sensors

MDPI

Article

NeuroDetect: Deep Learning-Based Signal Detection in Phase-Modulated Systems with Low-Resolution Quantization

Chanula Luckshan [1], Samiru Gayan [1], Hazer Inaltekin [2,*], Ruhui Zhang [3] and David Akman [4]

[1] Department of Electronic and Telecommunication Engineering, University of Moratuwa, Moratuwa 10400, Sri Lanka; luckshangwcm.20@uom.lk (C.L.); samirug@uom.lk (S.G.)
[2] School of Engineering, Macquarie University, North Ryde, NSW 2109, Australia
[3] Institute of Advanced Study in Mathematics, Harbin Institute of Technology, Harbin 150001, China; ruhui_zhang@hit.edu.cn
[4] Lifelong Learning, University of New South Wales, Kensington, NSW 2052, Australia; d.akman@unsw.edu.au
* Correspondence: hazer.inaltekin@mq.edu.au

Abstract: This manuscript introduces NeuroDetect, a model-free deep learning-based signal detection framework tailored for phase-modulated wireless systems with low-resolution analog-to-digital converters (ADCs). The proposed framework eliminates the need for explicit channel state information, which is typically difficult to acquire under coarse quantization. NeuroDetect utilizes a neural network architecture to learn the nonlinear relationship between quantized received signals and transmitted symbols directly from data. It achieves near-optimum performance, within a worst-case 12% margin of the maximum likelihood detector that assumes perfect channel knowledge. We rigorously investigate the interplay between ADC resolution and detection accuracy, introducing novel penalty metrics that quantify the effects of both quantization and learning errors. Our results shed light on the design trade-offs between ADC resolution and detection accuracy, providing future directions for developing energy-efficient high-speed and wideband wireless systems.

Keywords: low-resolution ADCs; machine learning; deep learning; model-free; signal detection; maximum likelihood detection; energy efficiency

check for updates

Academic Editor: Tiejun Lv

Received: 8 April 2025
Revised: 15 May 2025
Accepted: 15 May 2025
Published: 19 May 2025

Citation: Luckshan, C.; Gayan, S.; Inaltekin, H.; Zhang, R.; Akman, D. NeuroDetect: Deep Learning-Based Signal Detection in Phase-Modulated Systems with Low-Resolution Quantization. *Sensors* **2025**, *25*, 3192. https://doi.org/10.3390/s25103192

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Background and Motivation

Next-generation wireless systems, including ultra-massive multiple-input multiple-output (MIMO) arrays, millimeter-wave (mmWave) networks, and emerging sub-THz technologies, are pushing the limits of data rate, bandwidth, reliability, and latency requirements [1–3]. As these systems grow in scale and complexity, receiver power consumption emerges as a critical design bottleneck. Among the receiver components, ADCs contribute a significant fraction of the total power budget, particularly in high-speed and wideband wireless systems [4,5].

The power consumed by ADCs scales exponentially with resolution and linearly with sampling rate [6]. As a result, employing high-resolution ADCs in systems with high sampling rates can lead to substantial energy overhead. For instance, a high-speed ADC operating at or above 20 GSamples/s with 8–12 bits of resolution typically consumes around 500 mW. Since each RF chain requires two ADCs, this results in approximately 1 W per RF chain. In a future mmWave massive MIMO system with 256 RF chains (512 ADCs in total), the aggregate power consumption can exceed 256 W, posing a major challenge for

practical deployment [7]. To address this issue, low-resolution ADCs have emerged as a promising solution, significantly reducing power consumption by limiting the number of quantization bits [8,9].

While low-resolution ADCs offer substantial energy savings, they also introduce significant quantization-induced nonlinearities that fundamentally alter the input–output relationship of the communication system [10]. These nonlinear distortions undermine the assumptions underpinning traditional coherent detection methods, which typically depend on accurate channel state information (CSI). Estimating CSI becomes considerably more difficult under coarse quantization, due to the limited information content in the quantized observations [11]. Consequently, any mismatch or inaccuracy in channel estimation can lead to severe performance degradation, particularly in systems employing higher-order modulation schemes or operating in fast-fading environments.

A range of prior studies, such as [12], have explored *model-based* detection strategies, including maximum likelihood (ML) detection and sphere decoding, to mitigate the effects of quantization noise. While these methods are theoretically near-optimum, they often rely on the assumption of perfect or near-perfect CSI and become computationally infeasible as system dimensions and modulation orders increase. To address the trade-off between detection performance and practical complexity, deep learning (DL)-based detection techniques have emerged as compelling alternatives, offering great potential as data-driven, low-complexity solutions without explicit reliance on channel models.

The DL-based techniques enable the development of parametric models that learn system dynamics and signal detection rules directly from data. This class of methods, known as *model-driven* approaches, incorporates a priori domain knowledge into the network architecture while learning model parameters from training data [13]. By leveraging model structure, these approaches can approximate ML detection, achieving reduced computational complexity and improved numerical stability, especially in high signal-to-noise ratio (SNR) regimes [14].

Despite their advantages, model-driven approaches often depend on partial or idealized model knowledge, limiting their ability to generalize when real-world conditions diverge from training assumptions. In practical wireless environments, where channel characteristics may be highly dynamic, nonlinear, or poorly understood, these methods may fail to deliver robust performance. This limitation motivates the need for a *model-free* approach to signal detection, in which the underlying channel input–output relationship is learned directly from data, without relying on prior model assumptions.

To address these challenges, we propose NeuroDetect, a model-free DL-based signal detection framework designed for receivers equipped with low-resolution ADCs. NeuroDetect operates entirely without explicit channel knowledge and learns to map quantized observations to transmitted symbols directly from data. Our approach capitalizes on the representational power of deep neural networks to capture the underlying nonlinear relationship introduced by low-resolution quantization. By training on a small number of labeled examples within each channel coherence interval, NeuroDetect effectively eliminates the need for channel estimation, thereby offering robust performance across a range of fading environments.

1.2. Main Contributions

The key contributions of this work are summarized as follows:

- **Model-Free Detection:** We introduce a deep neural network architecture that learns to detect symbols under phase quantization without requiring explicit CSI. The proposed solution complements existing model-based signal detection approaches by offering

a robust alternative for scenarios where CSI is unavailable at the receiver with low-resolution quantization.

- **Near-Optimum Performance:** We demonstrate that NeuroDetect achieves symbol error rates within 12% of the ML detector, which assumes perfect CSI. This is the worst-case performance gap, and we show that the gap is much smaller in almost all scenarios we studied. This result is significant as it shows that near-optimum detection is attainable without the overhead of channel estimation. A key merit of NeuroDetect is to show that a lightweight and fully-connected architecture, if trained model-free, can close the gap to optimal ML detection under severe quantization errors, without the need for custom layers or specialized activation functions.
- **Asymptotic Optimality:** We show that the proposed DL-based detector achieves the optimum asymptotic error decay rate for different quantization levels, revealing a characteristic ternary behavior in diversity order.
- **Penalty Metrics:** We develop new metrics to quantify the learning and quantization penalties, offering insights into how the number of bits affects detection accuracy and power requirements.

1.3. Related Work

This section provides an overview of prior signal detection strategies for low-resolution ADC-based receivers. We first discuss a range of classical approaches and their inherent limitations. We then review DL-based solutions.

1.3.1. Classical Approaches

Over the past few years, numerous classical schemes have been developed for signal detection in low-resolution ADC-based wireless systems. Early work in [15] introduced a linear minimum-mean-square-error (LMMSE) receiver, which outperforms the conventional Wiener filter. However, this approach has the drawback of having error floors in the high SNR regime.

To overcome this limitation, subsequent research proposed the Bussgang LMMSE (BLMMSE) estimator, extending classical linear data detectors such as maximum ratio combining (MRC) and zero-forcing (ZF) via Bussgang decomposition [16]. This method was later adapted to the multi-bit scenario in [17], demonstrating improved performance over LMMSE. Nonetheless, error floors persist in the high SNR regime.

Assuming perfect CSI at the receiver, ML data detection has been investigated for one-bit quantized MIMO systems in [18,19] and extended to one-bit MIMO-OFDM systems in [20]. However, the computational complexity of the method in [18] increases exponentially with the signal constellation size, number of transmit antennas, and overall network dimensions, making it impractical for real-world implementations. To mitigate this issue, ref. [21] introduced a near-optimal ML detector based on convex optimization techniques. Although this method outperforms ZF, MRC, and LMMSE detectors in terms of SEP, simulation results indicate that it still suffers from an error floor at high SNR.

In [22], a sphere decoding-based near-ML detection method, termed one-bit sphere decoding (OSD), was presented for massive MIMO systems with one-bit ADCs. OSD reduces complexity by decomposing the received signal vector into multiple sub-vectors, thus balancing performance and computational demands. However, the complexity still grows exponentially with the number of transmit antennas and the constellation size.

Subsequent work in [23] addressed this scalability challenge by developing signal detection algorithms for one-bit MIMO systems under both perfect and statistical CSI, achieving near-optimum performance compared to the oracle lower bound in [24]. This method uses the Laplace approach to estimate the channel model with a Gaussian distribution.

However, it does not work well for all channel models, and the accuracy drops in real-world situations with more complex models.

Several studies, including our previous work, have investigated low-resolution phase quantization for *M*-PSK modulation schemes [11,12,25–27]. These works propose various detection techniques and demonstrate their effectiveness through Monte Carlo simulations. Notably, the simulations in [26] reveal that, under QPSK modulation with 8-bit phase quantization, a non-coherent block-fading channel can retain approximately 80–85% of the mutual information attainable with fully unquantized observations, incurring only a modest 2–3 dB increase in SNR. Further, ref. [27] shows that PSK achieves capacity under various channel models when phase quantization is employed at the receiver output.

