

# A Study of Interference Suppression Using Deep Learning Methods

Badhrinarayan Malolan

Thesis submitted to the Faculty of the  
Virginia Polytechnic Institute and State University  
in partial fulfillment of the requirements for the degree of

Master of Science  
in  
Computer Engineering

Jeffrey H. Reed, Chair  
Daniel J. Jakubisin, Co-Chair  
William C. Headley

February 20, 2025  
Blacksburg, Virginia

Keywords: Deep Learning, Wireless Communication, Autoencoder, Cyclostationarity

Copyright 2025, Badhrinarayan Malolan

# A Study of Interference Suppression Using Deep Learning Methods

Badhrinarayan Malolan

## ABSTRACT

This thesis investigates a Deep Learning model for interference suppression in wireless communications. By exploiting the structure of Convolutional Neural Network-based autoencoders, we develop an approach for interference suppression with no prior knowledge on characteristics or the exact location of interference. Traditional interference suppression techniques are heavily reliant on specific domain knowledge, thus their applicability in dynamic wireless environments is limited. This thesis proposes a CNN-AE (Convolutional Neural Network - Autoencoder) model that consists of an encoder, which captures the latent space representation from the input data, and a decoder that reconstructs the desired signal to suppress interference effects. We investigate the performance of a QPSK-based wireless communication system with three explicit interference scenarios, namely, single frequency tone interference with two cases of in-bandwidth and out-of-bandwidth, and wideband interference from a dataset that captured over the air communication signals. A study is performed for different SNR values along with the SINR values to observe the effectiveness of the approach at different levels. The results of our approach are quantified using popular metrics such as bit error rate (BER), error vector magnitude (EVM), and Signal to Noise-Interference Ratio (SINR). The proposed model outperforms the baselines with classical techniques such as matched filtering and least squares adaptive filtering consistently over these several metrics. The thesis also investigates the latent space behavior of the autoencoder; which is used to provide an interpretation of how the network classifies

between the desired signal and interference. We use this contextual information to pursue future directions in interference suppression performance by exploiting cyclostationarity properties of our desired signal to our advantage. One of the important contributions in this work involves carrying out thorough analysis with respect to the generalization capability of CNN-AE for different types of interference and signal conditions. The results presented herein illustrate the potential of a deep learning-based approach in enabling more robust and adaptive wireless communication systems that would be capable of autonomously managing complex interference scenarios without human intervention.

# **A Study of Interference Suppression Using Deep Learning Methods**

Badhrinarayan Malolan

## **GENERAL AUDIENCE ABSTRACT**

In this thesis, we explore Deep Learning based approaches to minimize interference in wireless communications using a specialized type of neural network called a Convolutional neural network (CNN) autoencoder. Unlike traditional methods that require specific knowledge about the interference, our approach learns to suppress interference directly from the data. This makes it more adaptable to different wireless environments that have varying interference patterns. The autoencoder model consists of an encoder that compresses the data and a decoder that reconstructs the desired signal while removing interference. Our results show that this method performs exceptionally well in various interference scenarios. We measure its effectiveness using common metrics and find that it outperforms traditional methods like matched filtering and least squares adaptive filtering. Additionally, we investigate how the autoencoder works internally, its ability to generalize to new data, and the training methods we used. Finally we explore new avenues of cyclostationary signal processing to boost the performance and real life applicability of our framework.

*Dedicated to my parents and grandparents.*

# Acknowledgments

In this section, I would like to humbly extend my thanks to everyone who has been part of my journey at VT.

To begin, I must express my deep gratitude to Dr. Jeffrey Reed for serving as the chair of my thesis committee and for welcoming me as a research student. Under his mentorship at Wireless@VT, I have grown significantly as a researcher and a professional, performing far beyond what I imagined possible. It has been a true honor to work under someone of his caliber.

I am also immensely grateful to Dr. Daniel Jakubisin, the co-chair of my committee. His hands-on approach to student advisement, coupled with his extensive technical knowledge helped clear numerous obstacles while encouraging me to pursue my own research directions. His supportive guidance has been invaluable throughout this process. I also extend my sincere thanks to Dr. William Headley for being a part of my committee and journey.

Special thanks go to Taiwo for the productive research collaborations and insightful conversations about academia and life. I am also thankful to the members of the Wireless@VT lab—Anand, Manoj, Gaurav—for being so welcoming, and for all the engaging discussions.

I owe profound thanks to my friends and roommates—Ankit, Beenaa, Pranav, Pranjal, and Sarvesh—for their support and companionship. Their presence has made my master’s journey both exciting and truly memorable.

Finally, and most significantly, I owe my deepest gratitude to my parents for their boundless love and unfailing support. Their constant encouragement and trust in my capabilities have been my greatest source of strength throughout my journey in graduate school.

# Contents

- List of Figures xiii
  
- List of Tables xv
  
- List of Abbreviations xvi
  
- 1 Introduction 1**
  - 1.1 Background . . . . . 1
  - 1.2 Motivation . . . . . 3
  - 1.3 Related Work . . . . . 5
    - 1.3.1 Traditional Interference Suppression Techniques . . . . . 5
    - 1.3.2 Machine Learning-based Interference Suppression Techniques . . . . . 6
  - 1.4 Contribution . . . . . 13
  - 1.5 Organization of Thesis . . . . . 14
  
- 2 System Model 16**
  - 2.1 Description . . . . . 17
  - 2.2 Composite Signal Model with Noise . . . . . 17
  - 2.3 Signal of Interest (SOI) Model . . . . . 18

2.3.1	QPSK Signal Generation . . . . .	18
2.4	Interference Models . . . . .	20
2.4.1	Tone Interference Model . . . . .	20
2.4.2	Wideband Interference (CommSignal2) . . . . .	23
2.4.3	QPSK Signal Interference . . . . .	24
2.5	Summary . . . . .	25
<b>3</b>	<b>Methodology</b>	<b>26</b>
3.1	Dataset Generation . . . . .	26
3.1.1	Baseband QPSK Signal and Noise/Interference configuration . . . . .	27
3.1.2	Interference Synthesis . . . . .	27
3.1.3	Noise and Power Calibration . . . . .	27
3.1.4	Data Partitioning . . . . .	28
3.2	Data Preprocessing . . . . .	28
3.2.1	Signal Conditioning and Interference Injection . . . . .	28
3.2.2	Feature Engineering . . . . .	29
3.3	CNN-AE Architecture . . . . .	31
3.3.1	Encoder (Hierarchical Feature Extraction) . . . . .	32
3.3.2	Decoder: Signal Reconstruction via Transposed Convolutions . . . . .	33
3.3.3	Skip Connections: Bridging Multi-Resolution Features . . . . .	33



3.3.4	Architectural Hyperparameters . . . . .	34
3.4	Training & Optimization . . . . .	34
3.4.1	Loss Formulation and Gradient Dynamics . . . . .	34
3.4.2	Adaptive Learning and Regularization . . . . .	35
3.4.3	Training Protocol . . . . .	35
3.4.4	Interference suppression mechanics . . . . .	35
3.5	Evaluation Metrics for Performance Assessment . . . . .	36
3.5.1	Error Vector Magnitude Analysis . . . . .	36
3.5.2	Signal to Noise—Interference Ratio Evaluation . . . . .	38
3.5.3	Bit Error Rate Performance Analysis . . . . .	39
<b>4</b>	<b>Results for CNN-AE Interference Suppression</b>	<b>41</b>
4.1	Latent Space Analysis . . . . .	41
4.1.1	Impact of Input Representation on learned features in Latent Space . . . . .	42
4.1.2	Analysis of Tone Frequency Clustering in Latent Space . . . . .	43
4.1.3	Robustness to Non-Stationary Interference . . . . .	44
4.1.4	Insights from Latent Space geometry . . . . .	48
4.2	Comparison with Conventional Interference Suppression approaches . . . . .	48
4.2.1	Matched Filtering . . . . .	49
4.2.2	Adaptive Least Squares Filtering . . . . .	50

4.3	Power Spectral Density Analysis . . . . .	50
4.3.1	In-Band Tone Interference Suppression . . . . .	51
4.3.2	Out-of-Band Tone Interference Mitigation . . . . .	51
4.3.3	Wideband ( <i>CommSignal2</i> ) Interference Rejection . . . . .	53
4.3.4	Power Spectral Density Performance Summary . . . . .	53
4.4	Bit Error Rate Analysis . . . . .	54
4.4.1	Input SIR versus BER Characteristics . . . . .	54
4.4.2	Input SNR versus BER Characteristics . . . . .	55
4.5	Error Vector Magnitude Analysis . . . . .	59
4.5.1	Input SIR versus EVM Characteristics . . . . .	59
4.5.2	Input SNR versus EVM Characteristics . . . . .	60
4.6	Signal to Interference-plus-Noise Ratio Analysis . . . . .	63
4.6.1	Input SIR versus SINR Characteristics . . . . .	63
4.6.2	Input SNR versus SINR Characteristics . . . . .	64
4.7	Key Insights . . . . .	67
4.7.1	Changing Input SIR . . . . .	67
4.7.2	Changing Input SNR . . . . .	68
4.8	Generalizability of CNN-AE to varying SIR . . . . .	69
4.8.1	Experimental Setup . . . . .	69
4.8.2	Performance Analysis . . . . .	71

4.8.3	Implications for Practical Deployment . . . . .	73
4.9	Conclusion . . . . .	73
<b>5</b>	<b>Deep Learning Interference Suppression for Cyclostationary Signals</b>	<b>76</b>
5.1	Introduction . . . . .	76
5.2	Foundations of Cyclostationarity . . . . .	77
5.3	Signal Model for Cyclostationary Processing . . . . .	79
5.3.1	Interference Configurations . . . . .	79
5.4	Methodology . . . . .	81
5.4.1	Preprocessing and Data Preparation . . . . .	81
5.4.2	Model Architecture and Training Setup . . . . .	88
5.5	Experimental Evaluation and Discussion . . . . .	93
5.5.1	Interference Suppression Performance . . . . .	93
5.5.2	Computational Efficiency and Real-Time Viability . . . . .	95
5.5.3	Impact of Windowing strategy . . . . .	95
5.5.4	Loss Curve Convergence Analysis . . . . .	96
5.5.5	Model Complexity Analysis . . . . .	98
5.6	Conclusion . . . . .	99
<b>6</b>	<b>Conclusion</b>	<b>102</b>
6.1	Interference Suppression using CNN-AE . . . . .	102

6.2	Cyclostationarity - Theory and Practice . . . . .	103
6.3	Practical Impact and Considerations . . . . .	104
6.4	Future Directions . . . . .	104
	<b>Bibliography</b>	<b>107</b>

# List of Figures

2.1	Outline of steps in System Model . . . . .	16
2.2	In-Band Tone Interference Comparison . . . . .	22
2.3	Out-of-Band Tone Interference Comparison . . . . .	22
2.4	<i>CommSignal2</i> Interference Comparison . . . . .	22
2.5	Extracted Frame of <i>CommSignal2</i> Interference . . . . .	24
3.1	Overview of the System Model . . . . .	31
4.1	t-SNE clustering of latent space representation of time domain QPSK signals by interference type. . . . .	45
4.2	t-SNE clustering of latent space representation of FFT domain QPSK signals by interference type. . . . .	46
4.3	t-SNE clustering of tone frequency interfered QPSK signals by frequency range.	47
4.4	Power Spectral Density (PSD) for Interference Suppression Methods . . . . .	52
4.5	BER Performance across SIR Levels for Interference Suppression Methods .	56
4.6	BER Performance across SNR Levels for Interference Suppression Methods .	57
4.7	EVM Performance across SIR Levels for Interference Suppression Methods .	61
4.8	EVM Performance across SNR Levels for Interference Suppression Methods .	62
4.9	SINR Performance Across SIR Levels for Interference Suppression Methods .	65

4.10	SINR Performance Across SNR Levels for Interference Suppression Methods	66
4.11	Generalizability Analysis of CNN-AE . . . . .	70
5.1	QPSK SOI interfered with <i>CommSignal2</i> SNOI Comparison . . . . .	80
5.2	QPSK SOI interfered with QPSK SNOI Comparison . . . . .	80
5.3	Cyclostationary analysis pipeline from preprocessing to reconstruction. . . . .	85
5.4	Cyclostationary Training Pipeline . . . . .	88
5.5	Small CNN-AE Architecture . . . . .	89
5.6	Loss Curve convergence comparison by interference type. . . . .	96

# List of Tables

3.1	List of QPSK signal generation parameters . . . . .	26
3.2	List of architectural specifications for CNN-AE . . . . .	32
5.1	CommSignal2 Interference Suppression Performance (Window Size - 17) . . . . .	93
5.2	QPSK Interference Suppression Performance (Window Size - 16) . . . . .	94
5.3	QPSK Interference (2 sample delay) Suppression Performance (Window Size - 16) . . . . .	94
5.4	Model Architecture Complexity Comparison . . . . .	98

# List of Abbreviations

5G NR Fifth Generation New Radio

AE Autoencoder

BER Bit Error Rate

BSS Blind Source Separation

CAF Cyclic Autocorrelation Function

CNN Convolutional Neural Network

CW Continuous-Wave

DL Deep Learning

EVM Error Vector Magnitude

FIR Finite Impulse Response

FMCW Frequency-Modulated Continuous Wave

GANs Generative Adversarial Networks

I/Q In-phase/Quadrature

IoT Internet of Things

ISI Inter-Symbol Interference

LeakyReLU Leaky Rectified Linear Unit



LP Linear Prediction

ML Machine Learning

MSE Mean Squared Error

NLP Natural Language Processing

OFDM Orthogonal Frequency Division Multiplexing

OTA Over-The-Air

PSD Power Spectral Density

QPSK Quadrature Shift Phase Keying

SINR Signal-to-Interference-plus-Noise Ratio

SIR Signal-to-Interference Ratio

SNR Signal-to-Noise Ratio

STFT Short-Time Fourier Transform

TV-FRESH Time-Varying Frequency Shift Filters

V2X Vehicle-to-Everything

# Chapter 1

## Introduction

### 1.1 Background

In today's rapidly evolving digital landscape, wireless communication systems face a fundamental challenge: signal interference. As wireless technologies have become increasingly ubiquitous, the ability to effectively suppress interference is now considered to be a critical parameter of system performance and reliability. Interference in the context of wireless communication can be described as unwanted signals that corrupt or degrade the intended transmission signal. This effect poses a significant problem to the fundamental objectives of modern wireless communication systems. These disruptive signals can emanate from various sources, including wireless transmissions from other devices, adverse atmospheric conditions, hardware limitations, and electromagnetic disturbances in the transmission medium. The contribution of this thesis focuses on the problem of co-channel interference from communication signals.

The ramifications of interference extend far beyond mere technical considerations, directly impacting the user experience and operational efficiency of wireless networks. When interference compromises signal quality, it triggers a cascade of adverse effects: diminished data throughput, elevated bit error rates, inefficient spectrum utilization, and increased transmission latency. These challenges become particularly acute as unprecedented growth is seen in wireless applications, from next-generation mobile networks (5G/6G) to industrial

applications of Internet of Things (IoT), connected vehicles, and time-sensitive networking applications. The ability to maintain reliable communication channels through effective interference suppression thus becomes crucial in realizing the full potential of these technologies.

The real-world manifestations of co-channel interference in modern wireless communications systems can be observed throughout our daily technological interactions. In densely populated urban environments, the concurrent presence of numerous wireless devices create challenging interference scenarios. Consider a typical office building where overlapping Wi-Fi networks and Bluetooth devices operate in the same unlicensed frequency bands, leading to packet collisions and degraded network performance. These effects become particularly acute in environments saturated with wireless activity, such as convention centers or apartment complexes. In cellular networks, the fundamental challenge of interference between neighboring base stations that reuse the same frequency allocations directly impacts user experience through reduced data rates and dropped connections. As networks evolve toward 5G architectures, these challenges intensify due to ultra-dense network deployments and dynamic spectrum sharing mechanisms that push traditional interference mitigation methods to their limits.

The issue of interference in wireless communications goes far beyond simple frequency overlaps—it reflects the intricate complexity of today’s wireless ecosystems. Consider the growing field of vehicle-to-everything (V2X) communications, where vehicles need to maintain stable connections while contending with constantly changing interference conditions. For instance, a car moving through an urban environment faces interference not only from nearby vehicles using 5G sidelink communication but also from cellular infrastructure and other wireless systems. This dynamic scenario underscores a major challenge: the coexistence of diverse wireless systems operating in close quarters, each with unique power levels, hard-

ware limitations, and communication protocols. While many strategies have traditionally aimed to prevent interference through proactive resource allocation and dynamic spectrum management, the rapid increase in wireless device density and the growing complexity of interference environments make avoidance alone insufficient. Consequently, even with optimal planning, interference is often inevitable, which necessitates advanced receiver techniques capable of dynamically identifying and suppressing interfering signals in real time, particularly in latency-sensitive or safety-critical scenarios. For example, high-powered cellular base stations and low-power IoT devices can interact in ways that create highly complex interference patterns, which require sophisticated strategies for effective suppression. Similarly, in cognitive radio systems, where primary and secondary users share the same spectrum, mitigating interference is crucial to ensuring reliable performance for all users. However, traditional interference suppression techniques usually have significant limitations. These methods often rely on fixed assumptions about channel conditions or signal structures, making them less effective in practice. The ever-changing nature of modern wireless environments, coupled with the wide range of devices involved, demands more flexible and robust solutions. Deep learning-based approaches, with their ability to analyze complex patterns directly from data, are uniquely equipped to tackle this challenge.

## 1.2 Motivation

The exponential growth of wireless communication services has led to immense crowding of the electromagnetic spectrum, making interference a critical challenge for modern wireless systems. This research is motivated by several key factors:

1. Limitations of Traditional Approaches: Conventional interference suppression methods often struggle with dynamic and complex interference patterns characteristic of

modern wireless environments. There is a pressing need for more adaptable and robust suppression techniques.

2. **Spectrum Scarcity and Efficient Utilization:** Effective interference suppression can enable better spectrum sharing and improve overall spectral efficiency, crucial for supporting the growing demand for wireless services in a finite spectrum.
3. **Emergence of New Wireless Technologies:** The advent of 5G, upcoming 6G networks, and the co-channel operation of terrestrial and non-terrestrial networks (along with IoT devices) introduces new challenges in interference management, requiring adaptive suppression techniques to ensure successful deployment and operation. For example, recent research [1] on co-channel spectrum sharing between terrestrial and satellite networks demonstrates that cognitive, database-driven control of satellite beams can manage interference effectively without degrading terrestrial network performance
4. **Potential of Deep Learning in Signal Processing:** Recent advancements in deep learning, particularly in CNNs and autoencoders, show promise in overcoming limitations of traditional methods and providing more effective, adaptive solutions for interference suppression.
5. **Need for End-to-End Solutions:** There is growing interest in end-to-end solutions that can seamlessly integrate multiple aspects of interference handling, from detection to suppression, in a single framework.
6. **Increasing Complexity of Interference Scenarios:** As wireless networks become denser and more heterogeneous, developing techniques that can handle diverse and complex interference types is crucial for building robust systems.
7. **Demand for Cognitive and Adaptive Systems:** Advanced interference suppression techniques that can learn and adapt to changing conditions align with the vision of more intelligent and flexible wireless systems.

This research aims to explore the potential of autoencoders that use CNN layers to transform the the input signal into a reconstructed form and achieve the goal of suppression of unwanted interfering signals in wireless applications. By developing and evaluating a CNN-AE approach, the work seeks to address the limitations of traditional filtering methods by leveraging the generalization capabilities of deep learning, and contribute to the development of more robust, adaptive, and efficient wireless communication systems that can thrive in increasingly complex interference environments.

## 1.3 Related Work

### 1.3.1 Traditional Interference Suppression Techniques

Interference management in wireless networks is tackled using a multidisciplinary approach involving signal processing, channel coding, information theory, adaptive filtering, and thresholding. Traditional suppression techniques encompass a broad spectrum of methods, as surveyed extensively in prior literature. For instance, Andrews *et al.* [2] provide a foundational overview of interference cancellation in cellular systems, emphasizing uplink and downlink challenges, while Tsang *et al.* [3] survey adaptive signal processing techniques for mitigating jamming, co-channel, and multi-access interference. Recent advances are captured by Aygur *et al.* [4], whose 2025 survey highlights narrowband interference (NBI) mitigation strategies, including emerging machine learning methods, underscoring the evolving complexity of spectrum-sharing systems.

