, T. Roy, and T. C. Clancy, "Over-the-Air Deep Learning Based Radio Signal Classification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 168–179, Feb. 2018.

[2] A. Swami and B. M. Sadler, "Hierarchical Digital Modulation Classification Using Cumulants," *IEEE Transactions on Communications*, vol. 48, no. 3, pp. 416–429, 2000.

[3] A. N. Tait, T. F. de Lima, E. Zhou, A. X. Wu, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Neuromorphic Photonic Networks Using Silicon Photonic Weight Banks," *Scientific Reports*, vol. 7, no. 1, article 7430, 2017.

[4] P. R. Prucnal and B. J. Shastri, *Neuromorphic Photonics*. Boca Raton, FL: CRC Press, 2017.

[5] J. C. Lederman, *Silicon Photonic Systems Applications*, Ph.D. dissertation, Princeton University, 2025.

[6] S. Xu *et al.*, "Analog spatiotemporal feature extraction for cognitive radio-frequency sensing with integrated photonics," *Light: Science & Applications*, vol. 13, article 50, 2024. DOI: 10.1038/s41377-024-01390-9.

[7] TorchDSP, "TorchSig: Open Synthetic RF Signal Dataset and DSP Library," `https://github.com/TorchDSP/torchsig`.

[8] Y. Bengio, N. Léonard, and A. Courville, "Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation," *arXiv preprint arXiv:1308.3432*, 2013.

[9] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.

[10] J. Duarte *et al.*, "Fast Inference of Deep Neural Networks in FPGAs for Particle Physics," *Journal of Instrumentation*, vol. 13, no. 7, p. P07027, 2018.

[11] DeepSig Inc., "Datasets," `https://www.deepsig.ai/datasets/`, 2024.

[12] Cyclostationary Signal Processing Blog, "More on DeepSig's RML Datasets," `https://cyclostationary.blog/2020/08/17/more-on-deepsigs-rml-data-sets/`, 2020.