1.3.2. DL-Based Approaches

DL-based methods offer a compelling alternative for the classical approaches, delivering superior performance in practical scenarios while reducing computational complexity [28]. These methods have already demonstrated remarkable capabilities in channel estimation, equalization, and signal detection, making them a viable solution for future wireless systems [29–32].

A supervised learning approach for blind signal detection was proposed in [33] to learn the empirical probability density function for each symbol with low-resolution quantization. Here, three blind detection methods were introduced: Empirical Maximum Likelihood Detection (e-MLD), Minimum Mean-Distance (MMD), and Minimum Center Detection (MCD). These methods are related to *K*-nearest neighbors classification and nearest centroid classification. It has been shown that MCD performs better than MMD and e-MLD. However, like classical schemes, the practical implementation is limited due to high complexity.

The same authors proposed a more efficient supervised learning-aided successive interference cancellation method to reduce both training overhead and complexity [34]. However, the efficiency of this method depends on the training sequence, which requires a long sequence of data to effectively learn the function approximator. To address this performance degradation, a supervised and semi-supervised framework has been introduced in [35], which uses a cyclic redundancy check for error control.

Several other model-driven DL approaches were proposed for signal detection and channel estimation with low-resolution ADCs. The Few-Bit Massive MIMO Channel Estimation Network (FBMCENet) and the Few-Bit Massive MIMO Data Detection Network (FBM-DetNet) [36] outperform the detector in [35] with QPSK and 16-QAM constellations with an SNR of 10 dB when perfect CSI is available. However, their complexity remains a major drawback. To address this, the same authors proposed the One-Bit Massive MIMO Data Detection Network (OBMNet), which integrates DL with conventional signal processing techniques [37]. By using a sigmoid approximation of the Gaussian cumulative distribution function, OBMNet achieves lower complexity compared to other proposed methods.

These model-driven approaches require the availability of CSI as well as prior knowledge of the system parameters. Hence, using model-driven approaches in a dynamically changing wireless communication system poses a challenge. By bridging the gap between model-free and model-driven approaches, the authors in [38] proposed LoRD-Net (Low Resolution Detection Network), a hybrid system that learns to carry out blind symbol detection from one-bit measurements.

Beyond signal detection, DL was also applied to pilot design and channel estimation in mixed-ADC massive MIMO systems [39,40]. These papers reduce pilot overhead and further highlight the potential of DL-based techniques in wireless communications with

low-resolution ADCs. Bidirectional LSTM (BiLSTM) models demonstrated significant improvements in channel estimation for MIMO-OFDM low-resolution ADC systems by leveraging bi-directional training [41,42].

2. System Setup

2.1. Channel Model and Signal Modulation

We focus on the point-to-point wireless channel with flat-fading. The received discrete-time baseband equivalent signal Y is given by

$$Y = \sqrt{\mathsf{SNR}} \cdot HX + W, \tag{1}$$

where $X \in \mathcal{C} \subset \mathbb{C}$ is the transmitted signal, $\mathcal{C}$ is the constellation set of information signals in $\mathbb{C}$ (complex plane), SNR is the ratio of the transmitted signal energy to the additive white Gaussian noise (AWGN) spectral density, $H \in \mathbb{C}$ is the unit power channel gain between the transmitter and the receiver, and W is the circularly-symmetric zero-mean unit-variance AWGN. In (1), the prefactor $\sqrt{\mathsf{SNR}}$ follows a standard normalization convention that keeps X unit-energy and W unit-variance for analytical tractability and simplicity of the exposition. We will assume that $\mathcal{C} = \left\{ e^{j\pi\left(\frac{2\gamma+1}{M}-1\right)} \right\}_{\gamma=0}^{M-1}$ throughout the paper. This represents the conventional M-ary phase shift keying (M-PSK) constellation. For simplicity, we restrict our analysis to cases where M is an integer power of 2. This setup is also adopted in our previous work [12], and we use those analytical results as a benchmark for our study.

2.2. Receiver Architecture

As illustrated in Figure 1, the receiver architecture is based on a low-resolution ADC.

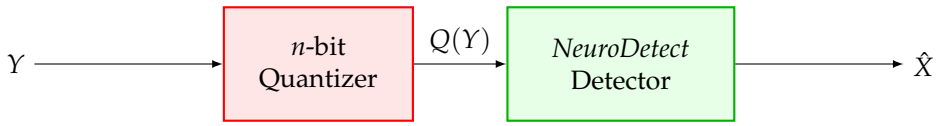


Figure 1. General receiver architecture with the low-resolution quantization. Here, NeuroDetect is the proposed DL-based framework for signal detection.

In this setup, the received signal Y is first sent through a low-resolution quantizer, and the resulting quantized received signal $Q(Y)$ is then used to estimate the transmitted signal X.

In this work, we use a phase quantizer. The adoption of phase quantization in our receiver architecture is driven by two key factors. First, since channel impairments typically manifest as phase rotations in transmitted signals, the quantization and decision regions for M-PSK modulation can be naturally represented as convex cones in the complex plane [12], eliminating the need for automatic gain control. Second, phase quantizers can be efficiently implemented using one-bit ADCs composed of simple comparators, which operate with minimal power consumption, typically in the milliwatt range.

If n bits are used to quantize Y, the quantizer $Q(\cdot)$ divides the complex domain $\mathbb{C}$ into 2^n quantization regions and outputs the index of the region in which Y lies as an input to the detector. Formally, we declare $Q(Y) = q$ if $Y \in \mathcal{R}_q$ for $q \in [0 : 2^n - 1]$, where $\mathcal{R}_q \subseteq \mathbb{C}$ is the qth quantization region. Since information is encoded in the phase of X with the above choice of constellation points, $\mathcal{R}_q$ is chosen to be the convex cone given by

$$\mathcal{R}_q = \left\{ z \in \mathbb{C} : \frac{2\pi}{2^n} q \leq \mathsf{Arg}(z) + \pi \leq \frac{2\pi}{2^n}(q+1) \right\}. \tag{2}$$

The receiver architecture outlined so far is similar to the one presented in [12]. The key distinction of our work lies in the implementation of a DL signal detector. Our DL architecture utilizes the quantized signal information $Q(Y)$ to predict the transmitted signal X, without having any channel knowledge. In the following section, we detail the design and functionality of the proposed DL-based signal detector.

3. Deep Learning-Based Signal Detection

This section presents our proposed DL-based signal detector designed to achieve near-optimum performance in low-resolution quantization scenarios, without relying on explicit channel knowledge. The ML detector proposed in [12] relies on channel information to generate ML estimates of the transmitted data symbols. On the other hand, our DL-based detector is a fully data-driven solution, utilizing a deep neural network to extract meaningful features from quantized signal inputs. We provide the details for the model-free DL framework for signal detection with low-resolution quantization, our DL-based signal detector, its neural network architecture, and the corresponding training methodology in the following subsections.

3.1. Deep Learning for Signal Detection with Low-Resolution Quantization

We adopt the model-free signal detection framework illustrated in Figure 2, originally proposed in [13], for symbol detection in low-resolution wireless communication systems. The input–output relationship is highly nonlinear with low-resolution quantization, which is a challenge for signal detection using traditional detection methods. Accurate channel estimation is also another important challenge, due to the low number of quantization bits used at the receiver to estimate the channel. The model-free signal detection framework overcomes these challenges by learning a near-optimum signal detector directly from training data, and, hence, eliminating the need for channel estimation while capturing input–output non-linearity.

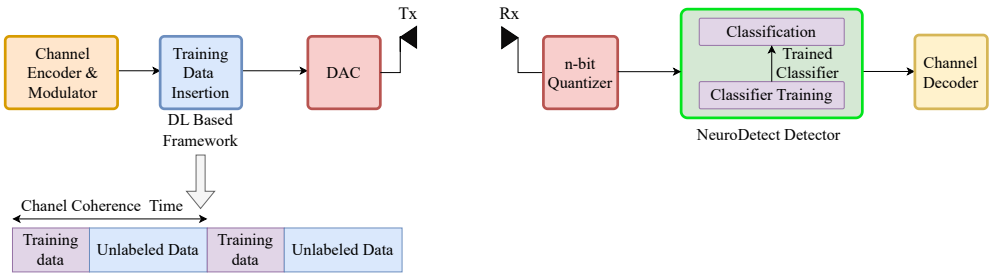


Figure 2. Model-free deep learning framework for signal detection with low-resolution quantization. Here, DAC is the digital-to-analog converter.

Within this framework, we divide the transmitted symbols during a channel coherence interval into two categories: *training data* and *unlabeled data*. The training data, which are known to both the transmitter and receiver, serve as labeled data for training a classifier to function as a signal detector at the receiver. We call the phase during which we train the classifier at the receiver the *Training Phase* (TP). We use the quantized received signals $Q(Y)$ as input features and the transmitted symbols X as target values during the TP to train the classifier.

The unlabeled data consist of transmitted symbols that are unknown to the receiver. To detect these symbols, the receiver uses the trained classifier with optimized parameters. We refer to this stage as the Data Transmission Phase (DP). During the DP, our signal detector processes the quantized received signal $Q(Y)$, assigns a likelihood to each possible transmitted symbol based on the observation $Q(Y)$, and selects the most probable symbol

as the estimate $\hat{X}$ of the transmitted symbol X. Unlike traditional methods or model-driven AI approaches [38,43], this framework operates without explicit channel state information.

We employ a DL-based classifier for signal detection at the receiver side of this framework. Our DL-based classifier consists of an input layer that gets $Q(Y)$ as an input, four hidden layers, and a SoftMax output layer (outputting the likelihoods for the transmitted data symbols). The architecture and implementation details of our proposed DL signal detector are given in Section 3.3.