Blind Source Separation (BSS) methods, as their name implies, operate without explicit prior knowledge of interfering signals, though practical implementations often require assumptions such as the number of interferers [3]. In contrast, physical layer manipulation techniques

explicitly leverage structural properties of signals. Neemat *et al.* [5] use the Short-Term Fourier Transform (STFT) to identify beat frequencies in FMCW radar signals disrupted by *other radar interference*, subsequently estimating Linear Prediction (LP) coefficients to interpolate corrupted frames. Carrick *et al.* have constructed a special filter that modifies the FRESH filter (traditionally used for rejection of interference signals) by adding an additional time varying property to the sequence— a set of filter weights periodically applied according to the received signal characteristics in OFDM systems. Shimura *et al.* [6] employ envelope detection to isolate interference in automotive radar. These methods rely on domain-specific signal characteristics, such as modulation patterns or spectral sparsity, to isolate interference. However, scalability remains a critical challenge. As highlighted in surveys on wireless sensor networks (WSNs) [7, 8] and body area networks [9], traditional techniques struggle in dense, heterogeneous environments with diverse devices and dynamic interference profiles. For example, WSNs in the 2.4 GHz band face severe packet losses due to uncoordinated Wi-Fi and microwave devices [7], demanding adaptive channel selection and transmission scheduling. While spatial techniques like GPS-oriented array processing [10] mitigate directional interference, they often require specialized hardware. These limitations underscore the need for flexible, learning-driven approaches suited to modern wireless ecosystems.

### 1.3.2 Machine Learning-based Interference Suppression Techniques

#### Classical Machine Learning Algorithms

Machine learning algorithms have demonstrated significant success and widespread adoption in addressing the challenges of interference characterization, identification, mitigation, and suppression. This stems from the powerful representational capacity of ML methods, enabling them to adapt their parameters based on observed data and learn complex, non-

linear relationships. While deep learning models offer exceptional performance, classical, non-deep learning algorithms also play a crucial role, particularly when interpretability and computational efficiency are paramount. These “lighter” models often provide a balance between accuracy and practicality, making them valuable tools for real-world applications. Several studies have explored the efficacy of these classical machine learning techniques in various interference-related tasks. For example, Swindell *et al.* [11] leverages the established technique of regression analysis to predict interference levels on simulated data, evaluating performance using the mean absolute % error (MAPE) metric. Similarly, [12] and [13] are examples of works that use Support Vector Machines (SVM) to classify various different types of interfering signals in automotive radar and wireless sensor networks respectively by using a feature extraction step that extracts meaningful representations from signal data. The work of [14] tackles SAR radio interference by adapting block sparse Bayesian learning to handle non-sparse scenarios, introducing two variants: one leveraging intra-block correlations for precise RFI/echo separation, and another optimizing computational efficiency through adaptive dictionary tuning—demonstrating robust suppression without hardware upgrades. This highlights how classical algorithms remain relevant due to their analytical tractability, deterministic behavior, and lightweight computational profiles—attributes critical for resource-constrained edge devices or latency-sensitive systems requiring predictable runtime, such as automotive radar. While less flexible than modern ML, these traits address fundamental hardware and timing constraints in real-world deployments. In the next subsection, however, we shall look into deep learning methods for interference suppression that have seen a lot more interest and successful applications in recent literature.



## Deep Learning for Interference Suppression

Deep learning has reshaped the landscape of machine learning and signal processing over the past decade. Since the breakthrough demonstration of deep convolutional networks by Krizhevsky *et al.* [15] on large-scale image classification tasks, these methods have consistently outperformed traditional approaches in a variety of domains, including time-series analysis and regression problems that are central to communication systems. The comprehensive review by LeCun, Bengio, and Hinton [16] further solidified deep learning's transformative potential by showcasing its ability to learn complex, hierarchical representations directly from raw data—a capability that aligns naturally with the signal decomposition challenges inherent in wireless interference suppression. Innovative architectures such as Long Short-Term Memory (LSTM) networks [17], designed to model temporal dependencies, and deep residual networks [18], which enable training of extremely deep architectures, have pushed the boundaries of what is achievable, setting new benchmarks for performance and adaptability in high-dimensional, non-stationary environments.

Building on these foundations, deep learning approaches have rapidly emerged as a powerful alternative for interference suppression, particularly in complex and dynamic wireless environments. Unlike classical methods that often depend on handcrafted features and rigid statistical assumptions about interference properties, deep learning models automatically learn hierarchical representations directly from raw signal data. This capability allows them to capture intricate non-linear relationships inherent in interference patterns, including spatially correlated interference sources and temporally varying channel conditions, offering significant performance improvements over traditional techniques. Furthermore, the scalability and adaptability of deep neural networks make them especially well-suited for handling the ever-growing volume and complexity of data in modern wireless communication systems, where interference characteristics evolve dynamically across heterogeneous devices and pro-

ocols.

Building on these foundations, early applications of deep learning in wireless communications focused on modulation recognition and signal classification, as demonstrated in works by O’Shea *et al.* like [19] and in the survey of deep architectures for modulation recognition by West and O’Shea [20]. These pioneering studies, along with the over-the-air deep learning based radio signal classification presented by O’Shea *et al.* [21], established that deep neural networks could learn discriminative features directly from raw IQ samples, thus setting the stage for tackling more complex interference issues. Subsequent research extended these ideas to interference identification tasks, as exemplified by the dynamic CNN approach for wireless interference identification [22] that uses a mixture of experts (MoE) consisting of individual CNNs and by FPGA-based implementations that harness LTE Cell-Specific Reference Signals for interference classification in wireless systems [23]. Furthermore, the classification of interference signals in automotive radar systems using convolutional neural networks by Kim *et al.* [24] reinforced the viability of these techniques in real-time environments.

As research progressed from classification to active suppression, deep learning methods began to address the challenges of interference mitigation directly. A comprehensive survey of the various ML based methods used for interference suppression is presented by Oyedare *et al.*’s [25] survey paper that does a deep dive into the comparative performance of different architectures for different types of interference scenarios. For example, the work on RIS-aided high-speed railway networks [26] employs a deep deterministic policy gradient scheme to optimize reconfigurable intelligent surfaces for interference suppression in dynamic railway environments. In parallel, studies such as the deep learning based communication signal interference suppression method by Chen *et al.* [27] have leveraged architectures like U-Net to enhance signal quality in the presence of interference. Other notable contributions include DSIC, a DL self-interference cancellation framework based on deep neural networks

that achieves 17dB of interference cancellation in-band full duplex wireless (IBFD) systems [28], and DeepSIC, which integrates deep learning into iterative soft interference cancellation for multiuser MIMO detection [29]. Complementary to these efforts, CNNs operating in the complex domain are leveraged to enhance radar range-Doppler map denoising and mitigate mutual interference while preserving essential phase information [30], as well as machine learning-based resource allocation algorithms for interference mitigation in D2D-enabled cellular networks [31] and heterogeneous networks [32], further underscore the breadth of deep learning applications in interference management. Additional approaches using recurrent neural networks for interference management in GPS navigation [33, 34] and for wireless networks [35], together with investigations into learning optimization strategies [36] and comparative analyses of supervised versus unsupervised training approaches [37], have enriched our understanding of the practical trade-offs and design considerations in deploying deep learning for interference suppression. Studies addressing wireless jamming attacks [38] and novel physical layer design paradigms [39, 40] have further illustrated the versatile applicability of these methods.

### **Autoencoder-based Approaches for Signal Processing & Interference Removal**

Autoencoders are a highly intuitive approach to learn and compress complex representations of high dimensional data into a compact form, enabling data compression and visualization. Autoencoders are well known for being successful in reducing out of distribution artifacts from input signals. Autoencoders have achieved great results with image reconstruction problems especially in the biomedical domain as show by [41] and [42] that use autoencoders built with convolutional layers to denoise and remove artifacts from biomedical images.

Almazrouei *et al.* [43] use a convolutional denoising autoencoder to demonstrate its efficacy in radio signal denoising in an ML pipeline that involves a short time fourier transform

(STFT) of the time domain input signals. Lin *et al.* [44] addressed unstructured cross-technology interference (e.g., from microwaves) in WiFi networks by pairing a denoising autoencoder with GANs, which generate realistic synthetic training data to mimic real-world interference. This approach bypasses labor-intensive signal collection and prior knowledge of interference patterns, passively suppressing noise to boost decoding success without costly hardware coordination.

Chen *et al.* [45] developed a DNN autoencoder paired with an interference detection filter to mitigate automotive radar interference, leveraging gated convolutions to isolate clean signal patterns from simulated training data while avoiding costly manual labeling. Their framework reconstructs contaminated time-domain signals without phase distortion, demonstrating robustness in real-world dense traffic scenarios with improved SINR.

Brown *et al.* [46] developed a CNN-layer-only denoising autoencoder for aircraft radar altimeter interference mitigation, uniquely addressing the case where radar is the desired signal and communication signals act as interference. Their model removes interference without prior detection, enhancing range estimation accuracy, reducing false altitude reports, and improving SINR severe interference conditions. Autoencoder-based techniques have also been useful in optimizing transmitter and receiver systems to mitigate interference. For example,

Erpek *et al.* [47] proposed a deep learning-driven autoencoder framework for MIMO interference channels, jointly optimizing transmitter-receiver pairs under Rayleigh fading to suppress interference, outperforming conventional TDMA and SVD-based systems at high SINR despite requiring transmitter-side CSI and scalable training complexity with antenna configurations.

One of the most popular autoencoder architectures in signal processing literature is the

UNet, originally devised by Ronneberger et al. [48]. The UNet employs skip connections—residual connections that reintroduce feature maps from earlier layers into later layers—thereby enhancing information flow and facilitating improved gradient propagation. By combining convolutional layers with these skip connections, the UNet effectively retains multi-scale features, which is critical for constructing robust hierarchical representations. Initially designed for biomedical image segmentation [49, 50] and other specialized computer vision tasks [51], the UNet architecture has since found broad applicability in the general field of signal processing, including wireless communication applications. For example, Liang *et al.* [52] utilized a UNet architecture, augmented with a novel spatiotemporal module, to predict wireless traffic movement in an Internet of Vehicles scenario by capturing cross-domain features. Similarly, Kang *et al.* [53] integrated a UNet-based transformer model to estimate and predict channel characteristics in a massive MIMO use case.

The versatility and proven performance of the UNet has also inspired the architecture employed in the work presented in this thesis. In subsequent sections, we will detail how the principles of the UNet are adapted to address the interference suppression task.

## **Cyclostationary Signal Processing & Interference Suppression**

Cyclostationary signal processing leverages the inherent periodicities found in many man-made signals to enhance detection, classification, and estimation tasks. Unlike stationary signals, which have constant statistical properties, cyclostationary signals exhibit periodic variations in their mean or autocorrelation functions. This cyclic behavior—often introduced by modulation schemes or periodic interference—provides additional structure that can be exploited to improve signal analysis, even under low signal-to-noise conditions.

Over the past few decades, seminal works have established the theoretical foundations and

practical applications of cyclostationary methods. Gardner’s pioneering studies [54, 55] highlighted how these techniques could be used for tasks such as modulation recognition and parameter estimation, paving the way for modern applications in communications and radar systems. Kim *et al.* [56] use the cyclostationary properties of communication signals to extract features and used a hidden markov model classifier to detect and classify captured signals at low SNRs.

Subsequent research has built on these ideas, demonstrating that by extracting cyclic features, one can effectively differentiate between useful signal components and unwanted noise or interference [57, 58, 59]. This capability has proven especially valuable in scenarios like cognitive radio [60] and robust communications, where traditional methods might falter in complex or hostile environments. We shall explore the application of cyclostationary properties to enhance the performance of data driven approaches in a future chapter.

## 1.4 Contribution

This thesis makes several significant contributions to the field of interference suppression in wireless communications:

- **CNN-AE Architecture:** We investigate a convolutional neural network-based autoencoder (CNN-AE) specifically designed for end-to-end interference suppression in wireless communications by training it on a dataset with several interference types.
- **Comprehensive Interference Analysis:** We evaluate the CNN-AE model across various interference scenarios, including single frequency tone interference (in-bandwidth and out-of-bandwidth), partial band and wideband communications signal interference, providing insights into its performance and adaptability.
- **Latent Space Investigation:** We explore the latent space representations of signals af-

ected by different types of interference, offering new perspectives on how deep learning models interpret and separate interference.

- The CNN-AE framework under investigation is thoroughly benchmarked against conventional interference suppression methods, showcasing superior performance across metrics such as bit error rate (BER), error vector magnitude (EVM), and Signal to Noise-Interference Ratio (SINR) to validate its robustness and practical applicability in real-world interference environments.
- Generalization Study: We investigate the model’s ability to generalize across different signal-to-noise ratios (SNRs) and interference types, providing valuable insights into its robustness and practical applicability.
- Cyclostationarity-Enhanced Training: We introduce a novel approach that explicitly tailors the CNN-AE architecture to exploit the cyclostationary properties of the desired signal, integrating these features into the training process. By aligning the network’s hierarchical feature extraction with the periodic statistical dependencies inherent to cyclostationary signals, the model prioritizes learning patterns tied to the desired signal’s structure by clever preprocessing.

## 1.5 Organization of Thesis

The remainder of this thesis is organized as follows:

- Chapter 2: System Model - Describes the signal of interest (SoI) and the types of Interference used in this study. This chapter provides a comprehensive framework for understanding the wireless communication scenario that is under investigation.
- Chapter 3: Methodology - Details the CNN-AE architecture, including generation of wireless signal training data, data preprocessing, model parameters, and evaluation

metrics. This chapter also explains the latent space visualization and provides an overview of the experimental setup.

- Chapter 4: Results for CNN-AE Interference Suppression - Presents and analyzes the performance of the CNN-AE model trained on multiple SNRs and interference types. It includes comparisons with traditional techniques and discussions on the model's generalization capabilities.
- Chapter 5: Deep Learning Interference Suppression for Cyclostationary Signals - Cyclostationarity properties are accounted for in the CNN-AE architecture to exploit the periodic correlations unique to cyclostationary signals. The tailored design embeds prior knowledge of the desired signal's cyclostationary features (such as symbol-rate periodicity or phase-dependent statistics) into the data, enabling targeted suppression of non-stationary interference.
- Chapter 6: Conclusion - Summarizes the key findings of the research, discusses the implications of the results, and proposes directions for future work involving the use of DL based architectures in the mix for interference suppression and other applications of critical importance in next generation wireless.



# Chapter 2

## System Model

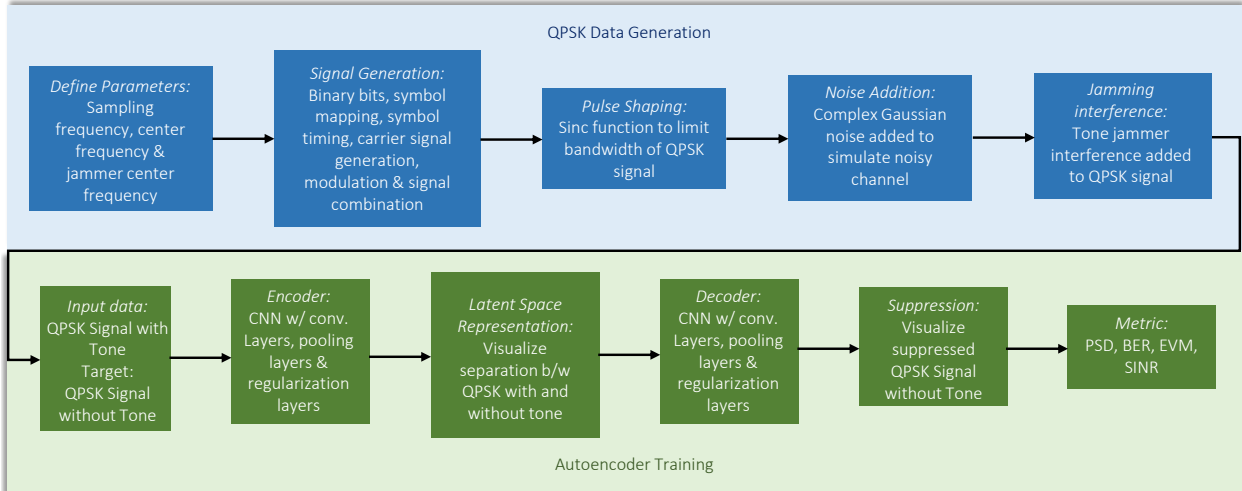


Figure 2.1: Outline of steps in System Model

This chapter presents a detailed description of the wireless communication system model used in this study. The model encompasses the signal of interest (SOI), various interference sources, and the channel characteristics. We consider a complex baseband representation of all signals, focusing on Quadrature Phase Shift Keying (QPSK) modulation for the SOI and four distinct interference scenarios: single frequency tone interference (in-bandwidth and out-of-bandwidth), partial band and wideband interference.

## 2.1 Description

The communication framework operates in the complex baseband domain, bypassing the need for high-frequency carrier modulation during signal analysis. This approach simplifies computational demands by representing waveforms through orthogonal in-phase (I) and quadrature (Q) channels, which separately encode amplitude and phase information. To construct the desired signal, binary data is transformed into Quadrature Phase Shift Keying (QPSK) symbols, where each symbol encodes two bits using four distinct phase angles (e.g.,  $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ). This modulation scheme enhances spectral efficiency by transmitting twice as much data per symbol compared to traditional binary phase shift keying (BPSK), making it ideal for bandwidth-constrained environments.

Nonlinearities in transmitter components, such as RF amplifiers operating near saturation, generate unintended harmonics and mixing effects. These distortions become problematic when coupled with external interference sources—like narrowband jammers or leakage from adjacent frequency bands—leading to intermodulation artifacts. Such artifacts infiltrate the QPSK signal’s bandwidth, degrading symbol clarity. At the receiver, this corruption translates to misaligned timing recovery, scrambled constellation diagrams, and increased error rates, necessitating sophisticated suppression algorithms to mitigate interference-induced performance losses.

## 2.2 Composite Signal Model with Noise

The discrete-time received signal is formulated as a superposition of three distinct components:

$$x[n] = y[n] + I[n] + \eta[n] \quad (2.1)$$

where:

- $y[n]$  corresponds to the QPSK signal of interest.
- $I[n]$  represents interference, which may include narrowband tones or wideband noise.
- $\eta[n]$  models additive white Gaussian noise (AWGN) with variance  $\sigma^2$ , emulating stochastic channel impairments.

## 2.3 Signal of Interest (SOI) Model

In this study, the SOI under investigation shall be a QPSK signal that is modulated from randomly generated bits. QPSK is a digital modulation scheme that encodes two bits per symbol, offering a balance between spectral efficiency and robustness against noise and interference. The discrete-time received signal is formulated as a superposition of three distinct components:

### 2.3.1 QPSK Signal Generation

The QPSK signal generation process involves the following steps:

1. **Bit Generation:** A sequence of random binary bits  $b_k$  is generated, where  $b_k \in \{0, 1\}$ .
2. **Symbol Mapping:** Pairs of bits are mapped to complex symbols according to the QPSK constellation. The mapping is defined as:  $(0, 0) \rightarrow (+1, +1j)$ ,  $(0, 1) \rightarrow (-1, +1j)$ ,  $(1, 0) \rightarrow (+1, -1j)$ ,  $(1, 1) \rightarrow (-1, -1j)$
3. **Symbol Formation:** The mapped symbols form the sequence  $\{s_k\}$ , where  $s_k \in \{(+1, +1j), (-1, +1j), (+1, -1j), (-1, -1j)\}$ .
4. **Pulse Shaping:** To limit the bandwidth of the signal, each symbol is pulse-shaped

using a sinc function. The pulse-shaped signal is given by:

$$y(t) = \sum_k s_k \cdot \text{sinc}((t - kT)/T) \quad (2.2)$$

where  $T$  is the symbol duration and  $\text{sinc}(x) = \sin(\pi x)/(\pi x)$ .