3.2. Deep Learning-Based Signal Detector

Figure 3 shows the architecture of our DL-based signal detector. The proposed DNN classifies the received quantized signal $Q(Y)$ into one of the M possible transmitted signals, given by

$$X_\gamma = e^{J\pi\left(\frac{2\gamma+1}{M}-1\right)} \quad \text{for} \quad \gamma = 0, 1, \ldots, M-1.$$

As the transmitted symbol depends only on γ, the proposed architecture is designed to predict γ by observing the quantized received signals.

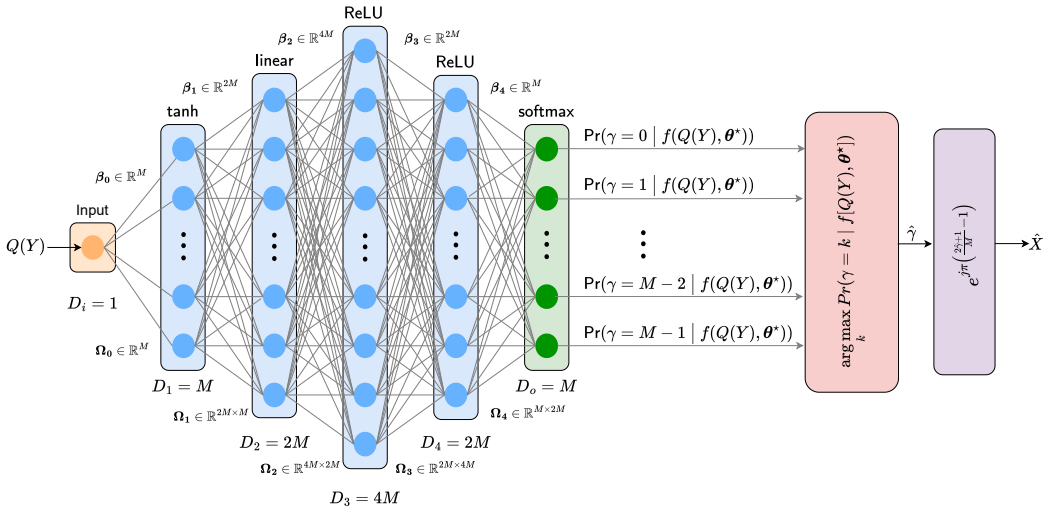


Figure 3. The proposed DL-based signal detector architecture for receivers with low-resolution quantization.

To build our DL-based detector, we define a model $f(Q(Y), \boldsymbol{\theta})$ with parameters $\boldsymbol{\theta}$, which processes an input $Q(Y)$ from a training dataset of I feature/label pairs $\left\{\left(Q(Y)^{[i]}, \gamma^{[i]}\right)\right\}_{i=1}^{I}$. This model produces a conditional probability distribution $\Pr(\gamma \mid \boldsymbol{\lambda})$, where $\boldsymbol{\lambda}$ are distribution parameters.

Because γ can take on any of the M discrete values $\{0, 1, \ldots, M-1\}$, we adopt a categorical distribution governed by $\boldsymbol{\lambda} = [\lambda_0, \lambda_1, \ldots, \lambda_{M-1}]^\top$. Hence,

$$\Pr(\gamma = k \mid \boldsymbol{\lambda}) \;=\; \lambda_k, \quad k = 0, 1, \ldots, M-1.$$

Our model learns to predict $\boldsymbol{\lambda}$ given $Q(Y)$ and, thus,

$$\Pr(\gamma \mid \boldsymbol{\lambda}) \;=\; \Pr\big(\gamma \mid f(Q(Y), \boldsymbol{\theta})\big).$$

The parameters $\lambda_1, \ldots, \lambda_{M-1}$ are restricted to lie between zero and one and must sum to one to form a valid probability distribution. Therefore, we pass the M outputs of the network through a softmax function that ensures these constraints are satisfied. Then, we have

$$\Pr(\gamma = k \mid f(Q(Y), \boldsymbol{\theta})) = \mathrm{softmax}_k(f(Q(Y), \boldsymbol{\theta})), \tag{3}$$

where $\mathsf{softmax}_k(\mathbf{z}) = \frac{e^{z_k}}{\sum_{m=0}^{M-1} e^{z_m}}$.

To train the model, we find the network parameters $\boldsymbol{\theta}$ that minimize the negative log-likelihood loss function $L(\boldsymbol{\theta})$ over the training dataset pairs $\left\{\left(Q(Y)^{[i]}, \gamma^{[i]}\right)\right\}_{i=1}^{I}$:

$$\begin{aligned} L(\boldsymbol{\theta}) &= -\sum_{i=1}^{I} \log\left(\mathsf{softmax}_{\gamma^{[i]}}\left(f\left(Q(Y)^{[i]}, \boldsymbol{\theta}\right)\right)\right), \\ &= -\sum_{i=1}^{I}\left(f_{\gamma^{[i]}}\left(Q(Y)^{[i]}, \boldsymbol{\theta}\right) - \log\left(\sum_{m=0}^{M-1} \exp\left(f_m\left(Q(Y)^{[i]}, \boldsymbol{\theta}\right)\right)\right)\right), \end{aligned} \tag{4}$$

where $f_m(Q(Y), \boldsymbol{\theta})$ denotes the output of the m-th branch of the softmax layer in Figure 3. This loss function is also known as the multiclass cross-entropy loss.

For a point estimate of γ, we take the most probable category

$$\hat{\gamma} = \arg\max_k \mathsf{Pr}(\gamma = k \mid f(Q(Y), \boldsymbol{\theta}^\star)), \tag{5}$$

where $\boldsymbol{\theta}^\star = \arg\min_{\boldsymbol{\theta}} L(\boldsymbol{\theta})$. Then, the estimated transmitted signal is given by

$$\hat{X} = \exp\left(\jmath\pi\left(\frac{2\hat{\gamma}+1}{M} - 1\right)\right). \tag{6}$$

3.3. NeuroDetect Architecture

The details of our DL-based signal detector, which we name *NeuroDetect*, are given below. The NeuroDetect architecture comprises a fully connected network with four hidden layers, as illustrated in Figure 3. The key components are as follows:

- **Input Layer:** Receives the quantized received signal $Q(Y)$, which is the output of the low-resolution ADC.
- **Hidden Layers:** Four fully connected layers, configured as follows:
  - Layer 1: Consists of neurons with a tanh activation function.
  - Layer 2: Consists of neurons with a linear activation function.
  - Layers 3 and 4: Consist of neurons with ReLU activation function.
- **Output Layer:** Consists of neurons with a softmax function that outputs

  $$\mathsf{Pr}(\gamma = k \mid f(Q(Y), \boldsymbol{\theta})),$$

  where $k = 0, 1 \ldots M-1$.
- **Inference:** To estimate the transmitted symbol $\hat{X}$, we first take the most probable category $\hat{\gamma} = \arg\max_k \mathsf{Pr}(\gamma = k \mid f(Q(Y), \boldsymbol{\theta}^\star))$, using the optimum parameters $\boldsymbol{\theta}^\star$. Then, we set the predicted symbol $\hat{X}$ to $\hat{X} = \exp\left(\jmath\pi\left(\frac{2\hat{\gamma}+1}{M} - 1\right)\right)$.

The main logic that underpins our proposed architecture is to have a design that is capable of handling the challenges posed by low-resolution quantization, ensuring robust performance without access to channel state information. In particular, unlike conventional detectors that depend on explicit channel knowledge, NeuroDetect is *model-free*, utilizing quantized received signals $Q(Y)$ to predict the transmitted signals $\hat{X}$.

A key feature of our DL detector is that it operates without access to channel information in either the training phase (TP) or the data detection phase (DP). Our method relies on the assumption that the wireless channel varies slowly relative to the signal transmission rate, which is common in practical scenarios where coherence times are long compared to symbol durations. This slow variation allows the channel to be treated as approximately

constant throughout both the training and detection phases of signal recovery, even though it remains random and unknown to the receiver.

3.3.1. Hidden Layers of NeuroDetect

We denote the input dimension as $D_i = 1$ and the output dimension as $D_o = M$. The dimensions of the hidden layers are $D_1 = M$, $D_2 = M$, $D_3 = 4M$, and $D_4 = 2M$. We define the vector of hidden units at layer l as $\mathbf{h}_l$, the vector of biases that contribute to the hidden layer $l+1$ as $\boldsymbol{\beta}_l \in \mathbb{R}^{D_{l+1}}$, and the weights that are applied to the l-th layer and contribute to the $(l+1)$-th layer as $\boldsymbol{\Omega}_l \in \mathbb{R}^{D_{l+1} \times D_l}$, where D_l is the dimension of the corresponding lth layer. Then, the NeuroDetect network $f(Q(Y), \boldsymbol{\theta})$ can be written as

$$\begin{aligned} \mathbf{h}_1 &= \mathbf{a}_1(\boldsymbol{\beta}_0 + \boldsymbol{\Omega}_0\, Q(Y)), \\ \mathbf{h}_2 &= \mathbf{a}_2(\boldsymbol{\beta}_1 + \boldsymbol{\Omega}_1\, \mathbf{h}_1), \\ \mathbf{h}_3 &= \mathbf{a}_3(\boldsymbol{\beta}_2 + \boldsymbol{\Omega}_2\, \mathbf{h}_2), \\ \mathbf{h}_4 &= \mathbf{a}_4(\boldsymbol{\beta}_3 + \boldsymbol{\Omega}_3\, \mathbf{h}_3) \\ f(Q(Y), \boldsymbol{\theta}) &= \boldsymbol{\beta}_4 + \boldsymbol{\Omega}_4\, \mathbf{h}_4, \end{aligned}$$

where $\mathbf{a}_1(\mathbf{z}) = \tanh(\mathbf{z})$, $\mathbf{a}_2(\mathbf{z}) = \mathbf{z}$, $\mathbf{a}_3(\mathbf{z}) = \mathrm{ReLU}(\mathbf{z})$, and $\mathbf{a}_4(\mathbf{z}) = \mathrm{ReLU}(\mathbf{z})$. Here, to enable the learning of complex, nonlinear relationships between the inputs and output, we use the rectified linear unit $\mathrm{ReLU}(\mathbf{z}) = (\max(0, z_i))_{i=1}^{\dim(\mathbf{z})}$, and $\tanh(\mathbf{z}) = \left(\frac{e^{z_i} - e^{-z_i}}{e^{z_i} + e^{-z_i}}\right)_{i=1}^{\dim(\mathbf{z})}$ and $\dim(\mathbf{z})$ denote the dimension of the input vector $\mathbf{z}$. The parameters $\boldsymbol{\theta}$ of the model comprise all the weight matrices and bias vectors $\boldsymbol{\theta} = \{\boldsymbol{\beta}_l, \boldsymbol{\Omega}_l\}_{l=0}^{4}$.

3.3.2. Training Dataset Generation

We generate I training samples $\left\{\left(Q(Y)^{[i]}, \gamma^{[i]}\right)\right\}_{i=1}^{I}$ synthetically, by generating $\gamma^{[i]}$ as a uniform random integer between 0 and $M-1$, substituting $X = \exp\left(\jmath\pi\left(\frac{2\gamma^{[i]}+1}{M} - 1\right)\right)$, and generating $Q(Y)^{[i]}$ by using the channel model (1) and quantizing the received signals according to (2). The channel is assumed to be fixed (but unknown) for the training samples.

3.3.3. One-Hot Encoding of Labels

While training NeuroDetect, we use one-hot encoding to generate the output label of the i-th training sample $\mathbf{1}_{\gamma^{[i]}} \in \mathbb{R}^M$. Then the m-th element of $\mathbf{1}_{\gamma^{[i]}}$ is given by

$$\mathbf{1}_{\gamma^{[i]}}(m) = \begin{cases} 1 & \text{if } m = \gamma^{[i]} + 1, \\ 0 & \text{otherwise,} \end{cases} \tag{7}$$

where $i = 1, \ldots, I$ and $m = 1, \ldots, M$.

3.3.4. Parameter Optimization

The parameter optimization in our model-free detection framework is performed using the Adaptive Moment Estimation (ADAM) algorithm. We adopt ADAM [44] because it offers several advantages that align well with the characteristics of our problem. Specifically, ADAM adaptively adjusts the learning rate for each parameter using first- and second-order moment estimates, making it particularly effective in training scenarios involving noisy gradients, such as those arising from limited training data within each channel coherence block. Moreover, the underlying objective function in our framework is highly non-convex, and ADAM's stable convergence behavior in such settings is well established both theoretically and empirically [45,46]. These properties make ADAM especially suitable for our model, which must train efficiently from small, rapidly changing data in fading channels.

The iterative process continues until convergence, ensuring optimal weights for all layers. The ADAM algorithm adds momentum to both the estimate of the gradient and the squared gradient according to

$$\mathbf{m}_{t+1} \leftarrow \alpha_1 \cdot \mathbf{m}_t + (1-\alpha_1)\frac{\partial L(\boldsymbol{\theta}_t)}{\partial \boldsymbol{\theta}}, \tag{8}$$

$$\mathbf{v}_{t+1} \leftarrow \alpha_2 \cdot \mathbf{v}_t + (1-\alpha_2)\left(\frac{\partial L(\boldsymbol{\theta}_t)}{\partial \boldsymbol{\theta}}\right)^2, \tag{9}$$

where $\alpha_1 \in [0,1)$ and $\alpha_2 \in [0,1)$ are the momentum coefficients for the two statistics and t is the iteration round. The statistics are modified as

$$\tilde{\mathbf{m}}_{t+1} \leftarrow \frac{\mathbf{m}_{t+1}}{1-\alpha_1^{t+1}}, \tag{10}$$

$$\tilde{\mathbf{v}}_{t+1} \leftarrow \frac{\mathbf{v}_{t+1}}{1-\alpha_2^{t+1}}. \tag{11}$$

As the terms with exponents of $t+1$ decrease with each time step, the denominators approach one, causing the modification to have a diminishing effect. Finally, we update the parameters as

$$\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \alpha \cdot \frac{\tilde{\mathbf{m}}_{t+1}}{\sqrt{\tilde{\mathbf{v}}_{t+1}} + \epsilon}. \tag{12}$$

3.3.5. Hyperparameters

In our framework, the learning algorithm and its associated parameters, such as the learning rate, are treated as hyperparameters. The process of selecting optimal values for these parameters is known as hyperparameter search. For this work, we set the ADAM optimizer's parameters as follows: learning rate $\alpha = 1 \times 10^{-3}$, $\alpha_1 = 0.9$, $\alpha_2 = 0.999$, and $\epsilon = 1 \times 10^{-7}$. A detailed exploration of hyperparameter search is beyond the scope of this paper.

3.4. NeuroDetect Algorithm

Our proposed NeuroDetect algorithm, to detect input data symbols with low-resolution quantization observations, is summarized in Algorithm 1. The algorithm optimizes the model parameters $\boldsymbol{\theta}$ using the input features and one-hot encoded labels during the TP through an iterative process that minimizes the loss in (4). During the DP, the optimized parameters $\boldsymbol{\theta}^\star$ are utilized to accurately detect signals, ensuring reliable performance.

Algorithm 1 Proposed NeuroDetect Algorithm for Signal Detection

Require: I training samples $\left\{\left(Q(Y)^{[i]}, \gamma^{[i]}\right)\right\}_{i=1}^{I}$, number of iterations T.
Ensure: $\boldsymbol{\theta}^\star = \underset{\boldsymbol{\theta}}{\operatorname{argmin}}\, L[\boldsymbol{\theta}]$ in (4).

1: Initialize model parameters $\boldsymbol{\theta}_1 = \{\boldsymbol{\beta}_l, \boldsymbol{\Omega}_l\}_{l=0}^{4}$.
2: Generate one-hot encoded labels from each $\gamma^{[i]}$ by using (7).
3: **for** $t = 1$ to T **do**
4: **for** $i = 1$ to I **do**
5: $\lambda^{[i]} \leftarrow \mathsf{softmax}_{\gamma^{[i]}}\left(f\left(Q(Y)^{[i]}, \boldsymbol{\theta}_t\right)\right)$
6: **end for**
7: Compute $L[\boldsymbol{\theta}_t]$ using $\left\{\lambda^{[i]}\right\}_{i=0}^{I}$.
8: Back propagate $L(\boldsymbol{\theta}_t)$ and update parameters using (12)
9: **end for**
10: $\boldsymbol{\theta}^\star \leftarrow \boldsymbol{\theta}_T$
11: **Return** $\boldsymbol{\theta}^\star$

Remark—Digital Implementation and Bit-Depth Considerations: NeuroDetect integrates seamlessly with the conventional baseband digital signal processing (DSP) pipelines, requiring no architectural modifications beyond replacing the signal detection module with the NeuroDetect detector. In the proposed receiver architecture, analog-to-digital conversion is performed solely at the RF front end, immediately after the antenna and analog filtering stage, and this is the only point in the receiver chain where quantization occurs. NeuroDetect operates entirely in the digital domain on quantized samples produced by this front-end ADC, and can be efficiently implemented on standard digital baseband processors using conventional fixed-point arithmetic. This design preserves the energy efficiency benefits of coarse RF quantization while maintaining compatibility with the existing DSP infrastructure.

Remark—Streaming Mode Operation: NeuroDetect operates as a feed-forward neural network that processes each quantized symbol independently for data detection. It does not require sequence-level buffering or batch processing, making it well suited for streaming (online) operation. A decision is produced with a single forward pass through the network, enabling low-latency symbol detection. The network has a lightweight architecture, enabling efficient implementation in digital baseband hardware.

4. Numerical Results

In this section, we present the Symbol Error Probability (SEP) results obtained using our DL-based signal detector, NeuroDetect, for M-PSK modulation with n-bit quantization. The channel follows a unit-power, circularly symmetric fading model with a Nakagami-m distributed magnitude, while the additive noise is complex Gaussian with zero mean and unit variance.

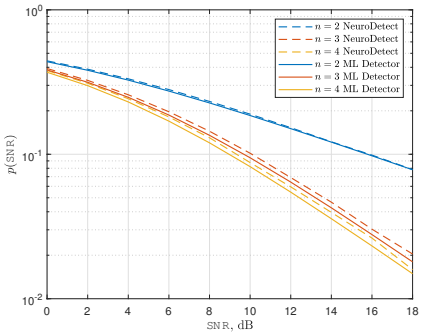
4.1. Performance Comparison with the ML Detector

Figure 4 compares the SEP performance of NeuroDetect with the ML detector from [12] for QPSK, 8-PSK, and 16-PSK modulation schemes. The comparison is made for phase quantization levels of $n = \log_2 M$, $n = \log_2 M + 1$, and $n = \log_2 M + 2$. The ML detector in [12] is the optimum detection rule if the receiver has access to full channel knowledge. Hence, it serves as the universal lower bound on the SEP performance of NeuroDetect.