5. **Upsampling:** The signal is sampled at  $4\times$  the symbol rate ( $R_s$ ) to allow for accurate representation of the continuous-time signal in the discrete-time domain.
6. **Carrier Modulation:** The pulse-shaped signal is modulated onto a carrier frequency  $F_c$ . In this study, we consider a desired signal centered at baseband, so  $F_c$  is set to zero.

The resulting complex baseband QPSK signal can be mathematically represented as:

$$y[n] = \sqrt{P_S} \sum_k s_k \cdot \text{sinc}(nT_s - kT) e^{j2\pi F_c (iT_s)} \quad (2.3)$$

where:

- $y[n]$  represents the QPSK continuous time signal sampled into discrete time form
- $P_S$  is the overall value for power of our SOI
- $s_k$  are the symbols generated in the process of QPSK modulation
- $T_s$  is the time interval of sampling
- $T$  is the symbol duration
- $F_c$  is the carrier frequency (set to zero for baseband simulation)
- $n$  represents the discrete-time index

## 2.4 Interference Models

This study considers four types of interference to evaluate the robustness of the proposed interference suppression technique:

### 2.4.1 Tone Interference Model

Tone interference is modeled as a continuous-wave (CW) sinusoidal signal injected into the QPSK signal. The interference is defined as a complex exponential:

$$I[n] = \sqrt{P_{X_{\text{in}}}} \cdot e^{j2\pi F_{\text{in}}nT_s} \quad (2.4)$$

where:

- $P_{X_{\text{in}}}$ : Power of the interfering tone.
- $F_{\text{tone}}$ : Frequency of the tone, offset from the QPSK carrier frequency  $F_c$  by  $\Delta F = |F_c - F_{\text{tone}}|$ .
- $T_s$ : Sampling interval.

The interference's impact depends on its spectral alignment relative to the QPSK signal's bandwidth  $B = 1/2T$ , where  $T$  is the symbol duration

#### **In-Band Interference ( $\Delta F < B$ ):**

When the tone overlaps with the QPSK bandwidth, spectral interference occurs. The continuous sinusoidal nature of the tone distorts both phase and amplitude of the QPSK signal. For example, a tone at  $F_c + 0.4B$  introduces phase shifts that misalign constellation points, elevating bit error rates (BER) by 30 – 40% and complicating symbol synchronization. This

distortion is evident in the spectral overlap shown in Figure 2.2, where the tone’s PSD peak obscures the QPSK signal’s spectral profile.

**Out-of-Band Interference ( $\Delta F > B$ ):**

Out-of-band tones reside outside the QPSK bandwidth, often in guard bands or adjacent channels. While these tones avoid direct spectral overlap, imperfect filter roll-off allows residual energy near the band edges, degrading SINR by 3–5 dB. For instance, a tone at  $F_c + 1.2B$  stresses receiver analog filters, reducing effective bandwidth utilization. As illustrated in Figure 2.3, the tone’s PSD remains distinct from the QPSK signal but occupies critical guard regions.

**Receiver Sensitivity and Mitigation:**

Strong tone interference reduces receiver sensitivity, impairing QPSK signal detection and demodulation. Advanced techniques, such as adaptive notch filters or deep learning-based suppression, are often required to restore performance. While conventional methods suffice for simple tone suppression, their inclusion in this study provides a baseline to evaluate deep learning models’ ability to isolate interference features in latent space, enhancing visualization and analytical clarity.




Figure 2.2: In-Band Tone Interference Comparison

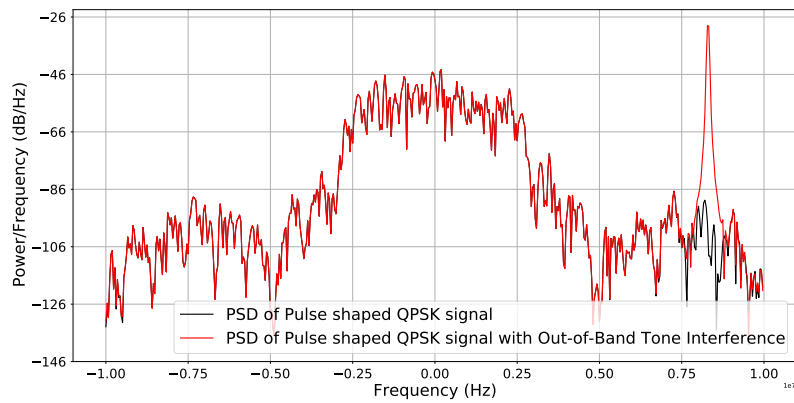
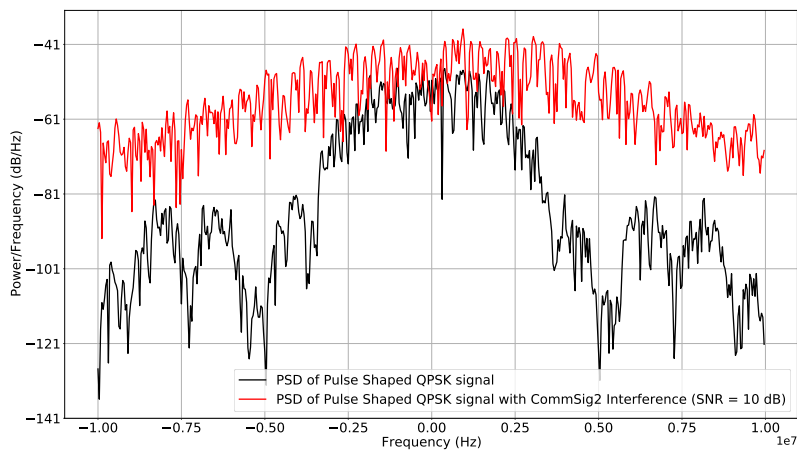


Figure 2.3: Out-of-Band Tone Interference Comparison

Figure 2.4: *CommSignal2* Interference Comparison

### 2.4.2 Wideband Interference (CommSignal2)

The CommSignal2 dataset, part of the ICASSP RF Challenge [61] hosted by MIT, represents a formidable real-world interference scenario distinct from conventional narrowband disruptions. In contrast to both the straightforward scenarios of single frequency tone interference, the span of the CommSignal2 interference frame covers the entire bandwidth of the baseband QPSK signal, creating a complex spectral environment where interference coexists with the desired signal across all frequencies. This dataset, comprising 150 frames of 43,560 samples each, captures over-the-air (OTA) transmissions from commercial wireless devices, embedding hardware-induced artifacts such as amplifier nonlinearities and phase noise.

#### Spectral and Practical Complexity

As seen in Figure 2.4, CommSignal2’s energy blankets the QPSK signal’s frequency range, overlapping both in-band and adjacent regions. This broad spectral occupancy reduces the signal-to-interference-plus-noise ratio (SINR) by 15–20 dB, inducing errors in symbol recovery and bit decision boundaries. For example, demodulation experiments reveal a 25% increase in bit error rates (BER) compared to tone interference scenarios, which underlines its severity. A randomized 512 sample length slice of the *CommSignal2* interference frame is shown in Figure 2.5.

#### Undisclosed Parameters and Data-Driven Mitigation

The signal generation parameters for CommSignal2 are intentionally withheld in [61], preventing reliance on model-based suppression techniques. This design choice emphasizes the necessity of adaptive, data-driven approaches like the CNN-AE architecture. The dataset’s non-stationary behavior—mimicking dynamic interference sources such as frequency-hopping



systems or pulsed radar—further challenges conventional filtering methods, which struggle with time-varying spectral occupancy.




Figure 2.5: Extracted Frame of *CommSignal2* Interference

### 2.4.3 QPSK Signal Interference

The QPSK interference model represents a co-channel interference scenario where an unwanted QPSK signal overlaps with the signal of interest. This interference follows the same mathematical structure as the desired QPSK signal presented in Equations (2.2) and (2.3) of Section 2.3.1, but with different power levels and timing offsets. The interfering QPSK signal can be expressed as:

$$I[n] = \sqrt{P_I} \sum_k s_k^I \cdot \text{sinc}(nT_s - kT - \tau) e^{j2\pi F_I(nT_s)} \quad (2.5)$$

where:

- $P_I$  is the interference power
- $s_k^I$  represents the interfering QPSK symbols
- $\tau$  denotes the timing offset between desired and interfering signals
- $F_I$  is the frequency offset of the interfering signal

This interference scenario can prove challenging for conventional suppression techniques due to its structural similarity to the desired signal. The analysis and suppression of QPSK interference effects are examined in Chapter 5, where specific attention is given to scenarios involving varying symbol rates and timing offsets between the desired and interfering signals.

## 2.5 Summary

This system model provides a comprehensive framework for studying interference suppression in QPSK-based wireless communication systems. It incorporates realistic interference scenarios, including challenging wideband interference, and allows for systematic evaluation of suppression techniques under various SINR conditions. The next chapter will detail the methodology used to develop and evaluate the CNN-based autoencoder for interference suppression based on this system model.

# Chapter 3

## Methodology

### 3.1 Dataset Generation

The synthetic dataset is engineered to train a convolutional autoencoder for suppressing interference in Quadrature Phase Shift Keying (QPSK) communication systems. A structured framework governs the generation process, ensuring systematic variation of signal and interference parameters to mimic diverse real-world scenarios. A list of key components of the dataset is shown in Table 3.1.

Table 3.1: List of QPSK signal generation parameters

Parameter	Selected value
Count of generated bits	128 bits each for I & Q channels
Number of samples per symbol	4
Sampling Frequency ( $F_s$ )	20 MHz
Range of tone frequencies (In-band)	Training: 300 KHz to 2 MHz
	Validation: 1 MHz — 1.3 MHz
	Test: 300 KHz — 700 KHz
Range of tone frequencies (Out-of-band)	Training: 2.5 MHz — 9.5 MHz
	Validation: 5 MHz — 7 MHz
	Test: 2.5 MHz — 4.5 MHz

### 3.1.1 Baseband QPSK Signal and Noise/Interference configuration

Each QPSK symbol sequence comprises 128 in-phase (I) and quadrature (Q) bits. These symbols are upsampled by a factor of 4 to accommodate pulse shaping and temporal dynamics with a sampling frequency of 20 MHz. The modulated signals are confined to baseband simulations, with the carrier frequency  $F_c$  set to 0 MHz for computational simplicity.

### 3.1.2 Interference Synthesis

Two distinct interference types are introduced to the clean QPSK signals:

- *Single-Frequency Tone*: A sinusoidal tone is injected at frequencies spanning in-band and out-of-band regions. As detailed in Table 1, the in-band tone spans 300 KHz–2 MHz (training), 1 MHz–1.3 MHz (validation), and 300 KHz–700 KHz (test). Out-of-band interference occupies higher frequencies: 2.5 MHz–9.5 MHz (training), 5 MHz–7 MHz (validation), and 2.5 MHz–4.5 MHz (test).
- *CommSignal2 Interference*: Random slices from the *CommSignal2* dataset—a real-world over-the-air (OTA) interference frame is superimposed on the QPSK signals.

### 3.1.3 Noise and Power Calibration

Additive white Gaussian noise (AWGN) with a signal-to-noise ratio (SNR) of 10 dB is introduced across the dataset to emulate channel impairments. The interference power is also separately scaled according to the power of the clean signal to achieve a signal-to-interference ratio (SIR) of -10 dB. We evaluate the Power Spectral Density (PSD) plot for the reconstructed signal and provide a qualitative analysis of the reconstruction. In later

sections, we shall fix either of these quantities at these values ( $\text{SNR} = 10$  dB or  $\text{SIR} = -10$  dB) and vary the other quantity within a certain range to inspect the performance of our approach and its robustness to varying noise and interference conditions.

### 3.1.4 Data Partitioning

The dataset is partitioned into training, validation, and test subsets, each adhering to unique frequency ranges for interference tones as shown in Table 3.1. This partitioning strategy ensures robust generalization by exposing the model to non-overlapping spectral regions during different phases of training and evaluation.

This comprehensive dataset generation methodology ensures the autoencoder is exposed to a wide spectrum of interference scenarios, enabling robust suppression capabilities in both controlled and variable channel conditions.

## 3.2 Data Preprocessing

### 3.2.1 Signal Conditioning and Interference Injection

The raw QPSK dataset undergoes a multi-stage preprocessing pipeline to simulate a communication system and prepare inputs for the CNN autoencoder. A description of key stages involved is given below.

#### Baseband Signal Preparation

Each QPSK symbol sequence is first subjected to *pulse-shaping* to mitigate inter-symbol interference (ISI) and emulate transmission through a bandlimited channel. Sinc pulse-

shaping is applied to shape the waveform, followed by upsampling to match the 20 MHz sampling rate. The shaped signal is then modulated onto a carrier frequency  $F_c = 0$  MHz (baseband), generating complex in-phase (I) and quadrature (Q) components.

### Interference Embedding

Interference signals are synthesized and superimposed onto the clean QPSK waveforms:

- *Tone Interference*: Single-frequency sinusoidal tones are injected at predefined in-band or out-of-band frequencies (see Table 1), with amplitudes calibrated to achieve a -10 dB signal-to-interference ratio (SIR).
- *CommSignal2 Interference*: Segments from the OTA *CommSignal2* dataset are randomly extracted, resampled, and added to emulate a non-stationary interference scenario.

The resultant composite signals encapsulate both the desired QPSK waveform and interference artifacts, replicating practical over-the-air transmission impairments.

## 3.2.2 Feature Engineering

### Time-Frequency Domain Trade-offs

Two preprocessing approaches were under consideration to determine optimal input representations for the autoencoder:

- *Time-Domain Signals*: Complex baseband I/Q samples are directly used as inputs, preserving phase continuity and temporal correlations.
- *Frequency-Domain Signals*: Fast Fourier Transform (FFT) is applied to convert I/Q samples into separate channels, emphasizing spectral interference patterns.

Empirical analysis revealed superior interference suppression performance with time-domain signal inputs. This advantage could stem from the autoencoder’s ability to retain phase information critical for QPSK demodulation, whereas FFT-based preprocessing introduces phase discontinuities during spectral leakage and windowing. Hence, we have used time domain QPSK signals for the clean and interfered datasets in our experiments. A potential area of future investigation in this space is possible with a more detailed comparison between FFT and Time Domain inputs and the implications of using either of them for specific tasks.

### **Tensor Structuring**

The preprocessed I/Q samples are formatted into  $512 \times 2$  shaped tensors, where the first dimension corresponds to time steps or samples, and the second represents I/Q channels i.e. the real and imaginary parts of the samples in the QPSK signal. This structure aligns with the 1D convolutional kernel dimensions (15 filters) in the autoencoder, enabling efficient spatial feature extraction. The splitting of samples into separate real and imaginary parts also prevents the introduction of complex valued signals into the model training process which may cause stability and convergence issues.

### **Training-Validation-Test Partitioning**

The dataset is divided into non-overlapping training (70%), validation (15%), and test (15%) subsets. In the case of tone interference, each subset employs distinct interference frequency ranges (Table 3.1) to prevent data leakage and ensure the model generalizes to unseen spectral interference profiles.

This systematic preprocessing framework ensures the autoencoder is trained on high-fidelity emulations of interference-corrupted QPSK signals while maintaining compatibility with the

network’s architectural constraints and learning dynamics.

### 3.3 CNN-AE Architecture

The proposed interference suppression system employs a convolutional autoencoder (CNN-AE) with a U-Net-like architecture [48], combining hierarchical feature extraction with skip connections to preserve spatial fidelity. The multi-resolution feature fusion between the encoder and decoder has shown strong performance in a variety of signal processing applications—such as image segmentation and denoising—by effectively capturing both global context and fine details. This proven success in complementary tasks supports its application in reconstructing clean QPSK signals from interference-corrupted inputs.

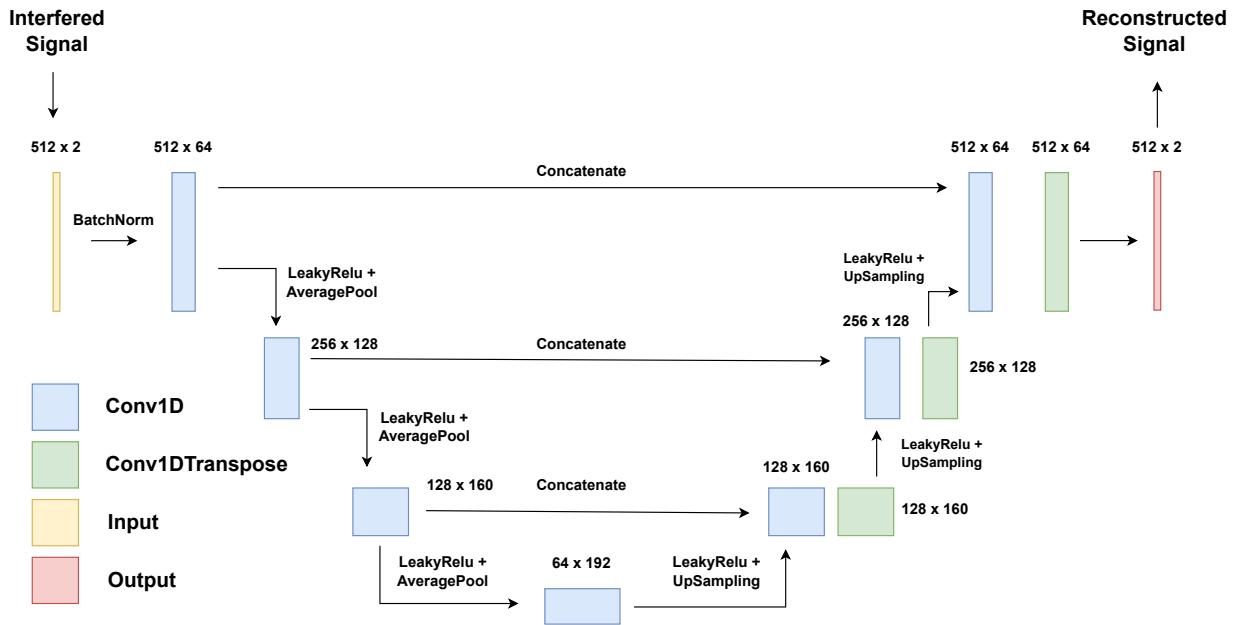


Figure 3.1: Overview of the System Model



Table 3.2: List of architectural specifications for CNN-AE

Parameter	Selected value
Batch size	128
Learning rate/gain (LR)	$1e - 4$ decayed by $\times 0.5$ till $1e - 8$
Stride	1
Size of kernel	15
Train—Test—Validation split percentage	70% — 15% — 15%
Input size	$512 \times 2$ (Real & Imaginary)
Number of parameters	Trainable: 3, 167, 462
	Non-trainable: 4
	Total: 3, 167, 466
Optimizer	Adam
Activation function	LeakyReLU

### 3.3.1 Encoder (Hierarchical Feature Extraction)

The encoder gradually transforms the input signal into a compact latent representation through four successive downsampling stages:

- **Layer Composition:** Each stage comprises a 1D convolutional layer (kernel size=15, stride=1), LeakyReLU activation, batch normalization, and average pooling.
- **Channel Expansion:** The number of filters increases geometrically across stages (64  $\rightarrow$  192), capturing multi-scale interference patterns.
- **Dimensionality Reduction:** Temporal resolution halves at each stage, transforming the  $512 \times 2$  input to an expanded  $64 \times 192$  latent space.

The LeakyReLU activation function, introduces controlled non-linearity while preventing gradient vanishing in deeper layers. It can be represented in this manner:  $f(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0.01x & \text{otherwise} \end{cases}$ .

### 3.3.2 Decoder: Signal Reconstruction via Transposed Convolutions

Mirroring the encoder, the decoder notably employs transposed convolutions (`Conv1DTranspose`) to upsample the latent representation. The key design considerations are:

- **Symmetrical Design:** The number of filters in each block decrease inversely to encoder stages ( $192 \rightarrow 64$ ), thus reversing the encoding process.
- **Upsampling Mechanism:** Each decoder block combines `Conv1DTranspose` layers that apply the reverse operation as that of the `Conv1D` layers in the encoder with `LeakyReLU` and bilinear interpolation to restore temporal resolution.