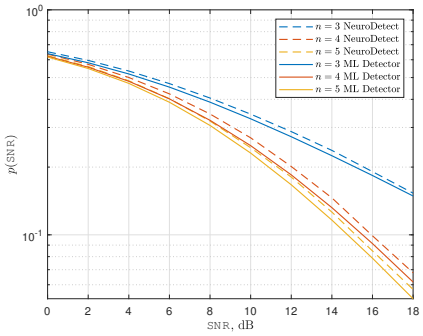
As shown in Figure 4, the proposed DL-based signal detector, NeuroDetect, achieves near-optimum performance compared to the ML detector. The main advantage of our NeuroDetect architecture is that this high accuracy is attained without any explicit channel knowledge, whereas the ML detector assumes perfect CSI. The performance gap remains below 12% across all modulation schemes and quantization levels we studied. This result is particularly significant, as accurate channel estimation poses a major challenge in low-resolution quantized receivers. NeuroDetect effectively mitigates this limitation, providing a robust solution for wireless communication systems with low-resolution quantization.

Remark—Training Overhead: To quantify the algorithmic overhead introduced by NeuroDetect, we evaluated the amount of labeled data required within each channel coherence interval. Empirically, fewer than 150 labeled symbols are sufficient to reach the near-optimum SEP reported in Figure 4. Assuming a representative urban micro-cell channel with a coherence bandwidth of $B_c = 1$ MHz and a coherence time of $T_c = 5$ ms [47], the resulting coherence block contains approximately

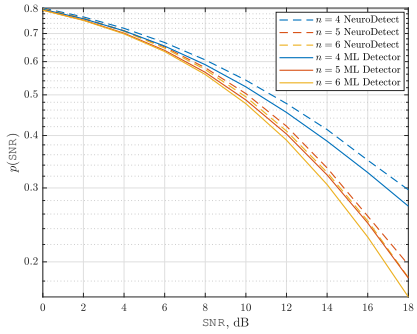
$$B_c T_c \;=\; 1 \times 10^6 \text{ Hz} \;\times\; 5 \times 10^{-3} \text{ s} \;=\; 5 \times 10^3 \text{ symbols.}$$



(**a**) QPSK



(**b**) 8-PSK



(**c**) 16-PSK

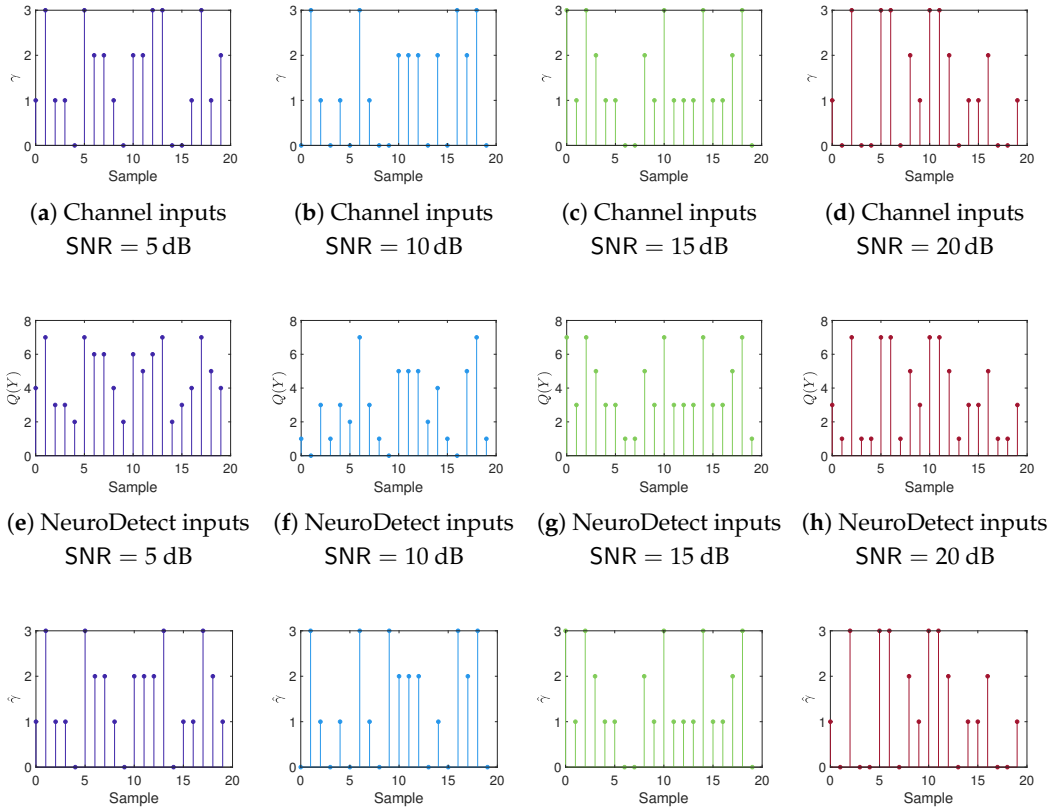
Figure 4. Comparison of the SEP performances of NeuroDetect (no channel knowledge) and the optimum ML detector (full channel knowledge). The number of quantization bits was set to $n = \log_2 M$, $\log_2 M + 1$, and $\log_2 M + 2$ for $M = 4, 8$, and 16.

Hence, the training overhead of NeuroDetect amounts to

$$\frac{150}{5000} \approx 3\%$$

of the available symbols per block. In other words, 97% of each coherence block remains available for payload data, indicating that the proposed model-free detector can be trained online with negligible rate loss while still achieving near-optimum performance, close to the ML benchmark. This modest overhead, combined with the detector's feed-forward streaming operation (see the remark on streaming mode operation above), indicates the potential of deploying NeuroDetect in latency-sensitive and bandwidth-constrained wireless systems.

To provide further insights into the time-domain behavior of NeuroDetect, we provide a graphical illustration of transmitted and received signals in discrete-time using QPSK modulation with 3-bit phase quantization in Figure 5. For the figure, the channel is fixed as $h = e^{j\frac{\pi}{6}}$. Transmitted symbols are randomly generated with a uniform distribution for each SNR level and are represented by symbol indices $\gamma \in \{0, 1, 2, 3\}$ in the first row of the figure. The second row shows the inputs to the NeuroDetect detector, which are the quantizer outputs indicating the sector of the received signal. With 3-bit quantization, these outputs take values in $\{0, 1, \ldots, 7\}$. The third row displays the outputs of the NeuroDetect detector, corresponding to the indices of the estimated transmitted symbols. A detection is considered successful when the estimated index matches the transmitted index. Otherwise, a symbol error occurs, for example, see sample 15 at SNR $= 5$ dB.



(**a**) Channel inputs SNR = 5 dB (**b**) Channel inputs SNR = 10 dB (**c**) Channel inputs SNR = 15 dB (**d**) Channel inputs SNR = 20 dB

(**e**) NeuroDetect inputs SNR = 5 dB (**f**) NeuroDetect inputs SNR = 10 dB (**g**) NeuroDetect inputs SNR = 15 dB (**h**) NeuroDetect inputs SNR = 20 dB

(**i**) NeuroDetect outputs SNR = 5 dB (**j**) NeuroDetect outputs SNR = 10 dB (**k**) NeuroDetect outputs SNR = 15 dB (**l**) NeuroDetect outputs SNR = 20 dB

Figure 5. Graphical illustration of time-domain signals under QPSK modulation and 3-bit phase quantization, assuming a fixed channel $h = e^{j\frac{\pi}{6}}$. The channel inputs are represented by transmitted symbol indices γ. The NeuroDetect inputs (or quantizer outputs) correspond to the quantized phase sectors $Q(Y)$, while the NeuroDetect outputs denote the estimated transmitted symbol indices $\hat{\gamma}$.

4.2. Effect of Quantization Bits on System Performance

In this part, we examine the impact of quantization bit levels on the SEP performance of the proposed NeuroDetect detector. Specifically, we focus on QPSK modulation and analyze the effect of quantization levels $n = 2, 3$, and 4 under Nakagami-m fading conditions, considering shape parameters $m = 1$ and $m = 2$.

A key structural property of data detection in low-resolution quantized systems under Nakagami-m fading is that the system diversity order (DVO), defined as $\mathrm{DVO} \triangleq \lim_{\mathsf{SNR}\to\infty} \frac{\log p(\mathsf{SNR})}{\log \mathsf{SNR}}$, follows a distinct behavior based on the quantization level. Specifically, $\mathrm{DVO} = \frac{1}{2}$ when $n = \log_2 M$, and $\mathrm{DVO} = m$ when $n > \log_2 M$. This result, established in [12], characterizes the optimum asymptotic decay rate of the SEP in the high-SNR regime. We will demonstrate that our proposed NeuroDetect detector is asymptotically optimum, as it attains the same diversity order as the ML detector.

Our results are presented in Figure 6. A significant performance improvement is observed for both $m = 1$ and $m = 2$ when the number of quantization bits increases from $n = 2$ to $n = 3$. This behavior aligns with the fundamental structural property of the system's DVO discussed earlier. Specifically, for $n = 2$, the diversity order remains at $\mathrm{DVO} = \frac{1}{2}$ in both cases for the QPSK modulation, leading to a relatively slow decay of SEP as a function of SNR. However, when n increases to 3, the DVO transitions from $\frac{1}{2}$ to m, and our proposed NeuroDetect detector achieves the optimum asymptotic SEP decay rate. Further increasing n from 3 to 4 yields performance improvements, though with diminishing returns, indicating a saturation effect in the benefits of additional quantization bits.