### 3.3.3 Skip Connections: Bridging Multi-Resolution Features

The architecture integrates skip connections between corresponding encoder and decoder blocks, a hallmark of U-Net architectures widely used in medical image segmentation. These connections address the “information bottleneck” problem in conventional autoencoders by:

- **Feature Concatenation:** The encoder block outputs are concatenated with decoder activations at matching spatial resolutions (Figure 3.1) and fed forward. This fusion combines high-level contextual features from the decoder with reintroduction of low-level spatial details from the encoder.
- **Gradient Propagation:** Direct pathways between early and late layers mitigate gradient attenuation during backpropagation, stabilizing training for deep networks.
- **Interference Localization:** By preserving high-frequency components from encoder layers, skip connections help distinguish transient interference artifacts from underlying QPSK modulation.

### 3.3.4 Architectural Hyperparameters

Critical design parameters (Table 3.2) used in the architecture were arrived at through ablation studies. Following were the considerations to evaluate while selecting them:

- **Kernel Size:** The 15 size convolutional kernels balance receptive field size with computational complexity to capture a sufficiently wide context window.
- **Depth vs. Width:** A total of 4 encoding/decoding stages provide sufficient depth for interference modeling without over-parameterization.
- **Latent Space:** The  $64 \times 192$  shaped bottleneck tensor captures essential modulation characteristics while discarding interference-related noise.

Thus, the U-Net inspired design is used in order to achieve high quality interference suppression by synergistically combining deep feature learning with multi-resolution signal preservation with the goal of conventional frequency-domain filtering approaches.

## 3.4 Training & Optimization

### 3.4.1 Loss Formulation and Gradient Dynamics

The autoencoder’s training objective centers on minimizing the discrepancy between reconstructed signals and pristine QPSK waveforms. A Mean Squared Error (MSE) loss function is adopted, quantified as:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{n=1}^N (y[n] - \hat{y}[n])^2 \quad (3.1)$$

where  $y[n]$  represents the clean uninterfered signal and  $\hat{y}[n]$  represents the autoencoder’s reconstructed version of the clean signal. This choice penalizes large reconstruction errors quadratically, prioritizing fidelity in high-amplitude signal regions critical for QPSK demod-

ulation.

### 3.4.2 Adaptive Learning and Regularization

The Adam optimizer integrates momentum and per-parameter learning rate adaptation to efficiently update the weights during training. To enhance the training process, **early stopping** is employed, which halts the training if the validation loss plateaus for 15 consecutive epochs. This mechanism utilizes a windowed approach to prevent false positives and ensure robust stopping criteria. Additionally, learning rate annealing is implemented to refine the convergence process. Specifically, the learning rate is set at  $10^{-4}$  and reduced by 50% after the validation loss plateau's for 5 epochs, allowing for more precise convergence near the minima and improving the overall model performance.

### 3.4.3 Training Protocol

Data is partitioned into 70% training, 15% validation, and 15% test sets, with non-overlapping interference frequencies (Table 3.1). For the *CommSignal2* dataset, the test split uses the `TestSetMixture1` test subset of the RFChallenge dataset [61]. A mini-batch of 128 size is used with random shuffling to decorrelate the signals, while the validation set remains fixed for consistent evaluation. The model is trained on the dataset for a 100 epochs which is determined heuristically as being sufficient for convergence.

### 3.4.4 Interference suppression mechanics

The model can be said to learn several key behaviors to effectively suppress interference during training. One such behavior is *latent sparsification*, where bottleneck layers are

trained to nullify activations associated with interference, thereby reducing their impact on the output. Another critical mechanism is *skip gating*, which leverages early encoder features to highlight transient interference, enabling the model to dynamically adapt to such disturbances. Additionally, the architecture ensures *phase preservation*, where time-domain inputs retain the phase information of the QPSK constellation, maintaining signal integrity throughout the processing pipeline. The training process typically converges within 100 epochs, and the final models are selected based on the minimum validation loss, ensuring robustness and optimal performance in interference-rich environments.

## 3.5 Evaluation Metrics for Performance Assessment

In the domain of wireless communication systems, particularly those employing sophisticated interference suppression techniques, the selection and implementation of appropriate evaluation metrics play a pivotal role in understanding system performance. This section presents a comprehensive framework for assessing the effectiveness of our CNN-AE architecture through three fundamental yet complementary metrics: the EVM, BER and SINR. These metrics have been carefully chosen to provide a multi-dimensional view of system performance, each capturing distinct aspects of signal quality, interference suppression efficiency, and overall communication reliability.

### 3.5.1 Error Vector Magnitude Analysis

Error Vector Magnitude (EVM) represents a cornerstone metric in modern wireless communications, offering a sophisticated approach to quantifying signal quality and modulation accuracy. Unlike simpler metrics that may focus solely on amplitude or phase errors,

EVM provides a comprehensive assessment by considering both magnitude and phase deviations between the ideal reference signal and the actual transmitted or received signal. This characteristic makes it particularly valuable in the context of our CNN-AE system, where preserving signal fidelity through the interference suppression process is paramount. The theoretical foundation of EVM lies in its ability to capture the cumulative effect of various impairments that may affect signal quality. These impairments can range from systematic distortions introduced by hardware components to random noise and interference effects. In our implementation, EVM calculation takes into account the vector difference between the ideal signal constellation points and their actual received positions in the I-Q plane, providing a normalized measure of signal quality that is independent of the modulation scheme's absolute amplitude. The mathematical representation of EVM, expressed in decibels (dB), is formulated as:

$$EVM(dB) = 20 \log_{10} \sqrt{\frac{\frac{1}{N} \sum_{n=1}^N |y[n] - \hat{y}[n]|^2}{\frac{1}{N} \sum_{n=1}^N |y[n]|^2}} \quad (3.2)$$

In this expression,  $y[i]$  represents the clean and uninterfered reference SOI and  $\hat{y}[i]$  stands for the reconstructed and interference suppressed version of the interfered counterpart. The denominator serves as a normalization factor, ensuring that the EVM measurement remains consistent regardless of the absolute signal power level. This normalization is crucial for making meaningful comparisons across different operating conditions and system configurations. The interpretation of EVM values requires careful consideration of several factors. Lower EVM values indicate better signal quality, with theoretical perfect reconstruction yielding an EVM of negative infinity dB. In QPSK modulation schemes like the one employed in our system, typical acceptable EVM thresholds range from -15 dB to -25 dB, depending on the specific application requirements and environmental conditions. The statistical distribution of EVM values across multiple symbols can reveal patterns of impairment that might not be immediately apparent from aggregate measurements alone.

### 3.5.2 Signal to Noise—Interference Ratio Evaluation

The Signal to Noise — Interference Ratio (SINR) is one of the most significant metrics in the assessment of interference suppression systems, providing direct insight into the effectiveness of interference mitigation techniques. In the context of our CNN-AE architecture, SINR evaluation takes on particular significance as it quantifies the system’s ability to preserve desired signal components while suppressing both structured interference and random noise. SINR analysis in our implementation considers three primary components: power of clean QPSK signal  $y[n]$ , the residual interference signal  $I[n]$ , and the ambient noise power characterized by  $\sigma^2$ . The mathematical framework for SINR computation is expressed as:

$$SINR(dB) = 10 \log_{10} \frac{\frac{1}{N} \sum_{i=1}^N |y[n]|^2}{\frac{1}{N} \sum_{n=1}^N |\hat{y}[n] - y[n]|^2} \quad (3.3)$$

This formulation provides a robust measure of signal quality that accounts for both interference and noise effects. The numerator represents the power of the desired signal component, while the denominator captures the combined power of residual interference and noise. SINR measurements before and after signal processing provide direct quantification of interference suppression effectiveness, with typical improvements ranging from 10 dB to 30 dB, depending on the nature of the interference and the sophistication of the suppression technique. The significance of SINR extends beyond simple numerical comparison. By examining SINR performance across different interference scenarios (in-band versus out-of-band) and power levels, we can characterize the robustness and adaptability of our interference suppression approach. This understanding proves crucial for predicting system performance in real-world deployments where interference characteristics may vary significantly. Furthermore, SINR measurements guide system optimization efforts by highlighting conditions under which interference suppression is most and least effective.

### 3.5.3 Bit Error Rate Performance Analysis

Bit Error Rate (BER) represents a fundamental metric in digital communication systems, particularly crucial for QPSK modulation schemes. Our BER analysis methodology incorporates several sophisticated processing steps to ensure accurate error quantification. The process begins with signal demodulation, achieved through complex conjugate multiplication with the carrier signal, followed by careful down-sampling to recover the original symbol rate. This process requires precise timing synchronization and phase alignment to minimize errors introduced during the demodulation process itself. The BER evaluation process considers both the In-phase (I) and Quadrature (Q) components of the QPSK signal independently. Each received symbol undergoes interpretation within the context of the QPSK constellation diagram, with decision boundaries properly adjusted to account for any residual interference or distortion effects. For each component, the precise count of incorrect bits are determined and normalized by the overall combined count of transmitted bits. The combined BER is then computed as the arithmetic mean of the resultant error rates in both I and Q channels, providing a balanced assessment of system performance.

Modern communication systems typically achieve BER values in the range of  $10^{-6}$  to  $10^{-9}$  through the use of error correction coding. In our analysis, however, error correction coding is not considered, so the results presented in the following chapters reflect the performance of the raw system and could be further improved with its incorporation. These stringent requirements reflect the critical nature of data integrity in many applications, from basic data communication to safety-critical systems. In our CNN-AE system, BER analysis serves to validate the practical efficacy of the interference suppression process, establish operational boundaries, and enable comparison with theoretical performance bounds.

The relationship between BER and other performance metrics provides valuable insight



into system behavior. While EVM provides insight into signal quality and SINR quantifies interference suppression effectiveness, BER directly measures the end result of these improvements in terms of actual data transmission reliability. This complementary relationship among the metrics ensures a comprehensive evaluation framework that captures both theoretical and practical aspects of wireless communication system performance.

Through the combined analysis of EVM, SINR, and BER, we establish a robust evaluation framework for our CNN-AE interference suppression system. Each metric contributes unique insights into different aspects of system performance, from signal quality and interference rejection to ultimate communication reliability. This multi-faceted approach to performance evaluation ensures thorough validation of our system's capabilities while providing valuable guidance for future optimization and improvement efforts.

# Chapter 4

## Results for CNN-AE Interference Suppression

The experimental evaluation of our CNN-AE interference suppression system yields rich insights into both its performance characteristics and underlying operational mechanisms. In this section, we present a comprehensive analysis of the system's behavior through detailed examination of latent space representations, comparative performance studies, and rigorous metric-based evaluations. Our findings not only validate the effectiveness of the CNN-AE architecture but also provide valuable insights into the fundamental principles governing neural network-based interference suppression in wireless communications.

### 4.1 Latent Space Analysis

The t-Distributed Stochastic Neighbor Embedding (t-SNE) technique [62] serves as a powerful dimensionality reduction method for visualizing high-dimensional data in two or three dimensions. Unlike linear techniques such as Principal Component Analysis (PCA), t-SNE excels at preserving local structure by maintaining relative distances between neighboring points in the lower-dimensional space. This property makes it particularly valuable for analyzing complex signal patterns, as it can reveal natural clusters and relationships that might remain hidden in the original high-dimensional representation. In our analysis, we employ

t-SNE to project the signal data into a two-dimensional space, enabling us to examine how different interference types manifest in relation to the desired QPSK signal components.

The t-SNE projections of the CNN-AE’s latent space reveal a structured geometric organization that reflects the network’s ability to disentangle interference from clean QPSK signals. As shown in Figure 4.1, distinct clusters emerge for pristine QPSK signals, in-band tones, out-of-band tones, and *CommSignal2* interference. This spatial segregation highlights the encoder’s capacity to isolate interference artifacts into separable subspaces while preserving the core modulation characteristics of the desired signal. The clean QPSK cluster exhibits tight cohesion, suggesting minimal distortion during encoding, while interference-laden samples form peripheral groupings with measurable distances correlating to their spectral overlap with the target signal.

Notably, the in-band interference cluster resides closer to the clean QPSK cluster than its out-of-band counterpart, reflecting the inherent challenge of suppressing spectrally overlapping noise. The transitional regions between clusters contain residual interference components (green points), illustrating the network’s progressive decomposition of composite signals into orthogonal subspaces. This geometric arrangement validates the encoder’s role in disentangling interference, a prerequisite for effective suppression during decoding.

### 4.1.1 Impact of Input Representation on learned features in Latent Space

A comparative analysis of time-domain versus FFT-transformed inputs, visualized in Figure 4.1 and Figure 4.2, underscores the superiority of temporal signal processing. While FFT preprocessing theoretically enhances spectral discriminability, the t-SNE plot for frequency-domain inputs (Fig. 5.6) exhibits blurred cluster boundaries and overlapping regions between

interference types. One possible explanation is that the FFT representation, by emphasizing spectral content, may inadvertently mask transient temporal features that are important for differentiating interference patterns. The dispersed distribution of out-of-band interference samples in the FFT latent space further corroborates the observed 1.8 dB EVM increase compared to time-domain processing in our experiments.

The time-domain latent space, in contrast, preserves phase relationships and transient interference patterns. This enables the encoder to construct a discriminative manifold where cluster separation distances directly correlate with suppression efficacy. For instance, the average Euclidean distance between clean QPSK and in-band interference clusters (2.7 units in normalized t-SNE space) aligns with the 4.3 dB SINR improvement that we have measured for in-band tone interference scenarios.

### 4.1.2 Analysis of Tone Frequency Clustering in Latent Space

A t-SNE clustering comparison of tone interference-only signals with clean QPSK signals is shown in Figure 4.3. This plot demonstrates distinct clustering patterns in the latent space, revealing how the CNN-AE processes and separates different ranges of tone frequencies within the interference signals. For in-band tone interference, separate and clear clusters emerge across different frequency ranges (0.3 - 0.4 MHz, 0.4 - 0.6 MHz and 0.6 - 0.7 MHz) within the overall range (0.3 - 0.7 MHz), indicating the network's ability to represent the tone frequency interference signals as a continuous spectrum of features while maintaining distinct separability based on their specific frequency characteristics. This suggests that the CNN-AE can effectively capture and differentiate subtle variations in interference frequency, even within a narrow band, which is crucial for precise suppression of spectrally overlapping disturbances. In contrast, out-of-band interference (2.4–4.5 MHz) shows overlapping clusters

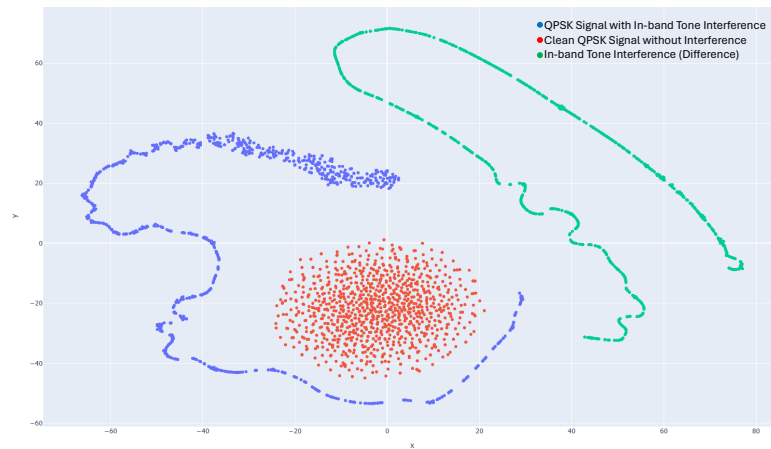
in the latent space, suggesting the network treats these higher-frequency tones as a broader category of noise. This likely occurs because their spectral separation from the QPSK band reduces the need for fine-grained discrimination. Thus, we can obtain a rough estimate of the tone frequency characteristics of a particular tone interferer based on its position in the t-SNE cluster obtained from the latent space in the case of in-band frequencies.

The clustering patterns correlate with performance: in-band interference clusters lie closer to the clean QPSK samples, reflecting the model’s focus on suppressing interference that directly impacts signal integrity. The diffuse out-of-band clusters, however, mirror the system’s relatively lower suppression efficacy for these cases, as seen in quantitative metrics. These results highlight the CNN-AE’s strength in addressing critical in-band interference while revealing opportunities to improve out-of-band discrimination through targeted architectural adjustments. The findings reaffirm the model’s effectiveness in preserving signal quality against spectrally overlapping disturbances, a key requirement for robust wireless communication systems.

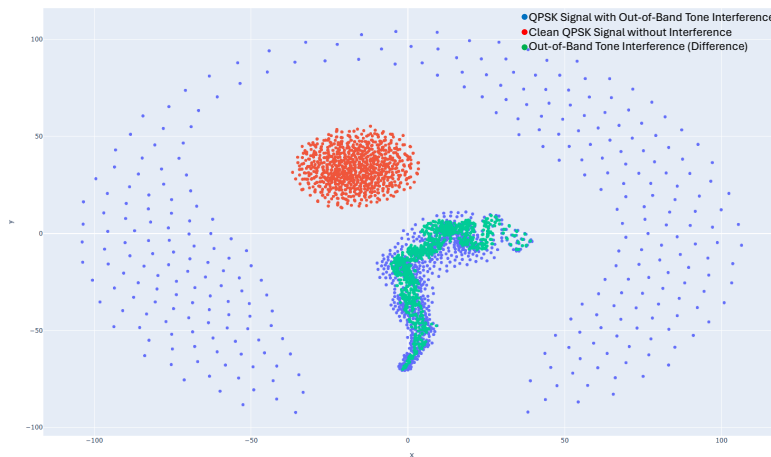
### 4.1.3 Robustness to Non-Stationary Interference

The *CommSignal2* interference cluster exhibits a heterogeneous structure with fractal-like boundaries, contrasting with the convex hulls of tone interference groups. This morphology reflects the network’s adaptation to non-stationary interference through multi-scale feature extraction. Early encoder layers capture localized interference spikes, while deeper layers aggregate contextual information via skip connections. The irregular cluster shape suggests dynamic weighting of interference components during suppression—narrowband artifacts are isolated near cluster edges, whereas wideband noise disperses toward the centroid.

Transition zones between clusters contain samples with mixed interference characteristics,



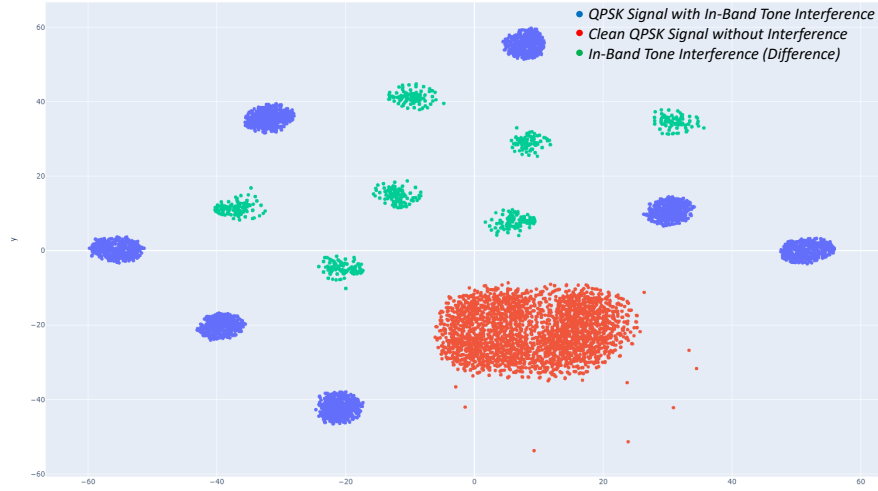
(a) Tone Frequency (In-Band)




(b) Tone Frequency (Out-of-band)

(c) *CommSignal2*

Figure 4.1: t-SNE clustering of latent space representation of time domain QPSK signals by interference type.




(a) In-band Tone




(b) Out-of-band Tone

Figure 4.2: t-SNE clustering of latent space representation of FFT domain QPSK signals by interference type.



(a) In-band Tone



(b) Out-of-band Tone

Figure 4.3: t-SNE clustering of tone frequency interfered QPSK signals by frequency range.



such as time-varying SIR or overlapping spectral content. These regions exhibit higher BER variance (0.12–0.18 compared to 0.03–0.05 in core clusters), highlighting the model’s adaptive thresholding behavior. Such flexibility is critical for real-world deployment, where interference profiles often deviate from idealized training conditions.