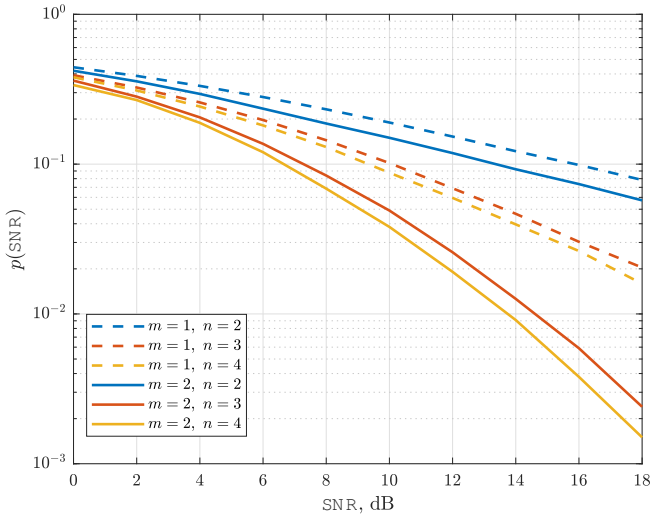


Figure 6. Average SEP curves for NeuroDetect (averaged over random channel realizations) as a function of SNR for QPSK modulation. $n = 2, 3, 4$ and $m = 1, 2$.

Figure 7 illustrates the average SEP performance of the NeuroDetect detector, averaged over random channel realizations, as a function of SNR. The results are presented for QPSK, 8-PSK, and 16-PSK modulation schemes, with the Nakagami-m shape parameter fixed at $m = 1$, corresponding to the classical Rayleigh fading scenario. The number of quantization bits is set to $n = \log_2 M$, $\log_2 M + 1$, and $\log_2 M + 2$.

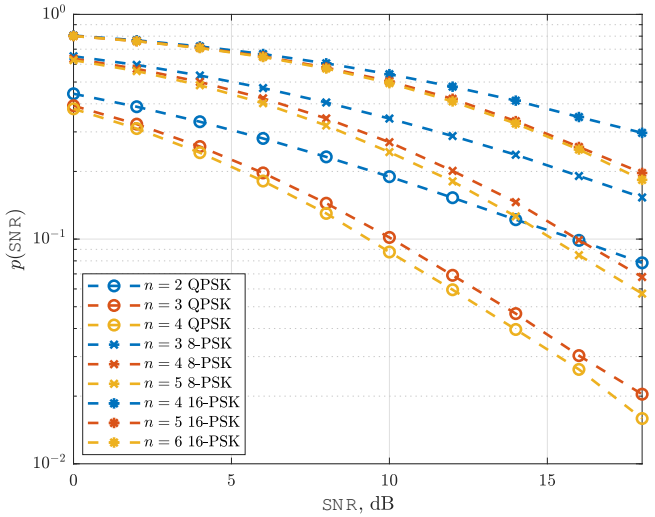


Figure 7. Average SEP curves for NeuroDetect (averaged over random channel realizations) as a function of SNR for different modulation schemes. $n = \log_2 M,\ \log_2 M + 1,\ \log_2 M + 2$, and $m = 1$.

From the Figure 7, it is evident that QPSK with 2-bit quantization, 8-PSK with 3-bit quantization, and 16-PSK with 4-bit quantization exhibit a diversity order DVO of $\frac{1}{2}$. Moreover, increasing the number of quantization bits beyond these levels, specifically QPSK with 3 or more bits, 8-PSK with 4 or more bits, and 16-PSK with 5 or more bits, results in a diversity order of 1, which is equal to m in this case. These findings confirm that the NeuroDetect detector is asymptotically optimum, as it attains the same diversity order as the ML detector in the high-SNR regime.

In Figure 8, we analyze the average SEP performance of the NeuroDetect detector as a function of SNR for 8-PSK modulation with 2-bit quantization and 16-PSK modulation with 2-bit quantization. The Nakagami-m fading parameter is set to $m = 0.5, 1$, and 2.

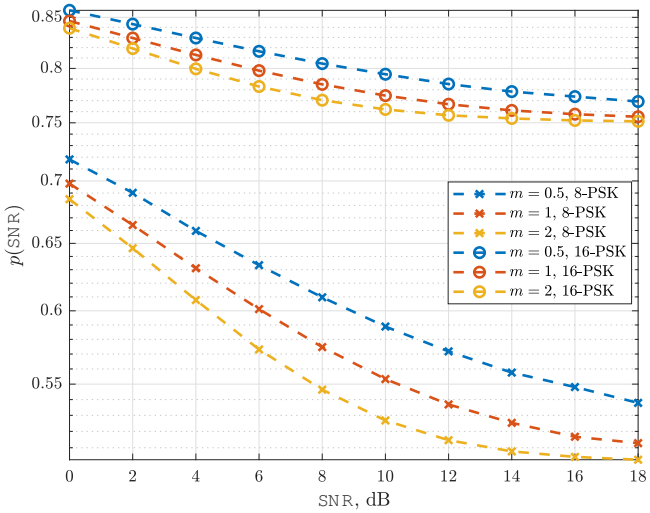


Figure 8. Average SEP curves for NeuroDetect (averaged over random channel realizations) as a function of SNR for 8-PSK and 16-PSK modulation schemes. $n = 2 < \log_2 M$, and $m = 0.5, 1$, and 2.

This scenario corresponds to a regime where the number of quantization bits is insufficient to fully resolve the transmitted symbols, leading to an inherent error floor in detection performance. This error floor phenomenon was analytically characterized in [12], which established that the average SEP for 8-PSK with 2-bit quantization is lower-bounded by 0.5, while for 16-PSK with 2-bit quantization, the lower bound is 0.75.

As shown in our results, NeuroDetect exhibits an error floor at high SNR values when $n < \log_2 M$, aligning with these theoretical limits. Moreover, in both cases, NeuroDetect SEP performance becomes close to these theoretical lower bounds.

4.3. Deep Learning Penalty Metrics

We define the increment of the average SEP as a *learning penalty* according to

$$\Psi_L(\mathrm{SNR}, n) \triangleq 10 \log\left(\frac{p(\mathrm{SNR}, n)}{p_{ML}(\mathrm{SNR}, n)}\right), \tag{13}$$

where $p(\mathrm{SNR}, n)$ is the average SEP of NeuroDetect with n number of quantization bits, and $p_{ML}(\mathrm{SNR}, n)$ is the average SEP of the ML detector with again n number of quantization bits.

Similarly, we define the increment of the average SEP for the ML detector as a quantization penalty according to

$$\Psi(\mathrm{SNR}, n) \triangleq 10 \log\left(\frac{p_{\mathrm{ML}}(\mathrm{SNR}, n)}{p_{\mathrm{ML}}(\mathrm{SNR}, \infty)}\right).$$

Using $\Psi_L(\mathrm{SNR}, n)$ and $\Psi(\mathrm{SNR}, n)$, the total penalty for the NeuroDetect detector, which is from both the learning penalty and the quantization penalty, is defined according to

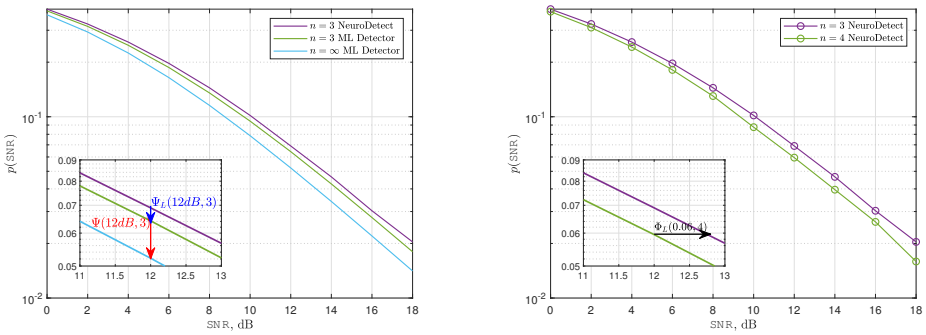
$$\begin{aligned} \Psi_T(\mathrm{SNR}, n) &\triangleq \Psi_L(\mathrm{SNR}, n) + \Psi(\mathrm{SNR}, n) \\ &= 10 \log\left(\frac{p(\mathrm{SNR}, n)}{p_{\mathrm{ML}}(\mathrm{SNR}, \infty)}\right). \end{aligned} \tag{14}$$

In addition to $\Psi_T(\mathsf{SNR}, n)$, we also define *power penalty* as an increase in transmission power to achieve the same average SEP with quantization levels of $n-1$ and n bits, which is given by

$$\Phi_L(\mathsf{SEP}, n) = 10 \log\left(\frac{\mathsf{SNR}_{n-1}}{\mathsf{SNR}_n}\right), \tag{15}$$

where SNR_n and SNR_{n-1} are the SNR values required to achieve a target average SEP with n and $n-1$ quantization bits using NeuroDetect, respectively.

We illustrate the learning penalty metric for SEP in (13) for NeuroDetect with a 3-bit phase quantizer at an SNR of 12 dB in Figure 9a. Specifically, the SEP learning penalty $\Psi_L(12\text{ dB}, 3)$ is 0.31 dB. Due to the quantization process, there is an additional quantization penalty of $\Psi(12\text{ dB}, 3) = 0.9$ dB. Therefore, the total penalty when using a 3-bit phase quantizer with NeuroDetect is $\Psi_T(12\text{ dB}, 3) = 1.21$ dB.



(**a**) SEP Penalty in (13) and (14) (**b**) Power penalty in (15)

Figure 9. Illustration of deep learning penalty metrics for QPSK.

We also illustrate the power penalty metric in (15) for NeuroDetect in Figure 9b. In particular, when the target average SEP is 0.06, we observe a power penalty of $\Phi_L(0.06, 4) \approx 0.25$ dB as we change the quantization level from $n = 3$ to 4 for the NeuroDetect. This result implies that to achieve an average SEP of 0.06 with 3-bit quantization, we need 0.25 dB more transmit power than what is required with 4-bit quantization when we use the NeuroDetect.

Figure 10 illustrates the two forms of SEP penalties: the learning penalty from (13) and the total penalty, which combines the learning and quantization penalties, from (14), as a function of SNR for QPSK. When considering the learning penalty, NeuroDetect performs similarly to the ML detector with a learning penalty smaller than 0.4 dB. The gap is smaller and bounded by 0.2 dB when the ADC has $n = \log_2 M$ bits. This is because, as the number of ADC bits increases, the number of input label categories for NeuroDetect increases to 2^n, while the classification task remains fixed at four classes for QPSK. As a result, mapping 2^n input labels to four classes becomes more complex, making it harder to learn the relationship between $Q(Y)$ and the most likely transmitted symbol.

We observe a different trend for the total penalty, which also depends on the quantization penalty. As the number of ADC bits increases, the estimation accuracy improves and the quantization penalty decreases. Consequently, the combined learning and quantization penalty decreases with higher ADC bit resolution, eventually converging to the high-resolution case.

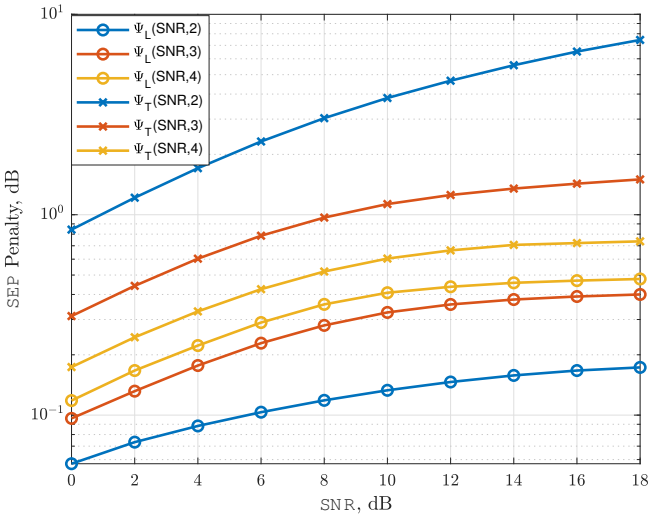


Figure 10. SEP Learning penalty in (13) and SEP learning penalty + quantization penalty in (14) as function of SNR for QPSK.

4.4. Impact of Channel Mismatch Between Training and Data Detection Phases

Finally, we evaluate the robustness of NeuroDetect by employing different channels for training and data detection. Specifically, we introduce phase deviations $\phi = \frac{\pi}{32}, \frac{\pi}{16}, \frac{\pi}{8}$, and $\frac{\pi}{4}$ relative to the channel used in the training phase, allowing us to assess the detector's performance under channel mismatches.

Figure 11 shows that NeuroDetect maintains performance comparable to the no-mismatch scenario for channel phase deviations up to $\frac{\pi}{8}$. Within this range, only a constant offset is observed relative to the no-mismatch case. However, for larger deviations, particularly $\phi = \frac{\pi}{4}$, an error floor emerges. In this setting, the mismatch between the channel used for training and the channel encountered during data detection constrains NeuroDetect's detection performance. In particular, when the channel phase during data detection deviates significantly from the conditions used in training, NeuroDetect encounters *out-of-distribution* samples for which its learned decision boundaries are no longer valid, ultimately resulting in the observed error floor phenomenon.

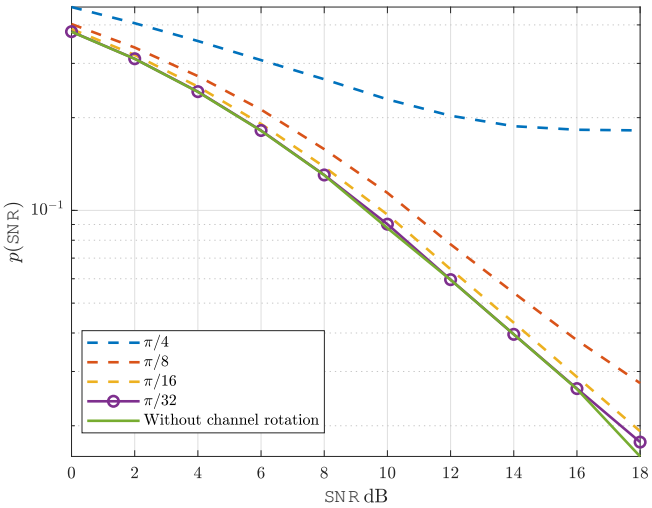


Figure 11. Impact of channel mismatch between training and data detection phases.

5. Conclusions and Future Work

In this paper, we introduced *NeuroDetect*, a deep learning-based signal detection framework for Nakagami-*m* fading channels with low-resolution quantization at the receiver. By operating in a model-free manner, NeuroDetect eliminates the need for explicit channel knowledge, which is often difficult to estimate under low-resolution constraints. Instead, it leverages training data acquired during each channel coherence interval to learn an optimum signal detection strategy, which is then used for symbol detection during the data transmission phase. Training requires fewer than 150 symbols to achieve near-optimum performance, which corresponds to 3% of a typical coherence block of 5000 symbols (assuming a bandwidth of 1 MHz and a coherence time of 5 ms).

Our numerical findings show that NeuroDetect achieves symbol error performance nearly on par with the maximum likelihood detector, which assumes full channel knowledge. In all examined scenarios, the performance gap remains below 12%. We also show that NeuroDetect attains the optimum asymptotic error decay rates, featuring a fundamental ternary behavior based on the number of quantization bits. These results demonstrate that our proposed model-free learning architecture can closely match the performance of optimum analytical detection strategies, while circumventing the challenge of accurate channel estimation in wireless communications systems with low-resolution quantization.

Our proposed NeuroDetect architecture is already lightweight, comprising five fully connected layers with modest widths determined by the modulation order. An important direction for future work is the development of even more compact architectures that preserve near-optimum detection performance under low-resolution quantization. Such designs would further enhance the practicality of model-free detection in resource-constrained hardware. Investigating this trade-off between model simplicity and detection accuracy is vital to enable scalable deployments of learning-based detectors. Future work may also include hardware-in-the-loop testing to quantify NeuroDetect's resilience to practical ADC impairments, such as comparator nonlinearities, offset drift, and timing jitter.

Author Contributions: Conceptualization, S.G. and H.I.; methodology, S.G., H.I., C.L., R.Z. and D.A.; validation, C.L., S.G., H.I., R.Z. and D.A.; formal analysis, C.L., S.G. and H.I.; writing—original draft preparation, C.L. and S.G.; writing—review and editing, S.G., H.I., R.Z. and D.A.; supervision, H.I. and S.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data are contained within the article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ADC	Analog-to-Digital Converter
DAC	Digital-to-Analog Converter
AWGN	Additive White Gaussian Noise
CSI	Channel State Information
DL	Deep Learning
DNN	Deep Neural Network
ML	Maximum Likelihood

MIMO	Multiple-Input Multiple-Output
PSK	Phase Shift Keying
QPSK	Quadrature Phase Shift Keying
SNR	Signal-to-Noise Ratio
LMMSE	Linear Minimum-Mean-Square-Error
BLMMSE	Bussgang LMMSE
MRC	Maximum Ratio Combining
ZF	Zero-Forcing
OSD	One-bit Sphere Decoding
e-MLD	Empirical Maximum Likelihood Detection
MMD	Minimum Mean-Distance
MCD	Minimum Center Detection
FBMCENet	Few-Bit Massive MIMO Channel Estimation Network
FBM-DetNet	Few-Bit Massive MIMO Data Detection Network
OBMNet	One-Bit Massive MIMO Data Detection Network
LoRD-Net	Low Resolution Detection Network
BiLSTM	Bidirectional Long Short-Term Memory
ReLU	Rectified Linear Unit
ADAM	Adaptive Moment Estimation
TP	Training Phase
DP	Data Transmission Phase
SEP	Symbol Error Probability
DVO	Diversity Order