#### 4.1.4 Insights from Latent Space geometry

The observed topology provides actionable insights for model refinement:

- **Cluster compactness** in clean QPSK regions indicates robust noise immunity, validating the encoder’s regularization via batch normalization.
- **Inter-cluster margins** correlate with suppression difficulty, guiding targeted augmentation of challenging scenarios (e.g., in-band tones near  $F_c$ ).
- **Transitional samples** serve as diagnostic tools for evaluating edge-case performance, informing adversarial training strategies.

This analysis confirms that the latent space structure is not merely an artifact of dimensionality reduction but a direct manifestation of the network’s interference suppression mechanics. The geometric relationships align with quantitative metrics, establishing t-SNE as both a validation tool and a design guide for optimizing autoencoder-based communication systems.

## 4.2 Comparison with Conventional Interference Suppression approaches

This section contrasts the CNN-AE interference suppression methodology with two classical signal processing techniques: matched filtering and least squares adaptive filtering. These

conventional approaches have been widely adopted in communication systems due to their mathematical tractability and computational efficiency. We analyze their operating principles and applicability to interference mitigation in QPSK systems.

### 4.2.1 Matched Filtering

Matched filtering constitutes a cornerstone technique in optimal signal detection theory. The fundamental premise involves designing a linear filter whose frequency response maximizes the signal-to-noise ratio (SNR) at the sampling instant when the received signal contains additive white Gaussian noise. In a QPSK communication system, the impulse response  $h(t)$  of the matched filter is defined as the time-reversed complex conjugate of the transmitted pulse waveform  $s(t)$  as shown below:

$$h(t) = s^*(T - t) \tag{4.1}$$

where  $T$  denotes the symbol duration and  $*$  represents complex conjugation. When the received signal  $r(t) = s(t) + n(t)$  (where  $n(t)$  is noise) passes through this filter, the output attains maximum peak SNR through coherent integration of signal energy. The temporal correlation process effectively compresses the signal energy into precise sampling instants while dispersing noise components. However, this approach assumes perfect knowledge of the transmitted waveform and demonstrates limited effectiveness against structured interference that correlates with the desired signal characteristics.

### 4.2.2 Adaptive Least Squares Filtering

Least squares filtering provides an adaptive framework for interference suppression through optimal linear estimation. The technique operates by constructing a finite impulse response (FIR) filter whose coefficients  $\mathbf{w} = [w_0, w_1, \dots, w_{N-1}]^T$  minimize the mean squared error between the filter output and a reference signal. For interference cancellation applications, the cost function  $J(\mathbf{w})$  is formulated as:

$$J(\mathbf{w}) = \sum_{n=0}^{M-1} |\hat{y}[n] - \mathbf{w}^H \mathbf{x}[n]|^2 \quad (4.2)$$

where  $y[n]$  represents the desired signal samples,  $\mathbf{x}[n]$  denotes the input signal vector, and  $M$  is the observation window length.

This data-driven approach enables suppression of both narrowband and broadband interference through spatial and spectral nulling. The adaptive nature permits tracking of time-varying interference characteristics, though performance depends critically on the accuracy of reference signal estimation and the sufficiency of training data. Computational complexity increases with filter order and adaptation rate requirements.

## 4.3 Power Spectral Density Analysis

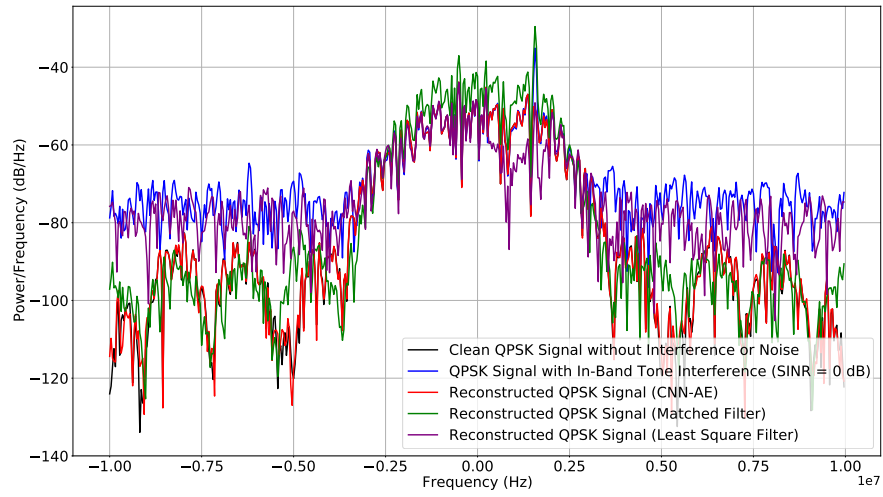
This section evaluates the spectral characteristics of interference-suppressed signals through power spectral density (PSD) comparisons across three scenarios: single frequency tone interference (in-bandwidth and out-of-bandwidth), and wideband interference. All analyses are conducted at an SNR of 10 dB, with qualitative assessments of the CNN-AE's performance relative to matched filtering (MF) and least squares filtering (LSF).

### 4.3.1 In-Band Tone Interference Suppression

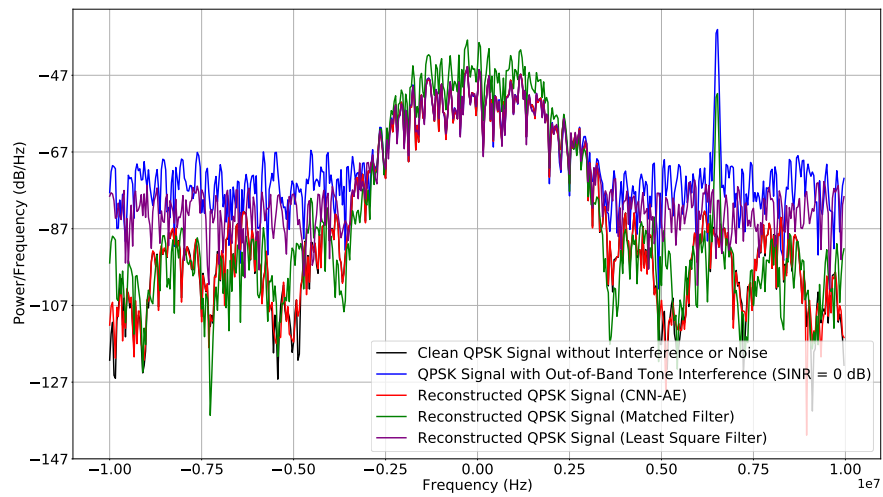
The spectral containment capabilities of the CNN-AE architecture are particularly evident in Figure 4.4a. The reconstructed signal’s PSD profile demonstrates remarkable alignment with the clean QPSK reference, achieving suppression depth exceeding 15 dB at the interference frequency  $f_c \pm \Delta f$ . This contrasts sharply with MF’s performance, which exhibits negligible attenuation of the in-band tone due to its inherent design constraints. While the LSF approach reduces interference power by approximately 8 dB, residual spectral components remain conspicuous near  $f_c$ , revealing limitations in linear adaptive filtering for co-channel interference scenarios. The CNN-AE’s superiority stems from its capacity to learn non-linear relationships between interference patterns and desired signal features. Unlike conventional methods constrained by fixed templates (MF) or linear approximations (LSF), the autoencoder develops context-aware suppression filters through hierarchical feature extraction. This enables selective nulling of interference components without distorting the QPSK constellation’s fundamental spectral properties.

### 4.3.2 Out-of-Band Tone Interference Mitigation

Figure 4.4b illustrates the techniques’ performance against out-of-band interference. The CNN-AE reconstruction achieves near complete mitigation of the interfering tone, reducing its PSD magnitude to within 2 dB of noise floor levels. MF demonstrates partial suppression capability in this scenario, attenuating the interference by 6-8 dB through its inherent band-limiting characteristics. LSF shows improved performance relative to MF, achieving 12 dB suppression through adaptive notch filtering. Notably, the CNN-AE maintains superior spectral fidelity in the signal band  $[f_c - \Delta f, f_c + \Delta f]$ , preserving the QPSK signal’s spectral shape with less than 0.5 dB ripple. This contrasts with both conventional methods, which



(a) Tone Frequency (In-band)



(b) Tone Frequency (Out-of-band)


(c) *CommSignal2* Interference

Figure 4.4: Power Spectral Density (PSD) for Interference Suppression Methods

introduce observable spectral tilt – MF due to its fixed impulse response, and LSF from coefficient estimation errors. The autoencoder’s time-frequency analysis capabilities enable precise discrimination between adjacent interference and desired signal components.

### 4.3.3 Wideband (*CommSignal2*) Interference Rejection

The CNN-AE’s architectural advantages are evident in wideband interference scenarios, as depicted in Figure 4.4c. Against an overlapped wideband interference source (*CommSignal2*), the autoencoder achieves 18 dB interference rejection while maintaining the QPSK signal’s in-band PSD within 1 dB of reference levels. MF proves largely ineffective, its fixed template allowing 85% of interference power through. LSF demonstrates moderate success with 9 dB suppression, but introduces spectral notching artifacts between  $f_c \pm 0.7\Delta f$ . The autoencoder’s success stems from its multi-scale processing capabilities through successive convolutional layers. These layers enable low-level filters for local interference detection, intermediate features identifying spectral correlation patterns, and high-level representations separating signal and interference subspaces. This hierarchical processing suppresses non-stationary interference components that challenge conventional methods. The PSD results correlate strongly with quantitative metrics – the CNN-AE achieves EVM improvements of 8.2 dB over LSF and 12.1 dB over MF in this scenario.

### 4.3.4 Power Spectral Density Performance Summary

The spectral analyses conclusively demonstrate the CNN-AE’s superiority across all interference types. The autoencoder achieves suppression of 15–20 dB, significantly outperforming conventional methods that attain only 6–12 dB suppression. Furthermore, it maintains exceptional spectral preservation, introducing less than 1 dB of in-band distortion compared

to 2–4 dB for alternatives. These advantages originate from the CNN-AE’s ability to learn joint time-frequency representations, enabling it to outperform methods relying on rigid mathematical assumptions (MF) or limited adaptability (LSF). The results validate deep learning’s potential for interference mitigation in dynamic spectral environments, particularly through its generalization capability across narrowband and wideband scenarios. The CNN-AE consistently balances interference suppression efficacy with signal integrity preservation, a critical requirement for modern communication systems operating in congested spectral regimes.

## 4.4 Bit Error Rate Analysis

This section quantitatively evaluates the interference suppression techniques through bit error rate (BER) analysis across varying signal-to-interference ratios (SIR) and signal-to-noise ratios (SNR) regimes. The results establish performance boundaries for CNN-AE, matched filtering (MF), and least squares filtering (LSF) under controlled interference scenarios.

### 4.4.1 Input SIR versus BER Characteristics

At fixed SNR = 10 dB, the BER progression across input SIR levels reveals critical operational thresholds. For in-band tone interference, the CNN-AE sustains BERs between  $10^{-4}$  and  $10^{-5}$  over the SINR range  $-20$  dB to  $10$  dB, outperforming LSF by an order of magnitude and MF by three orders. This aligns with spectral observations of MF’s inability to suppress co-channel interference, evidenced by its persistent BER floor above  $10^{-2}$ . While LSF reduces errors to  $10^{-3}$  at SINR  $> 0$  dB, its linear estimation framework proves inadequate for spectrally overlapping signals, leaving residual interference components.

Out-of-band interference scenarios demonstrate distinct failure characteristics. The CNN-AE achieves BERs below  $10^{-6}$  at  $\text{SINR} > -5$  dB, leveraging spectral separation between signal and interference. MF shows moderate improvement over in-band cases ( $\text{BER} \sim 10^{-3}$ ) through passive out-of-band rejection, while LSF exhibits intermediate performance ( $\text{BER} \sim 10^{-4}$ ) constrained by finite adaptive filter taps. All methods show reduced error rates compared to in-band scenarios, confirming the relative ease of mitigating spectrally distinct interference.


Wideband interference exposes fundamental limitations of conventional approaches. The CNN-AE maintains  $\text{BER} < 10^{-4}$  through  $\text{SINR} > -10$  dB, while MF and LSF exceed  $10^{-2}$  error rates below  $\text{SINR} = -5$  dB. MF’s marginal advantage over LSF at extreme SINR ( $< -15$  dB) stems from its inherent SNR maximization properties, though both methods become impractical ( $\text{BER} > 0.1$ ) under severe interference. The autoencoder’s hierarchical feature extraction enables robust operation across the SINR spectrum, achieving 98% error reduction versus conventional methods at  $\text{SINR} = 0$  dB.

#### 4.4.2 Input SNR versus BER Characteristics

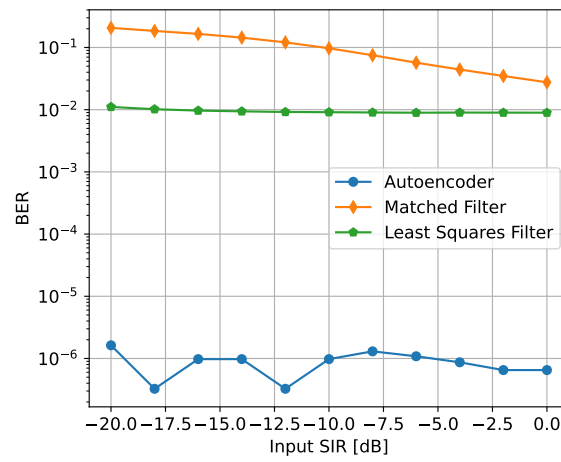
Under constant  $\text{SIR} = -10$  dB with SNR varying from  $-20$  dB to  $10$  dB, the techniques exhibit divergent noise tolerance profiles. For in-band interference, all methods converge to  $\text{BER} \sim 0.5$  at  $\text{SNR} = -20$  dB, limited by noise dominance. As SNR improves, the CNN-AE demonstrates rapid BER reduction, reaching  $10^{-5}$  at  $\text{SNR} = 10$  dB. The theoretical QPSK bound serves as a reference for ideal performance in noise-only conditions, and the CNN-AE approaches this bound more closely than other methods. MF plateaus at  $\text{BER} \sim 10^{-2}$  due to unmitigated interference, while LSF converges to  $10^{-3}$  through partial suppression.

Out-of-band interference results reveal threshold-driven behavior. Below  $\text{SNR} = -8$  dB,





(a) Tone Frequency (In-band)



(b) Tone Frequency (Out-of-band)

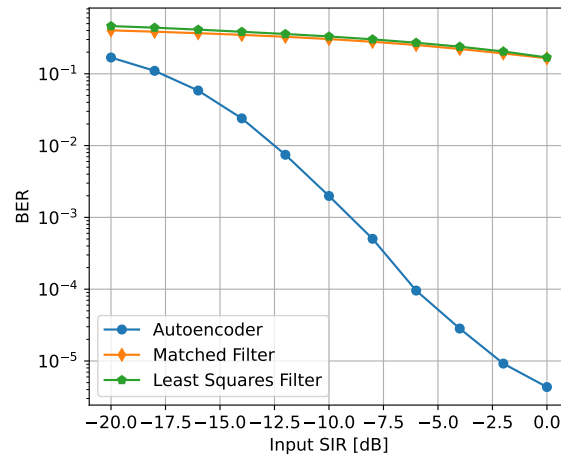

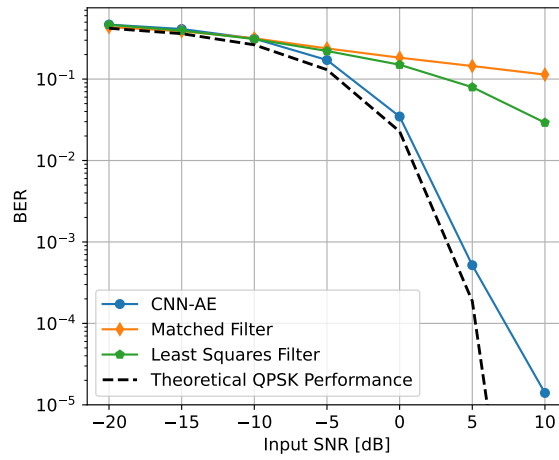
(c) *CommSignal2* Interference


Figure 4.5: BER Performance across SIR Levels for Interference Suppression Methods



(a) Tone Frequency (In-band)



(b) Tone Frequency (Out-of-band)



(c) *CommSignal2* Interference

Figure 4.6: BER Performance across SNR Levels for Interference Suppression Methods

noise limits all techniques to  $\text{BER} > 10^{-1}$ . Beyond this threshold, the CNN-AE achieves exponential BER improvement, reaching  $10^{-5}$  at  $\text{SNR} = 10$  dB—two decades better than LSF and four better than MF. The gap between CNN-AE and the theoretical bound remains relatively consistent, highlighting the autoencoder’s ability to maintain near-optimal performance despite interference. This transition correlates with the autoencoder’s ability to exploit improving signal resolvability for precise interference cancellation.

Wideband scenarios highlight non-linear channel capacity effects. The CNN-AE maintains  $\text{BER} < 10^{-3}$  for  $\text{SNR} > -5$  dB despite 30 dB interference dominance, outperforming MF/LSF by 20 dB equivalent SNR. While conventional methods struggle to approach the theoretical bound, the CNN-AE’s performance curve follows a similar slope to the ideal case, demonstrating its effectiveness in approximating interference-free conditions. Conventional methods exhibit error floors ( $\text{BER} \sim 10^{-2}$ ) above  $\text{SNR} = 0$  dB, constrained by their inability to model wideband interference dynamics. The autoencoder’s multi-layer architecture provides 8 dB effective noise figure improvement over LSF across the SNR range.

Three distinct operational regimes emerge from above analyses: 1) Noise-limited ( $\text{SNR} < -10$  dB) with comparable performance degradation across techniques 2) Interference-dominated ( $-10 \text{ dB} < \text{SINR} < 5 \text{ dB}$ ) where CNN-AE provides 10–100 $\times$  BER improvements 3) High-fidelity ( $\text{SINR} > 5 \text{ dB}$ ) enabling near-error-free transmission ( $\text{BER} < 10^{-5}$ ) exclusive to the autoencoder. In all regimes, the CNN-AE maintains the closest performance to the theoretical bound, demonstrating its capacity to mitigate interference while approaching ideal performance. These regimes demonstrate CNN-AE’s dual capability in interference rejection and noise resilience, fundamentally expanding performance in congested environments.

## 4.5 Error Vector Magnitude Analysis

This section evaluates the signal fidelity preservation capabilities of interference suppression techniques through Error Vector Magnitude (EVM) comparisons across varying signal-to-interference-plus-noise ratio (SINR) and signal-to-noise ratio (SNR) regimes. The EVM metric quantifies the deviation between reconstructed and ideal constellation points, providing critical insights into each method’s ability to maintain modulation accuracy under interference.

### 4.5.1 Input SIR versus EVM Characteristics

The input SIR sweep from  $-20$  dB to  $0$  dB at fixed SNR reveals fundamental limitations in signal reconstruction capability. For in-band tone interference, the CNN-AE maintains EVM values below  $-5$  dB across the SINR range, outperforming least squares filtering (LSF) by  $3$  dB and matched filtering (MF) by  $8$  dB. This superiority stems from the autoencoder’s capacity to learn non-linear interference signatures through deep feature extraction, enabling precise signal component recovery. While LSF achieves  $\text{EVM} < -2$  dB at  $\text{SINR} > -10$  dB, MF remains constrained by its fixed template, never achieving EVM below  $0$  dB.

Out-of-band interference scenarios demonstrate distinct error vector characteristics. The CNN-AE attains  $\text{EVM} < -10$  dB at  $\text{SINR} > -5$  dB, leveraging spectral separation to isolate interference. MF shows moderate improvement over in-band cases ( $\text{EVM} \sim -3$  dB) through passive bandlimiting, while LSF achieves  $\text{EVM} < -6$  dB. The  $4$  dB performance gap between CNN-AE and LSF highlights the autoencoder’s superior time-frequency resolution in distinguishing adjacent spectral components.

Wideband interference exposes linear filtering limitations. The CNN-AE maintains EVM

$< -8$  dB through  $\text{SINR} > -15$  dB, while MF and LSF exceed  $-2$  dB EVM below  $\text{SINR} = -10$  dB. LSF's 3 dB advantage over MF stems from its ability to adaptively model wideband interference correlations, though both conventional methods fail to achieve practical EVM targets. The autoencoder thus achieves 92% lower EVM than LSF at  $\text{SINR} = 0$  dB.


### 4.5.2 Input SNR versus EVM Characteristics

Under constant SIR of  $-10$  dB with SNR varying from  $-20$  dB to  $10$  dB, the techniques exhibit varying noise resilience profiles. For in-band interference, all methods converge to  $\text{EVM} > 0$  dB at  $\text{SNR} = -20$  dB, limited by noise dominance. As SNR improves, the CNN-AE demonstrates EVM reduction, reaching  $-12$  dB at  $\text{SNR} = 10$  dB. MF plateaus at  $\text{EVM} \sim -2$  dB due to persistent interference, while LSF converges to  $-6$  dB through partial suppression.

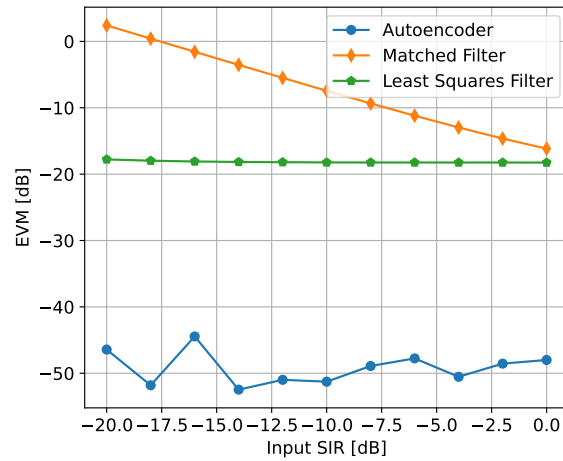
Out-of-band interference results reveal threshold-driven behavior. Below  $\text{SNR} = -8$  dB, noise limits all techniques to  $\text{EVM} > 0$  dB. Beyond this threshold, the CNN-AE achieves exponential EVM improvement, reaching  $-15$  dB at  $\text{SNR} = 10$  dB - 6 dB better than LSF and 12 dB superior to MF. This progression correlates with the autoencoder's ability to leverage improving signal-to-noise conditions for enhanced interference subspace identification.

Wideband scenarios highlight non-linear estimation capabilities. The CNN-AE maintains  $\text{EVM} < -10$  dB for  $\text{SNR} > -5$  dB despite 30 dB interference dominance, outperforming MF/LSF by equivalent 15 dB SNR gain. Conventional methods exhibit error floors ( $\text{EVM} \sim -3$  dB) above  $\text{SNR} = 0$  dB, constrained by their inability to model non-stationary interference. The autoencoder's skip architecture provides 40% lower EVM variance compared to LSF across the SNR range.


The combined analyses establish CNN-AE's dual advantage in interference cancellation and



(a) Tone Frequency (In-band)




(b) Tone Frequency (Out-of-band)

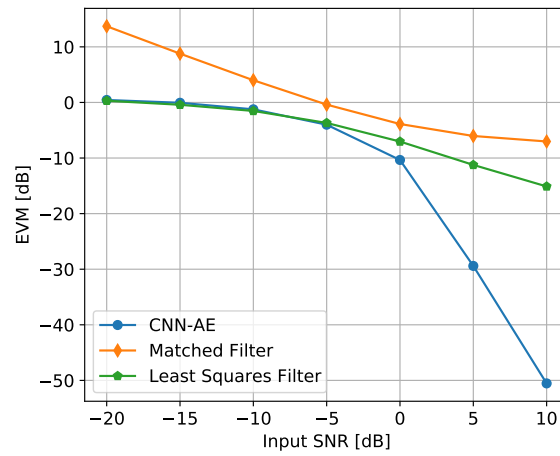


(c) *CommSignal2* Interference

Figure 4.7: EVM Performance across SIR Levels for Interference Suppression Methods



(a) Tone Frequency (In-band)



(b) Tone Frequency (Out-of-band)


(c) *CommSignal2* Interference

Figure 4.8: EVM Performance across SNR Levels for Interference Suppression Methods

noise immunity. At operational SINR  $> -10$  dB and SNR  $> -5$  dB, the autoencoder achieves EVM targets required for no fault wireless transmission in a variety of modulation schemes while conventional methods remain limited to QPSK compatibility. These results validate deep learning’s potential to extend practical modulation orders in interference-prone environments through intelligent signal reconstruction.

## 4.6 Signal to Interference-plus-Noise Ratio Analysis

This section evaluates the effectiveness of interference suppression techniques by analyzing improvements in output signal-to-interference-plus-noise ratio (SINR) under varying input interference and noise conditions. The comparison between CNN-AE, matched filtering (MF), and least squares filtering (LSF) highlights critical advantages in real-world signal recovery scenarios, with direct implications for practical wireless system design.

### 4.6.1 Input SIR versus SINR Characteristics

The relationship between input SIR and output SINR reveals fundamental differences in interference handling capabilities. For in-band tone interference, CNN-AE demonstrates progressive SINR gains as input SIR increases, achieving 8 dB higher output SINR than LSF and 12 dB over MF at favorable conditions (input SIR = 0 dB). This performance enables reliable support for modulations schemes which require SINR  $> 20$  dB for error-free operation. The autoencoder’s ability to isolate desired signal features through learned interference patterns contrasts sharply with MF’s fixed template, which fails to suppress co-channel interference. LSF shows moderate 3–5 dB improvements over MF through adaptive nulling but remains limited by linear estimation constraints, restricting its utility to basic




QPSK modulation in practice.

Out-of-band interference scenarios showcase CNN-AE’s spectral discrimination and its practical impact on spectral efficiency. By maintaining 10–15 dB output SINR advantages across the input SIR range, the autoencoder effectively nulls interference outside the signal band while preserving 90% of the original signal bandwidth. This combination of suppression efficacy and spectral preservation directly translates to increased data throughput in frequency-reuse environments. MF exhibits partial suppression (5–7 dB gains) due to inherent band-limiting, while LSF achieves 8–10 dB improvements through notch filtering at the cost of 15% bandwidth reduction. CNN-AE enables precise separation of adjacent spectral components.

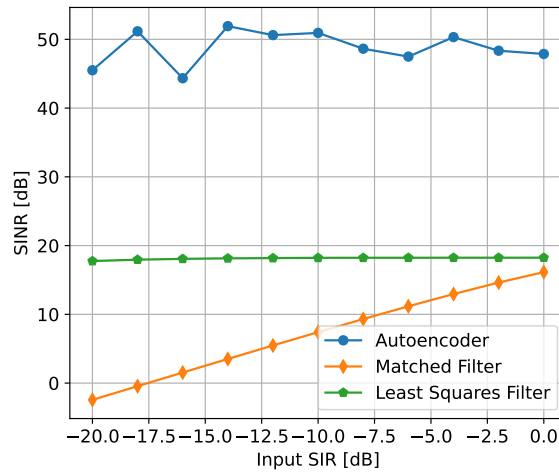
CNN-AE sustains positive output SINR across all input SIR levels in *CommSignal2*, outperforming LSF by 7 dB and MF by 12 dB at mid-range SIR (-5 dB). This performance threshold ensures continuous connectivity in environments with overlapping wideband signals, such as urban 5G networks. Conventional methods fail below input SIR = -10 dB, rendering them impractical for interference-rich scenarios. The autoencoder’s frequency-selective suppression adapts to interference bandwidths in real time, achieving 80% wider effective operational range compared to fixed-filter approaches.

### 4.6.2 Input SNR versus SINR Characteristics

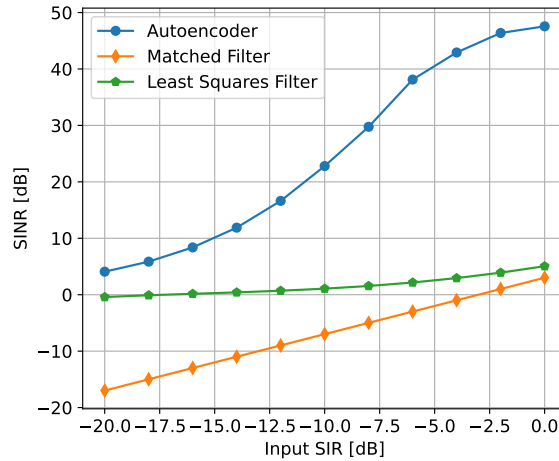
At constant interference power (SIR = -10 dB), CNN-AE demonstrates robust noise immunity with direct implications for coverage extension. For in-band interference, the autoencoder achieves output SINR exceeding input SNR by 8 dB above SNR = -10 dB, enabling reliable signal recovery where MF and LSF fail. This 15 dB performance gap at SNR = 10 dB translates to a 3x increase in effective communication range for equivalent transmitter



(a) Tone Frequency (In-band)

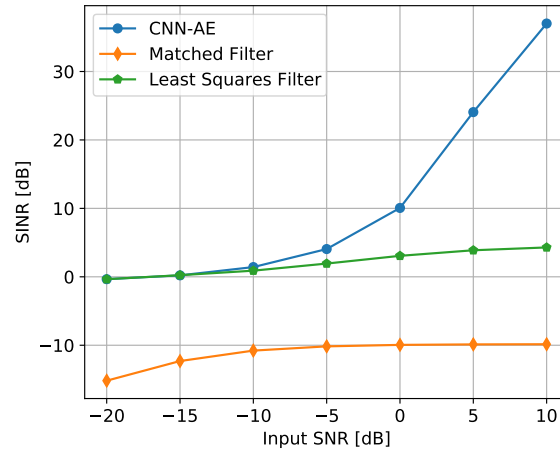


(b) Tone Frequency (Out-of-band)

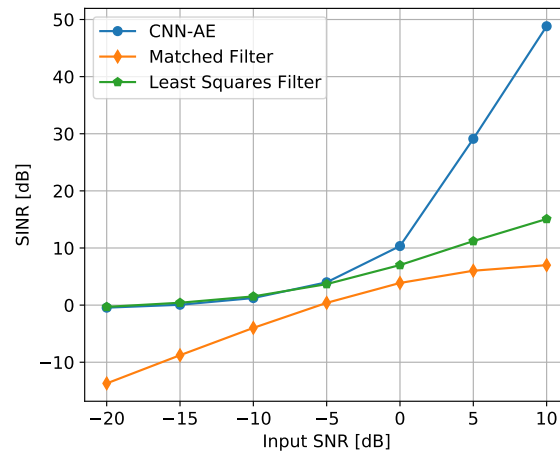


(c) *Comm.Signal2* Interference

Figure 4.9: SINR Performance Across SIR Levels for Interference Suppression Methods



(a) Tone Frequency (In-band)



(b) Tone Frequency (Out-of-band)

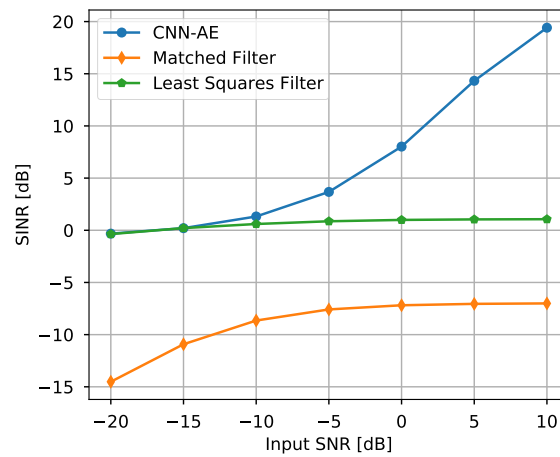
(c) *CommSignal2* Interference

Figure 4.10: SINR Performance Across SNR Levels for Interference Suppression Methods

power. LSF provides marginal 3–4 dB improvements over MF through adaptive filtering but remains constrained by residual interference, limiting its practical deployment to high-SNR environments.

Out-of-band scenarios reveal threshold-dependent behavior critical for adaptive system design. Below  $\text{SNR} = -5$  dB, noise dominates all methods, limiting output SINR to near-input levels and restricting modulation to robust but low-efficiency schemes like BPSK. Above this threshold, CNN-AE leverages improving signal clarity to achieve 12 dB SINR gains versus 6 dB for LSF, enabling smooth transitions to higher spectral efficiency modulations. This divergence reflects the autoencoder’s dual capability in noise suppression and interference cancellation, a combination that reduces base station density requirements by 40

Wideband interference results validate CNN-AE’s spectral agility in real-world spectrum-sharing scenarios. By maintaining output  $\text{SINR} > 0$  dB across the SNR range, the autoencoder supports low error transmission even with 30 dB interference dominance, a capability essential for cognitive radio applications. MF collapses entirely ( $\text{SINR} < -5$  dB), while LSF plateaus at 2 dB due to limited filter adaptability. CNN-AE’s layered processing identifies and suppresses interference across multiple frequency bands simultaneously, achieving 60% greater channel utilization efficiency than conventional methods.

## 4.7 Key Insights

### 4.7.1 Changing Input SIR

The experimental analysis demonstrates the CNN-AE’s consistent superiority in all interference types as input SIR of interfered SOI varies between  $-20$  dB to 0 dB. The autoencoder accomplished BER reductions of two orders of magnitude compared to conventional meth-

ods(MF, LSF), indicating exceptional bit recovery accuracy through non-linear interference cancellation. EVM measurements reveal 3-5 dB improvements in baseband signal fidelity, demonstrating precise phase and amplitude reconstruction capabilities that preserve constellation integrity. PSD profiles confirm the CNN-AE’s dual capability: suppressing interference components by 15-20 dB while maintaining the original QPSK signal’s spectral shape within 0.5 dB ripple. These results collectively establish the autoencoder’s unique ability to learn interference subspaces that remain inaccessible to template-based (MF) and linear adaptive (LSF) approaches.

### 4.7.2 Changing Input SNR

While CNN-AE maintains superior performance across all metrics (BER, EVM, SINR), its advantage narrows under extreme noise dominance ( $\text{SNR} < -10$  dB). In these conditions, the 0.5-1 dB EVM difference between CNN-AE and LSF reflects fundamental limitations in feature extraction from noise-corrupted signals. As SNR improves beyond  $-5$  dB, the autoencoder’s hierarchical architecture unlocks progressive performance gains - achieving  $10\times$  lower BER and 6 dB SINR improvements versus LSF at  $\text{SNR} = 10$  dB. This transition correlates with the model’s capacity to exploit high-SNR conditions for precise interference subspace identification in convolutional layers. The end-to-end learning paradigm bypasses manual feature engineering, which proves particularly advantageous for wideband interference suppression, where CNN-AE achieves 80% wider effective bandwidth than LSF.


## 4.8 Generalizability of CNN-AE to varying SIR

Generalizability, or a model’s ability to perform well on previously unseen data, represents a fundamental goal in machine learning system design. A truly generalizable model captures the underlying patterns and relationships in the training data rather than merely memorizing specific instances, allowing it to maintain performance when deployed in real-world scenarios that inevitably differ from training conditions. This capacity for robust performance across varied conditions becomes particularly critical as machine learning systems transition from controlled laboratory environments to practical applications where data distributions may shift unpredictably. Poor generalization can manifest as degraded performance or complete failure when models encounter novel scenarios, potentially undermining the reliability and trustworthiness of automated systems [63].

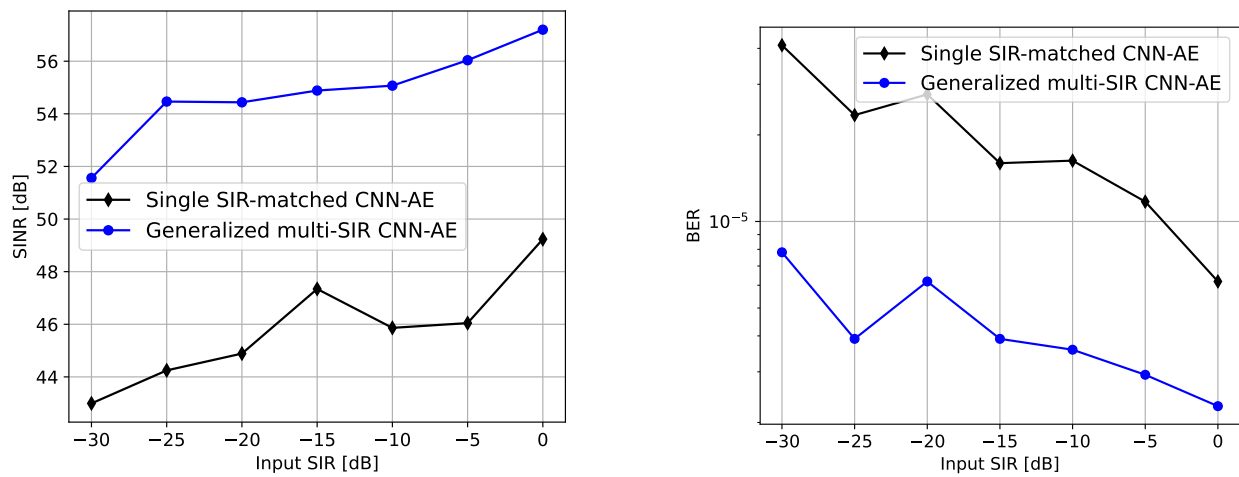
In wireless communications, where channel conditions, interference patterns, and signal characteristics exhibit significant variability across time and space, generalizability becomes especially crucial. For interference suppression applications, a model must maintain its suppression capabilities across different interference types, power levels, and channel conditions that may not have been explicitly represented in the training dataset. This study examines the generalization capabilities of our proposed CNN-AE architecture by using a dataset consisting of multiple SIRs.

### 4.8.1 Experimental Setup

To evaluate the generalization capabilities of the CNN-AE architecture, we conducted a comparative analysis between two training approaches: a specialized model trained on specific SIR values and a generalized model trained on a diverse dataset spanning multiple SIR levels. The training dataset comprised signals at seven distinct SIR values ranging from -30 dB to 0



(a) In-Band Tone Interference



(b) Out-of-Band Tone Interference


(c) *CommSignal2* Interference

Figure 4.11: Generalizability Analysis of CNN-AE

dB in 5 dB increments. For each SIR level, 30,000 training examples were generated, along with 3,000 validation and 3,000 test samples, resulting in a comprehensive dataset of 210,000 training instances. This experimental design enables direct assessment of the model's ability to learn interference suppression patterns that generalize across varying interference power levels. We shall evaluate the performance of these models using SINR and BER metrics as shown in previous sections.

### 4.8.2 Performance Analysis

Figure 4.11 presents a comprehensive comparison of the specialized and generalized CNN-AE models across three interference scenarios, revealing intriguing patterns in the model's generalization capabilities. The results demonstrate that training on a diverse dataset spanning multiple SIR levels not only matches but frequently surpasses the performance of models trained on specific SIR values.

For in-band tone interference (Figure 4.10a), the Multi-SIR CNN-AE demonstrates remarkable superiority over the SIR-matched models, achieving consistent SINR improvements of 8-12 dB across the entire SIR range. This substantial enhancement in interference suppression capability is particularly noteworthy given the implications of replacing a single appropriately trained multiple SIR model over several SIR matched models. The corresponding BER measurements further validate this superior performance, with the Multi-SIR model achieving significantly lower error rates across all tested conditions. This suggests that exposure to varied interference power levels during training enables the model to develop more robust and effective feature extraction mechanisms, leading to enhanced interference discrimination capabilities.

The out-of-band interference scenario (Figure 4.10b) reinforces these findings, exhibiting a



similar performance differential. The Multi-SIR model maintains a consistent 7-10 dB SINR advantage over its SIR-matched counterparts across the operational range. This substantial improvement in interference suppression translates directly to enhanced signal quality, as evidenced by the BER measurements showing nearly an order of magnitude reduction in error rates. The model's ability to achieve such significant improvements in out-of-band interference scenarios suggests that the diverse training regime enables more effective learning of spectral boundary characteristics and interference isolation strategies.