References

1. Lin, X.; Li, J.; Baldemair, R.; Cheng, J.F.T.; Parkvall, S.; Larsson, D.C.; Koorapaty, H.; Frenne, M.; Falahati, S.; Grovlen, A.; et al. 5G new radio: Unveiling the essentials of the next generation wireless access technology. *IEEE Commun. Stand. Mag.* **2019**, *3*, 30–37. [CrossRef]
2. Wang, C.X.; You, X.; Gao, X.; Zhu, X.; Li, Z.; Zhang, C.; Wang, H.; Huang, Y.; Chen, Y.; Haas, H.; et al. On the road to 6G: Visions, requirements, key technologies, and testbeds. *IEEE Commun. Surv. Tutor.* **2023**, *25*, 905–974. [CrossRef]
3. Ning, B.; Tian, Z.; Mei, W.; Chen, Z.; Han, C.; Li, S.; Yuan, J.; Zhang, R. Beamforming technologies for ultra-massive MIMO in terahertz communications. *IEEE Open J. Commun. Soc.* **2023**, *4*, 614–658. [CrossRef]
4. Bai, Q.; Nossek, J.A. Energy efficiency maximization for 5G multi-antenna receivers. *Trans. Emerg. Telecommun. Technol.* **2015**, *26*, 3–14. [CrossRef]
5. Gayan, S.; Senanayake, R.; Inaltekin, H.; Evans, J. Reliability Characterization for SIMO Communication Systems with Low-Resolution Phase Quantization Under Rayleigh Fading. *IEEE Open J. Commun. Soc.* **2021**, *2*, 2660–2679. [CrossRef]
6. Walden, R.H. Performance trends for analog to digital converters. *IEEE Commun. Mag.* **1999**, *37*, 96–101. [CrossRef]
7. Zhang, J.; Dai, L.; Li, X.; Liu, Y.; Hanzo, L. On low-resolution ADCs in practical 5G millimeter-wave massive MIMO systems. *IEEE Commun. Mag.* **2018**, *56*, 205–211. [CrossRef]
8. Ivrlac, M.T.; Nossek, J.A. On MIMO channel estimation with single-bit signal-quantization. In Proceedings of the ITG Smart Antenna Workshop, Vienna, Austria, 26–27 February 2007.
9. Mo, J.; Alkhateeb, A.; Abu-Surra, S.; Heath, R.W. Hybrid architectures with few-bit ADC receivers: Achievable rates and energy-rate tradeoffs. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 2274–2287. [CrossRef]
10. Gayan, S.; Inaltekin, H.; Senanayake, R.; Evans, J. The Cost of Noncoherence: Avoiding Channel Estimation Through Differential Encoding in Phase Quantized Systems. *IEEE Open J. Commun. Soc.* **2023**, *4*, 2578–2595. [CrossRef]
11. Gayan, S.; Senanayake, R.; Inaltekin, H.; Evans, J. Selection Combining for Multi-Antenna Communication with Low-Resolution ADCs. In Proceedings of the 2021 IEEE International Symposium on Information Theory (ISIT), Melbourne, Australia, 12–20 July 2021; pp. 3343–3348.
12. Gayan, S.; Senanayake, R.; Inaltekin, H.; Evans, J. Low-resolution quantization in phase modulated systems: Optimum detectors and error rate analysis. *IEEE Open J. Commun. Soc.* **2020**, *1*, 1000–1021. [CrossRef]
13. Jeon, Y.S.; Kim, D.; Hong, S.N.; Lee, N.; Heath, R.W. Artificial intelligence for physical-layer design of MIMO communications with one-bit ADCs. *IEEE Commun. Mag.* **2022**, *60*, 76–81. [CrossRef]
14. Nguyen, N.T.; Lee, K. Deep learning-aided tabu search detection for large MIMO systems. *IEEE Trans. Wirel. Commun.* **2020**, *19*, 4262–4275. [CrossRef]

15. Mezghani, A.; Koufi, M.; Nossek, J.A. A modified MMSE receiver for quantized MIMO systems. In Proceedings of the ITG/IEEE Workshop on Smart Antennas, Vienna, Austria, 26–27 February 2007.
16. Li, Y.; Tao, C.; Seco-Granados, G.; Mezghani, A.; Swindlehurst, A.L.; Liu, L. Channel estimation and performance analysis of one-bit massive MIMO systems. *IEEE Trans. Signal Process.* **2017**, *65*, 4075–4089. [CrossRef]
17. Jacobsson, S.; Durisi, G.; Coldrey, M.; Gustavsson, U.; Studer, C. Throughput analysis of massive MIMO uplink with low-resolution ADCs. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 4038–4051. [CrossRef]
18. Choi, J.; Love, D.J.; Brown, D.R.; Boutin, M. Quantized distributed reception for MIMO wireless systems using spatial multiplexing. *IEEE Trans. Signal Process.* **2015**, *63*, 3537–3548. [CrossRef]
19. Yu, C.Y.; Shao, M.; Chen, W.K.; Liu, Y.F.; Ma, W.K. An Efficient Global Algorithm for One-Bit Maximum-Likelihood MIMO Detection. In Proceedings of the 2023 IEEE 24th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Shanghai, China, 25–28 September 2023; pp. 231–235.
20. Shao, M.; Ma, W.K. Divide and Conquer: One-bit MIMO-OFDM Detection by Inexact Expectation Maximization. In Proceedings of the ICASSP 2021—2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Toronto, ON, Canada, 6–11 June 2021; pp. 4890–4894.
21. Choi, J.; Mo, J.; Heath, R.W. Near maximum-likelihood detector and channel estimator for uplink multiuser massive MIMO systems with one-bit ADCs. *IEEE Trans. Commun.* **2016**, *64*, 2005–2018. [CrossRef]
22. Jeon, Y.S.; Lee, N.; Hong, S.N.; Heath, R.W. One-bit sphere decoding for uplink massive MIMO systems with one-bit ADCs. *IEEE Trans. Wirel. Commun.* **2018**, *17*, 4509–4521. [CrossRef]
23. Safa, K.; Combes, R.; De Lacerda, R.; Yang, S. Data detection in 1-bit quantized MIMO systems. *IEEE Trans. Commun.* **2024**, *72*, 5396–5410. [CrossRef]
24. Combes, R.; Yang, S. An approximate ML detector for MIMO channels corrupted by phase noise. *IEEE Trans. Commun.* **2017**, *66*, 1176–1189. [CrossRef]
25. Gayan, S.; Inaltekin, H.; Senanayake, R.; Evans, J. Phase modulated communication with low-resolution ADCs. In Proceedings of the ICC 2019—2019 IEEE International Conference on Communications (ICC), Shanghai, China, 20–24 May 2019; pp. 1–7.
26. Singh, J.; Madhow, U. Phase-quantized block noncoherent communication. *IEEE Trans. Commun.* **2013**, *61*, 2828–2839. [CrossRef]
27. Bernardo, N.I.; Zhu, J.; Evans, J. On the Capacity-Achieving Input of Channels With Phase Quantization. *IEEE Trans. Inf. Theory* **2022**, *68*, 5866–5888. [CrossRef]
28. Ye, H.; Li, G.Y.; Juang, B.H. Power of deep learning for channel estimation and signal detection in OFDM systems. *IEEE Wirel. Commun. Lett.* **2017**, *7*, 114–117. [CrossRef]
29. O'shea, T.; Hoydis, J. An introduction to deep learning for the physical layer. *IEEE Trans. Cogn. Commun. Netw.* **2017**, *3*, 563–575. [CrossRef]
30. Gruber, T.; Cammerer, S.; Hoydis, J.; Ten Brink, S. On deep learning-based channel decoding. In Proceedings of the 2017 51st annual conference on information sciences and systems (CISS), Baltimore, MD, USA, 22–24 March 2017; pp. 1–6.
31. Qin, Z.; Ye, H.; Li, G.Y.; Juang, B.H.F. Deep learning in physical layer communications. *IEEE Wirel. Commun.* **2019**, *26*, 93–99. [CrossRef]
32. Ye, H.; Liang, L.; Li, G.Y.; Juang, B.H. Deep learning-based end-to-end wireless communication systems with conditional GANs as unknown channels. *IEEE Trans. Wirel. Commun.* **2020**, *19*, 3133–3143. [CrossRef]
33. Jeon, Y.S.; Hong, S.N.; Lee, N. Blind detection for MIMO systems with low-resolution ADCs using supervised learning. In Proceedings of the 2017 IEEE International Conference on Communications (ICC), Paris, France, 21–25 May 2017; pp. 1–6.
34. Jeon, Y.S.; Hong, S.N.; Lee, N. Supervised-learning-aided communication framework for MIMO systems with low-resolution ADCs. *IEEE Trans. Veh. Technol.* **2018**, *67*, 7299–7313. [CrossRef]
35. Nguyen, L.V.; Ngo, D.T.; Tran, N.H.; Swindlehurst, A.L.; Nguyen, D.H. Supervised and semi-supervised learning for MIMO blind detection with low-resolution ADCs. *IEEE Trans. Wirel. Commun.* **2020**, *19*, 2427–2442. [CrossRef]
36. Nguyen, L.V.; Nguyen, D.H.; Swindlehurst, A.L. DNN-based detectors for massive MIMO systems with low-resolution ADCs. In Proceedings of the ICC 2021-IEEE International Conference on Communications, Montreal, QC, Canada, 14–23 June 2021; pp. 1–6.
37. Nguyen, L.V.; Swindlehurst, A.L.; Nguyen, D.H. Linear and deep neural network-based receivers for massive MIMO systems with one-bit ADCs. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 7333–7345. [CrossRef]
38. Khobahi, S.; Shlezinger, N.; Soltanalian, M.; Eldar, Y.C. LoRD-Net: Unfolded deep detection network with low-resolution receivers. *IEEE Trans. Signal Process.* **2021**, *69*, 5651–5664. [CrossRef]
39. Doshi, A.; Andrews, J.G. One-bit mmWave MIMO channel estimation using deep generative networks. *IEEE Wirel. Commun. Lett.* **2023**, *12*, 1593–1597. [CrossRef]
40. Rahman, M.H.; Sejan, M.A.S.; Aziz, M.A.; Tabassum, R.; Baik, J.I.; Song, H.K. Deep learning based one bit-ADCs efficient channel estimation using fewer pilots overhead for massive MIMO system. *IEEE Access* **2024**, *12*, 64823–64836. [CrossRef]
41. Le, H.A.; Van Chien, T.; Nguyen, T.H.; Choo, H.; Nguyen, V.D. Machine learning-based 5G-and-beyond channel estimation for MIMO-OFDM communication systems. *Sensors* **2021**, *21*, 4861. [CrossRef]

42. Helmy, I.; Tarafder, P.; Choi, W. LSTM-GRU model-based channel prediction for one-bit massive MIMO system. *IEEE Trans. Veh. Technol.* **2023**, *72*, 11053–11057. [CrossRef]
43. Kim, S.; Chae, J.; Hong, S.N. Machine learning detectors for MU-MIMO systems with one-bit ADCs. *IEEE Access* **2020**, *8*, 86608–86616. [CrossRef]
44. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
45. Reddi, S.J.; Kale, S.; Kumar, S. On the convergence of adam and beyond. *arXiv* **2019**, arXiv:1904.09237.
46. Loshchilov, I.; Hutter, F. Decoupled weight decay regularization. *arXiv* **2017**, arXiv:1711.05101.
47. Tse, D.; Viswanath, P. *Fundamentals of Wireless Communication*; Cambridge University Press: New York, NY, USA, 2005.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.