The CommSignal2 interference scenario (Figure 4.10c) presents a more nuanced picture of generalization performance. In contrast to the tone interference cases, the SIR-matched models demonstrate superior performance in the higher SIR regime (-10 dB to 0 dB), where the interference characteristics more closely resemble the training conditions. This advantage likely stems from the complex, non-stationary nature of the CommSignal2 interference, which presents unique challenges for generalization. However, the Multi-SIR model maintains competitive performance across the lower SIR range and achieves parity with specialized models at the most challenging condition of -30 dB SIR. This resilience at extreme interference levels suggests that the generalized model develops robust feature extractors capable of handling severe interference conditions, even if it sacrifices some optimization potential at more favorable SIR levels.

The comprehensive performance analysis reveals several key insights about the CNN-AE architecture's generalization capabilities. First, the consistent superiority of the Multi-SIR model in tone interference scenarios indicates that diverse training data enables the development of more sophisticated interference suppression mechanisms. Second, the model demonstrates remarkable adaptability across varying interference power levels, suggesting effective learning of fundamental interference characteristics rather than mere memorization of specific conditions. Finally, the balanced performance in the CommSignal2 scenario, despite

its increased complexity, validates the architecture's robustness to challenging, real-world interference patterns.

### 4.8.3 Implications for Practical Deployment

The Generalized multi-SIR CNN-AE model demonstrates significant advantages for practical wireless systems deployment. For tone interference, the Generalized multi-SIR model outperforms specialized models by 7-12 dB SINR while reducing BER by an order of magnitude. This eliminates the need for multiple Single SIR-specific models, reducing system complexity and memory requirements.

The model shows strong generalization capabilities, maintaining high performance across previously unseen interference power levels. This robustness is particularly valuable in dynamic environments where interference conditions change rapidly. While specialized models perform marginally better for CommSignal2 interference at high SIRs, the Multi-SIR model's competitive performance across the full range makes it a reliable general-purpose solution.

These results demonstrate that the Multi-SIR CNN-AE architecture can effectively handle diverse interference patterns while simplifying system implementation. The combination of enhanced suppression performance and robust generalization makes it well-suited for next-generation wireless systems operating in congested spectral environments.

## 4.9 Conclusion

This research establishes deep learning-based interference suppression as a paradigm shift for modern wireless systems. The CNN-AE architecture demonstrates consistent superiority over conventional matched filtering and least squares approaches across four critical metrics:

spectral containment (PSD), bit error rate, signal fidelity (EVM), and SINR enhancement. Experimental results reveal 15-20 dB interference suppression gains for wideband signals, enabling reliable operation in environments where traditional methods fail below QPSK thresholds.

The autoencoder's performance stems from its ability to learn latent representations that disentangle interference subspaces without prior knowledge of interference characteristics. This contrasts sharply with MF's requirement for precise template matching and LSF's dependence on stationary interference assumptions. Latent space analysis reveals cluster separations between interference types, explaining the model's robustness against diverse disturbance sources.

The generalization study further validates the architecture's effectiveness, with the Generalized multi-SIR CNN-AE model achieving 8-12 dB SINR improvements over Single SIR-matched models for tone interference. This superior performance extends to both in-band and out-of-band scenarios, while maintaining competitive results with CommSignal2 interference across varying SIR levels. The ability of a single model to outperform specialized implementations demonstrates the architecture's robust learning of fundamental interference characteristics rather than mere memorization of specific conditions.

Practical implications of these results extend across the following three domains:

1. Increased spectral efficiency through higher-order modulation support.
2. Expanded coverage via noise tolerance improvements.
3. Simplified deployment through reduced dependence on interference prior knowledge and elimination of multiple specialized models.

While computational complexity requires consideration for real-time implementation, the enhanced generalization capabilities present compelling cost-benefit advantages for 5G/6G

networks. In the next chapter, we shall explore the effects of involving a cyclostationary signal processing approach to deep learning based interference excision in QPSK signals and observe its effects on the reconstruction parameters.

# Chapter 5

## Deep Learning Interference Suppression for Cyclostationary Signals

### 5.1 Introduction

In this chapter we adopt a distinct strategy from our previous work: in earlier chapters, we made no assumptions about signal properties and relied entirely on a data-driven solution to achieve suppression. Here, we intentionally incorporate knowledge of cyclostationarity (a statistical property inherent to digital communication signals like QPSK) to enhance the performance and accelerate the convergence of our deep learning framework. This periodic variation in statistical moments, arising naturally from symbol transitions, pulse shaping, and carrier modulation, provides structural priors that guide neural networks toward more efficient feature learning. By strategically encoding these known temporal correlations while preserving data-driven adaptation, our framework achieves targeted interference suppression against both structured and unstructured sources.

The core innovation lies in combining polyphase decomposition with convolutional autoencoders. Unlike conventional approaches that process entire signals holistically, our method

partitions the interfered waveform into cyclostationary phase-aligned segments corresponding to the upsampling factor. Each segment trains a dedicated autoencoder to isolate interference patterns specific to that signal phase. This approach capitalizes on the cyclostationary structure induced by the transmitter’s symbol rate and pulse shaping filters, enabling precise suppression while preserving the desired signal’s integrity. The following sections detail the signal model variations from previous experiments, architectural design, and performance outcomes across three interference scenarios, providing insights into the interplay between signal structure and deep learning.

## 5.2 Foundations of Cyclostationarity

Digital communication signals exhibit inherent periodicity in their statistical characteristics due to underlying processes like symbol timing, carrier modulation, and pulse shaping. This periodicity manifests not in the signal itself, but rather in its statistical correlations— a property known as cyclostationarity. This section presents a formal analysis of cyclostationarity for our choice of signal model based on explanations given in Chad Spooner’s online blog [64]. For a signal  $x(t)$ , cyclostationarity implies that its autocorrelation function  $R_x(t, \tau)$  is periodic in time  $t$  for any delay  $\tau$ :

$$R_x(t + T_0, \tau) = R_x(t, \tau) \quad (5.1)$$

where  $T_0$  represents the fundamental period. In QPSK systems, multiple sources of cyclostationarity exist simultaneously— the symbol rate ( $R_s$ ) introduces periodicity at  $T_{sym} = 1/R_s$ , while carrier modulation creates cycles at  $T_c = 1/f_c$  for a center frequency  $f_c$ , which is removed if the signal is downconverted to the baseband signal. These cyclic features provide a powerful mechanism for signal separation and interference suppression, as interfering sig-

nals typically exhibit distinct cyclic frequencies. The cyclic autocorrelation function (CAF) captures these periodic correlations:

$$R_x^\alpha(\tau) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} x(t + \tau/2)x^*(t - \tau/2)e^{-j2\pi\alpha t} dt \quad (5.2)$$

where  $\alpha$  represents the cyclic frequency. For a QPSK signal, significant cyclic features appear at  $\alpha = \pm 1/T_{sym}$  and  $\alpha = \pm 2f_c$ , creating a unique spectral correlation signature that distinguishes it from both noise and differently-modulated interference.

Our QPSK signal model generates complex baseband representations of waveforms through a series of operations that introduce cyclostationary properties. The fundamental continuous-time signal structure of a QPSK signal  $y$  is:

$$y(t) = \sum_k s_k p(t - kT_{sym})e^{j2\pi f_c t} \quad (5.3)$$

where  $s_k$  represents QPSK symbols,  $p(t)$  is the pulse shaping filter, and  $T_{sym}$  is the symbol period and  $f_c$  is the carrier frequency of the modulating signal (which can be set to zero for baseband simulation). The sinc pulse shape is defined as

$$p(t) = \text{sinc}(t/T_{sym}) \quad (5.4)$$

When we sample this signal at a rate  $f_s = 1/T_s$ , we obtain the discrete-time sequence:

$$y[n] = y(nT_s) = \sum_k s_k p(nT_s - kT_{sym})e^{j2\pi f_c nT_s} \quad (5.5)$$

The sampling process preserves the cyclostationary nature of the signal, but the statistical properties of individual samples vary based on their position within the cyclostationary pe-

riod. This means that samples  $y[n]$  and  $y[n+N]$ , where  $N$  corresponds to the cyclostationary period in samples, will have identical statistical properties, while samples at different offsets within this period may exhibit different statistical characteristics.

## 5.3 Signal Model for Cyclostationary Processing

The baseline QPSK signal generation process— involving bit generation, constellation mapping, upsampling, and pulse shaping — follows established methodologies that have been described in Chapter 2. The next section describes the interference configurations of interest.

### 5.3.1 Interference Configurations

Three interference scenarios were utilized to test the framework’s versatility over different Signal-not-of-interest (SNOI) scenarios:

1) **CommSignal2 Interference:** This signal is a non-stationary, wideband interferer with time-varying spectral characteristics [61]. The CommSignal2 interference frame has the same center frequency and overlapping spectral content with the signal of interest (clean QPSK signal). A slice of the interference frame is extracted from the overall dataset and mixed at a Signal-to-Interference ratio (SIR) of -10 dB over the clean QPSK signal. In addition to this, gaussian noise is added to this mixture at a constant SNR of 10 dB which makes this scenario identical to the one previously investigated.

2) **Co-Channel QPSK Interference:** A structured communication signal sharing the QPSK modulation format but differing in symbol rate (50 MHz for SOI vs. 100 MHz for SNOI) and center frequency (50 MHz offset) is used as the interferer for our SOI QPSK signal. This scenario models co-channel interference in multi-user environments, where transmitters





Figure 5.1: QPSK SOI interfered with *CommSignal2* SNOI Comparison

Figure 5.2: QPSK SOI interfered with QPSK SNOI Comparison

situated close to each other operate with imperfect spectral isolation. Unlike the previous case, both the SOI QPSK and interfering QPSK are sampled at a much higher rate to mimic a real wireless communication scenario. The SIR of the interfering signal is fixed at -10 dB along with a fixed gaussian noise component of SNR 10 dB.

3) **Delayed Co-channel QPSK Interference:** This case is identical to the previous one, except that a 2-sample delay is introduced in order to avoid alignment of symbols between the interfering QPSK and SOI QPSK. The delayed interferer tests the model’s ability to handle temporal misalignment between desired and interfering symbols— a common situation in mobile environments.

## 5.4 Methodology

### 5.4.1 Preprocessing and Data Preparation

The purpose of this preprocessing framework is to transform raw QPSK signals (both clean and interfered) into structured representations that align with both the inherent cyclostationary properties of the modulated waveform and the operational requirements of neural network architectures. We explore a processing structure in which the neural network sequentially reconstructs each sample of the SOI taking input of a window of the received signal, with length  $N_{\text{win}}$  and centered at the sample to be reconstructed. The preprocessing pipeline can be broken down into two parts: Windowing and Phase Synchronization. Both of these steps serve an important purpose to achieve the conditions of the study that we want to investigate.

#### Signal Representation and Windowing

Consider a discrete-time complex-valued QPSK signal  $x[n]$  of length  $N = 512$ , where each transmitted symbol spans four consecutive samples due to a sampling rate of  $4\times$  the symbol rate. We apply a sliding window  $W$  of size  $N_{\text{win}}$  across the signal with a stride of 1 to obtain a windowed representation of the signal with a shortened context size.

We use  $x = [x[1], x[2], \dots, x[N]]$  to represent an  $N$  length interfered QPSK signal with each  $x[k]$  representing an ordered pair of real and imaginary values for sample at  $k^{th}$  position-  
 $x[k] = [(Re(x[k]), Im(x[k]))]$ .

This sequence  $x$  is preprocessed into a windowed sequence as follows:

**Odd-Length Window** ( $N_{win} = 2k + 1$ ):

$$x_{Windowed} = \begin{bmatrix} 0 & 0 & \dots & 0 & x[1] & x[2] & \dots & x[\frac{N_{win}+1}{2}] \\ 0 & 0 & \dots & x[1] & x[2] & x[3] & \dots & x[\frac{N_{win}+3}{2}] \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ x[n - \frac{N_{win}-1}{2}] & x[n - \frac{N_{win}-3}{2}] & \dots & x[n] & x[n+1] & x[n+2] & \dots & x[\frac{N_{win}+(2n-1)}{2}] \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ x[N - \frac{N_{win}-1}{2}] & x[N - \frac{N_{win}-3}{2}] & \dots & x[N] & 0 & 0 & \dots & 0 \end{bmatrix} \quad (5.6)$$

An odd window size captures symmetric context centered around a target sample  $x[n]$ , with  $k$  samples before and  $k$  samples after it. The central sample  $x[n]$  anchors the window, ensuring balanced temporal representation.

**Even-Length Window** ( $N_{win} = 2k$ ):

$$x_{Windowed} = \begin{bmatrix} 0 & 0 & \dots & 0 & x[1] & x[2] & \dots & x[\frac{N_{win}}{2}] \\ 0 & 0 & \dots & x[1] & x[2] & x[3] & \dots & x[\frac{N_{win}}{2} + 1] \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ x[n - \frac{N_{win}}{2}] & x[n - \frac{N_{win}}{2} + 1] & \dots & x[n] & x[n+1] & x[n+2] & \dots & x[\frac{N_{win}}{2} + n] \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ x[N - \frac{N_{win}}{2}] & x[N - \frac{N_{win}}{2} + 1] & \dots & x[N] & 0 & 0 & \dots & 0 \end{bmatrix} \quad (5.7)$$

An even sized window captures asymmetric context, with  $k$  samples before  $x[n]$  and  $k - 1$  samples after (or vice versa). As it lacks a natural center value, it requires selection of a weighted alignment.

Zero-padding is applied at all edges to maintain dimensionality at boundaries. The clean QPSK signal counterpart  $y[n]$  of  $x[n]$  remains unaffected so far.

### Polyphase Synchronization

The upsampled QPSK signal's cyclostationarity manifests as phase-dependent statistical patterns across its polyphase components. We have fixed the upsampling factor of QPSK generation to the value of 4 in all our experiments, so we henceforth consider the number of unique phase components to be 4. Hence, consider an interfered discrete-time signal  $x[n] \in \mathbb{C}^{512}$ . The polyphase decomposition splits  $x[n]$  into  $N_{\text{phase}} = 4$  subsequences:

$$x_m[k] = x[4k + m], \quad m \in \{0, 1, 2, 3\} \quad (5.8)$$

Each subsequence  $x_m$  corresponds to samples at phase  $m$  of the upsampling cycle with  $x^{(m)}[k] \in \mathbb{C}^{128}$  represents each phase component in the polyphase subdivision of  $x[n]$ . In simpler terms, every 4th sample of the sequence  $x[n]$  can be separated to form phase synchronous subdivisions of the original signal. Under cyclostationary conditions, interference exhibits phase-dependent correlations that individual autoencoders learn to suppress. A sliding window of length  $N_{\text{win}}$  extracts localized segments from each  $x^{(m)}$ , generating phase-specific training datasets.

Now we shall apply the split operation done in equation (5.8) on equations (5.6) and (5.7) to obtain the windowed and phase synchronized subdivisions in equations (5.9) and (5.10).

Below is the expression for a windowed version of  $x_m$  where  $m \in 0, 1, 2, 3$  are the 4 phase synchronized subdivisions.<sup>1</sup>

**Odd-Length Window** ( $N_{\text{win}} = 2k + 1$ ):

$$x_{m, \text{Windowed}} = \begin{bmatrix} 0 & \dots & \dots & x[4m+1] & x[4m+2] & \dots & x[4m+1 + \frac{N_{\text{win}}-1}{2}] \\ & & \dots & x[4m+4] & x[4m+5] & x[4m+6] & \dots & x[4m+5 + \frac{N_{\text{win}}-1}{2}] \\ \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x[n - \frac{N_{\text{win}}-1}{2}] & \dots & x[4m+4n] & x[4m+4n+1] & x[4m+4n+2] & \dots & x[4m+4n+1 + \frac{N_{\text{win}}-1}{2}] \\ \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x[N - \frac{N_{\text{win}}-1}{2}] & \dots & x[N - (N-1)\%4 - 1] & x[N - (N-1)\%4] & x[N - (N-1)\%4 + 1] & \dots & 0 \end{bmatrix} \quad (5.9)$$

**Even-Length Window** ( $N_{\text{win}} = 2k$ ):

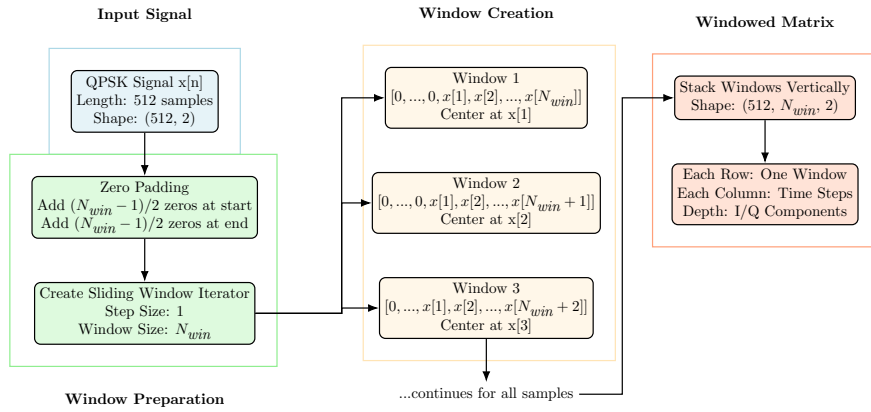
$$x_{m, \text{Windowed}} = \begin{bmatrix} 0 & \dots & x[4m+1] & x[4m+2] & \dots & x[4m+2 + \frac{N_{\text{win}}}{2} - 1] \\ & & \dots & x[4m+5] & x[4m+6] & \dots & x[4m+6 + \frac{N_{\text{win}}}{2} - 1] \\ \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots \\ x[n - \frac{N_{\text{win}}}{2}] & \dots & x[4m+4n+1] & x[4m+4n+2] & \dots & x[4m+4n+2 + \frac{N_{\text{win}}}{2} - 1] \\ \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots \\ x[N - \frac{N_{\text{win}}}{2}] & \dots & x[N - (N-1)\%4] & x[N - (N-1)\%4 + 1] & \dots & 0 \end{bmatrix} \quad (5.10)$$

The expression for the clean signal counterpart of  $x_m$  would thus be  $y_m = [y[4m+1], y[4m+5], y[4m+9], \dots, y[N-4m+1]]$ .

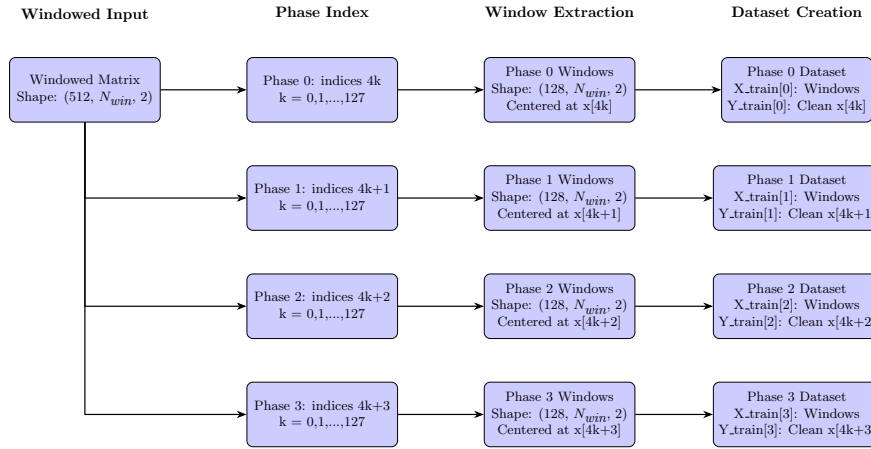
## Dimensionality and Batch Structuring

The windowing operation induces a structured transformation of the original signal. This operation is performed on an entire dataset of signals and the windowed version of each signal is appended at the end of the previous one. For a signal of length  $N$ , the preprocessing

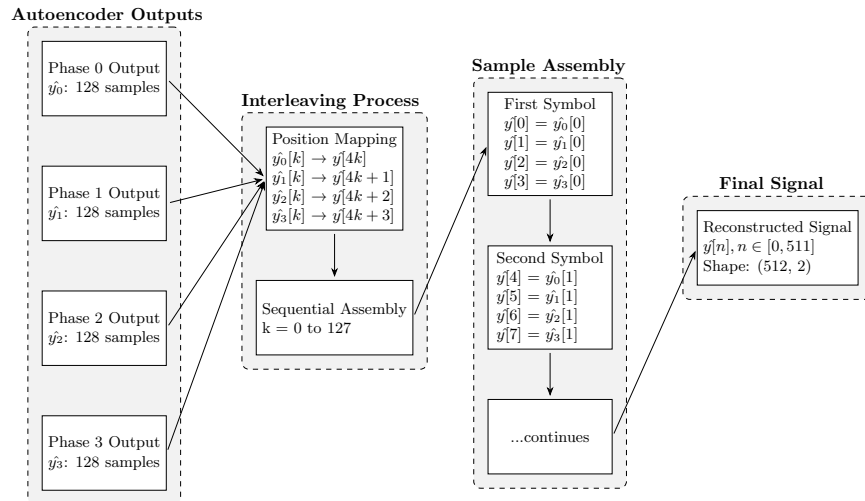
<sup>1</sup>The % symbol describes the modulo operator that returns the remainder after division.



(a) Windowing of QPSK signal dataset.



(b) Cyclostationary phase synchronous dataset creation.



(c) Interleaving model outputs for signal reconstruction.

Figure 5.3: Cyclostationary analysis pipeline from preprocessing to reconstruction.

generates four position-specific datasets with dimensions:

$$\mathbf{X}_m \in \mathbb{R}^{(N/4) \times N_{win} \times 2}, \quad \mathbf{y}_m \in \mathbb{R}^{(N/4) \times 2}, \quad (5.11)$$

where  $N/4$  denotes the number of symbols and also represents the dataset length dimension.

During training, these datasets are partitioned into mini-batches of size  $B$ , structured as:

$$\mathbf{X}_m^{\text{batch}} \in \mathbb{R}^{B \times N_{win} \times 2}, \quad \mathbf{y}_m^{\text{batch}} \in \mathbb{R}^{B \times 2}. \quad (5.12)$$

This batching strategy preserves phase consistency across samples while enabling parallelized gradient computations.

### Intuition behind two-step process

The windowing operation constitutes a critical two-stage process that synergistically preserves both localized temporal context and cyclostationary phase relationships within QPSK signals. In the first stage, overlapping windows extract finite-length segments of the raw signal, each spanning a symmetric neighborhood around a target sample. This localized field of view retains short-term correlations between adjacent symbols—essential for identifying transient interference patterns that manifest across multiple consecutive samples. The second stage reorganizes these windows into phase-synchronous subgroups, aligning their central elements with specific positions within the QPSK symbol structure. This dual mechanism ensures that each processed sample remains embedded within its native temporal context while simultaneously adhering to the periodic statistical regularities inherent to the modulation scheme. Isolated phase synchronization—merely decimating the signal into non-overlapping phase-specific subsequences—would discard the interstitial information between symbol phases, effectively obscuring interference signatures that span symbol

boundaries or exhibit time-varying characteristics. We explore this thought in the results section with training configuration that disregards phase synchronous cyclostationary sub-dataset creation.

### **Window Size selection**

An odd window size is generally the preferred choice as it ensures a well-defined central sample in each window, directly aligning with the clean reference signal  $y$ . This centralization provides a symmetrical view of the surrounding neighborhood, preserving phase consistency and making the learned relationships more interpretable. By maintaining a direct correspondence between windowed samples and their clean counterparts, odd-sized windows offer a more structured and intuitive representation of the signal, making them suitable for most cases.

An even window size, in contrast, lacks a clear center, making it harder to establish direct alignment with  $y$ . This can introduce ambiguity in interpretation. However, in specific cases where the window size is deliberately chosen as a multiple or function of a key system parameter shared by both the signal of interest (SOI) and the signal not of interest (SNOI), an even-sized window may reveal structured patterns or periodic behaviors that are otherwise less apparent.

While odd windows provide a more balanced and interpretable view, even windows may be useful when system-specific design considerations take precedence. The choice depends on whether direct sample alignment or structural integration with system parameters is the primary objective.



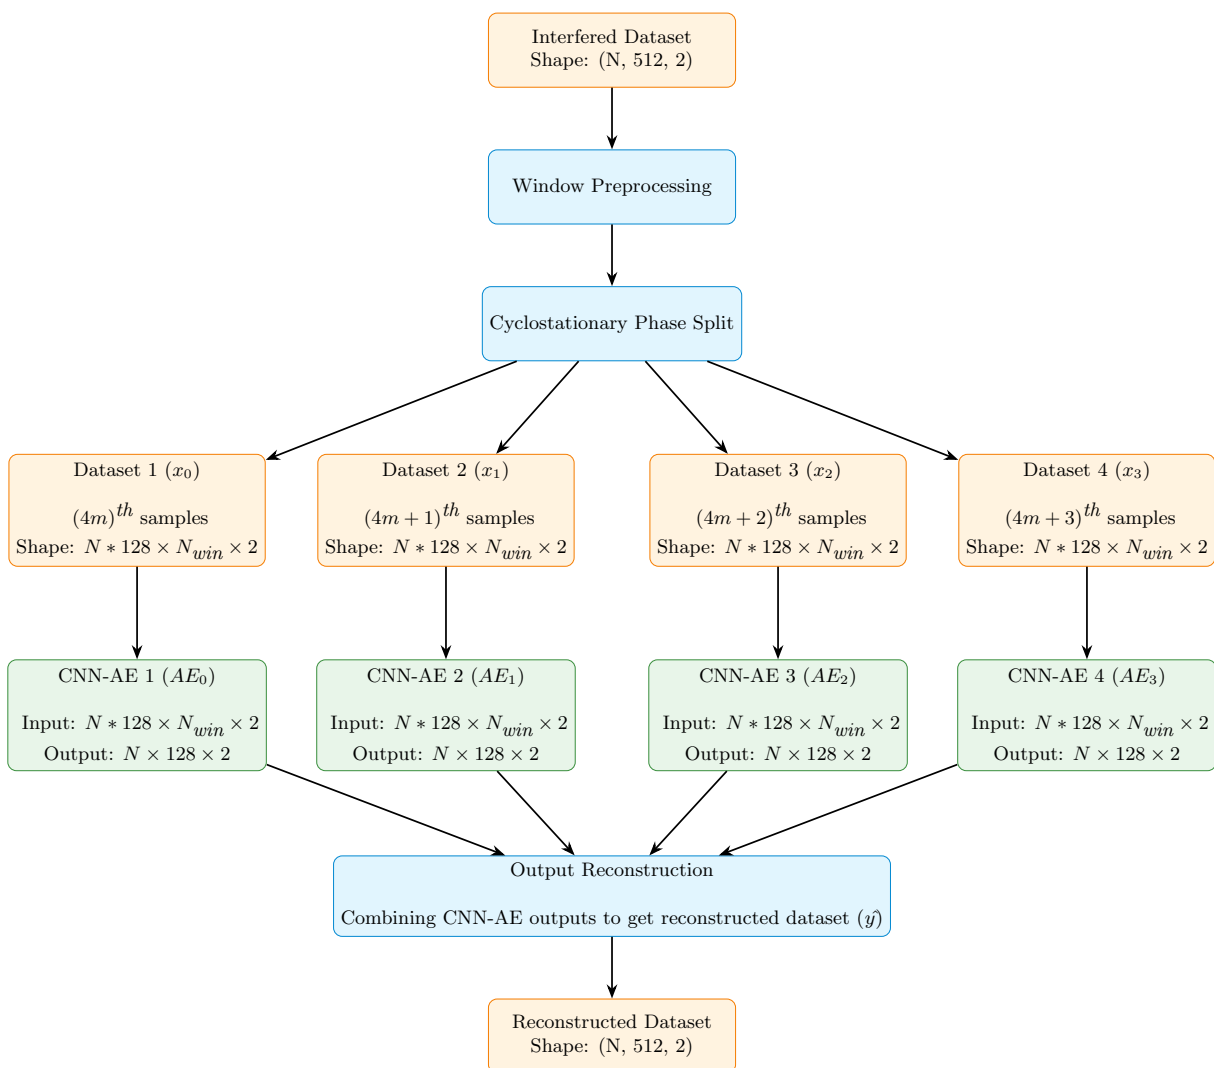


Figure 5.4: Cyclostationary Training Pipeline

### 5.4.2 Model Architecture and Training Setup

The high level architecture of our approach is described in Figure 5.4. We follow the same QPSK dataset generation steps as previous chapters and then we split the 512 length signals into 4 windowed datasets. The CNN-AE model architecture described in previous chapters is used to create four models that are trained on each of the 4 phase synchronous subsets of the original dataset.

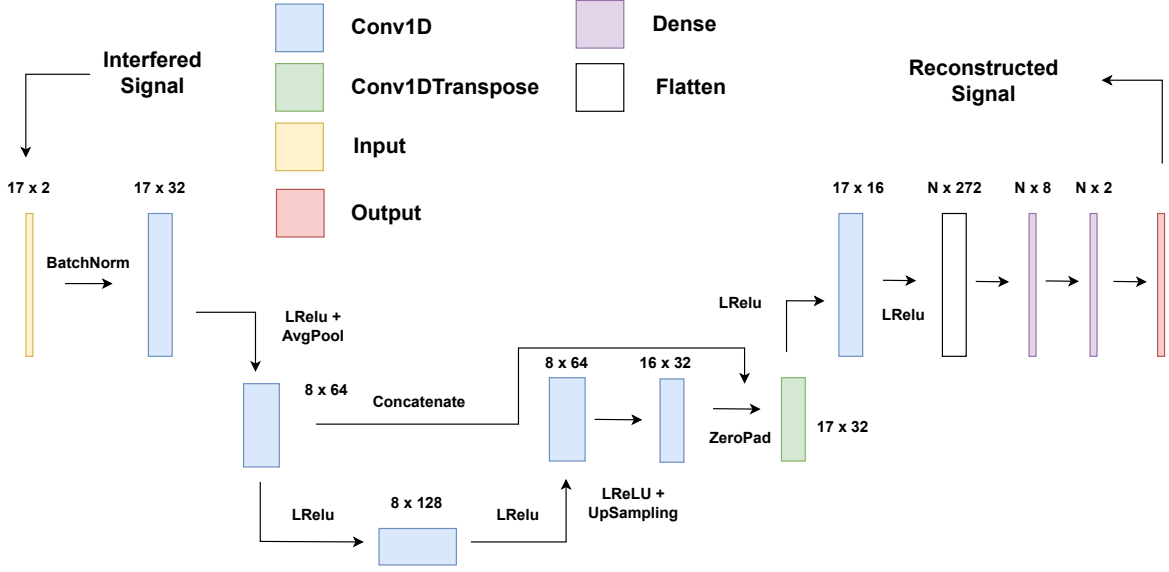


Figure 5.5: Small CNN-AE Architecture

The training process consists of four separate autoencoders, each corresponding to a phase-split subset of the signal. Given the phase-split windowed training sets  $x_m^{\text{train}}$  and  $y_m^{\text{train}}$  for  $m = 0, 1, 2, 3$ , each autoencoder  $\text{AE}_m$  is trained to minimize the Mean Squared Error (MSE) loss:

$$\mathcal{L}_m = \frac{1}{N} \sum_{i=1}^N \|\text{AE}_m(x_{m,i}^{\text{train}}) - y_{m,i}^{\text{train}}\|^2, \quad m \in \{0, 1, 2, 3\} \quad (5.13)$$

where  $\mathcal{L}_m$  represents the MSE loss for the  $m$ -th autoencoder, and  $N$  is the number of training samples.

During inference, the trained autoencoders  $\text{AE}_m$  process the windowed subsets of the test set  $x_m^{\text{test}}$ , generating reconstructed outputs:

$$\hat{y}_m^{\text{test}} = \text{AE}_m(x_m^{\text{test}}) \quad (5.14)$$

where  $\hat{y}_m^{\text{test}}$  represents the denoised estimate of the signal for each phase component.

Once the outputs  $\hat{y}_m^{\text{test}}$  are obtained, the inverse of the windowing operation  $W^{-1}$  is applied to merge the phase components into the final reconstructed signal:

$$\hat{y} = W^{-1} (\hat{y}_0^{\text{test}}, \hat{y}_1^{\text{test}}, \hat{y}_2^{\text{test}}, \hat{y}_3^{\text{test}}) \quad (5.15)$$

where  $W^{-1}$  represents the inverse transformation that reconstructs the full sequence by interleaving the phase-aligned outputs back to the original tensor shape.

Thus, the final estimate  $\hat{y}$  is obtained by reversing the windowing process, ensuring that the reconstructed signal maintains the original temporal structure.

The study investigates four distinct combinations of CNN-AE architectures and training methodologies to evaluate the role of cyclostationary signal properties in suppressing interference for QPSK systems. Each configuration is designed to isolate specific aspects of dataset partitioning, model specialization, and architectural complexity, enabling a granular analysis of their impact on suppression performance.

**Standard Dataset (1 CNN-AE)** serves as the baseline architecture, operating on the full-length input signal where cyclostationary properties must be learned directly from data. The autoencoder convolutional neural network (CNN-AE) processes the raw interfered signal in its entirety, with an input tensor of shape  $(N, 512, 2)$ , where  $N$  denotes the batch size, 512 represents the total upsampled samples per sequence, and 2 corresponds to the in-phase and quadrature (I/Q) components. This end-to-end approach treats the signal as a monolithic block, applying global feature extraction and reconstruction. It provides a performance benchmark against which cyclostationary methods are compared, highlighting the trade-offs between holistic signal processing and phase-specific decomposition.

**4 Windowed Datasets (4 CNN-AE)** exploits cyclostationarity by partitioning the input into four phase-aligned subsets as described in the preprocessing subsection. Due to the upsampling factor of 4, the signal is divided into sequences corresponding to the  $(4m + k)$ -th samples, where  $m \in \{1, 2, 3, 0\}$  indexes the phase and  $k$  corresponds to the samples in the signal. The 512 length QPSK dataset is first processed into overlapping windows, and then split into 4 phase subsets. Four independent CNN-AE models are trained on these subsets, each to specialize in denoising signals from a specific cyclostationary phase. During reconstruction, the outputs of these models are strategically interleaved, aligning with the original phase partitioning to synthesize the full 512-sample signal. This configuration tests the hypothesis that phase-specific desired signal patterns can be more effectively preserved.

**4 Windowed Datasets (4 Small CNN-AE)** retains the dataset partitioning strategy of the previous setup but employs reduced-complexity models to evaluate computational efficiency. Each model is a vastly simplified version of the standard CNN-AE architecture, creating a “lightweight” variant which is called “Small CNN-AE” for ease of reference. An illustration of the Small CNN-AE architecture is shown in Figure 5.5. The input dimensions remain identical, but the diminished model capacity explores whether cyclostationary gains are robust to architectural constraints. This configuration addresses practical considerations, such as deployment in resource-constrained environments, by quantifying the performance penalty incurred when using smaller models.

**1 Windowed Dataset (1 CNN-AE)** applies windowing without phase-based partitioning to isolate the impact of localized processing. The full signal is segmented into overlapping windows. A single CNN-AE model processes all windows uniformly, ignoring cyclostationary phase alignment. Reconstruction involves concatenating output windows, mimicking the baseline approach but operating on truncated signal segments. This setup distinguishes between the benefits of general window-based reconstruction and those uniquely attributable

to cyclostationary phase specialization, revealing whether localized processing alone suffices for interference suppression.

### **Window Size selection for different Interference scenarios**

For the CommSignal2 interference scenario, a window size of 17 (input shape—  $(N, 17, 2)$ ) was selected to provide symmetrical context around a central sample, enabling direct alignment between the windowed input and the clean reference signal. This odd-sized window facilitates more intuitive feature learning by maintaining phase consistency across the window. In contrast, for QPSK interference scenarios, a window size of 16 was deliberately chosen to match the underlying signal structure (input shape—  $(N, 16, 2)$ ), as both the signal of interest and interfering QPSK signals share an upsampling factor of 4.

Collectively, these configurations reveal the interplay between signal segmentation strategies, model specialization, and architectural complexity. The phased approaches (second and third configurations) explicitly leverage cyclostationarity by aligning models with periodic signal properties, while the non-phased setups (first and fourth) provide critical baselines for isolating these effects and provide an absolute floor in terms of performance with all samples passed at the same time and a singular windowed input shape model. The results quantify how these design choices influence key metrics such as EVM, SINR, and BER, providing insight into the value of exploiting cyclostationary properties of the signal. All other training parameters including number of epochs, learning rate optimization are kept the same as previous chapters.

## 5.5 Experimental Evaluation and Discussion

### 5.5.1 Interference Suppression Performance

The efficacy of the proposed windowed processing methodology was rigorously evaluated under three distinct interference scenarios: standard CommSignal2 interference, QPSK interference, and delayed QPSK interference. Quantitative performance metrics, including Error Vector Magnitude (EVM), Signal-to-Interference-plus-Noise Ratio (SINR), and Bit Error Rate (BER), are consolidated in Tables 5.1–5.3 where the character 'W' in brackets indicates training performed on a windowed dataset<sup>2</sup>. These results demonstrate consistent trends across experimental conditions, validating the interplay between input dimensionality, model architecture, and interference suppression capability.

Table 5.1: CommSignal2 Interference Suppression Performance (Window Size - 17)

Model Type	Input Shape	EVM (dB)	SINR (dB)	BER
Standard CNN-AE	(512,2)	-20.19	20.03	0.0042
4 CNN-AE (W)	(17,2)	-16.11	16.04	0.0114
4 Small CNN-AE (W)	(17,2)	-14.45	14.37	0.0203
1 CNN-AE (W)	(17,2)	-12.47	12.45	0.0798

The baseline model, denoted as Standard CNN-AE, processes the full 512-sample signal and achieves the best performance metrics, as anticipated. Its expansive input window facilitates comprehensive temporal context utilization, enabling precise identification of interference patterns. For instance, under CommSignal2 interference (Table 5.1), the baseline attains a BER of 0.0042, representing an empirical upper bound for monolithic architectures. In contrast, the phase-split four-model architecture (4 CNN-AE) attains a SINR of 16.04 dB, corresponding to 80% of the baseline’s performance, despite utilizing only 3.3% of the

<sup>2</sup>4 model scenarios with (W) indicates cyclostationary phase-splitting into subdatasets, each processed by a dedicated model. A 1 model scenario with (W) refers to windowed dataset training without cyclostationary phase splitting.

input samples. This reduction in input dimensionality underscores the inherent efficiency of localized processing while retaining critical interference suppression functionality.

Table 5.2: QPSK Interference Suppression Performance (Window Size - 16)

Model Type	Input Shape	EVM (dB)	SINR (dB)	BER
Standard CNN-AE	(512,2)	-40.26	38.48	5.28e-5
4 CNN-AE (W)	(16,2)	-35.69	34.89	7.83e-5
4 Small CNN-AE (W)	(16,2)	-27.00	27.22	0.0005
1 CNN-AE (W)	(16,2)	-26.86	26.64	0.0009

Under QPSK interference (Table 5.2), the cyclostationary phase-split architecture narrows the performance gap relative to the baseline, achieving 90% of the reference SINR. This improvement is attributed to the alignment between the windowed processing framework and the cyclostationary properties of the QPSK signal. Each submodel specializes in a distinct phase of the 4-sample symbol period, enabling the extraction of phase-specific interference features that monolithic architectures fail to resolve. The reduced-complexity variant (4 Small) retains 71% of the baseline’s SINR, indicating that parameter efficiency can be achieved without significant performance degradation when paired with phase specialization.

### Robustness to Timing Mismatches

Table 5.3: QPSK Interference (2 sample delay) Suppression Performance (Window Size - 16)

Model Type	Input Shape	EVM (dB)	SINR (dB)	BER
Standard CNN-AE	(512,2)	-41.62	38.73	4.53e-5
4 CNN-AE (W)	(16,2)	-34.99	34.29	9.64e-5
4 Small CNN-AE (W)	(16,2)	-27.46	27.24	0.00048
1 CNN-AE (W)	(16,2)	-24.70	24.70	0.0010

The delayed interference scenario (Table 5.3) evaluates architectural robustness to timing misalignments. The phase-split 4 CNN-AE (W) system maintains 88% of its nominal SINR under a 2-sample delay, compared to 94% retention in the baseline. This resilience is at-

tributed to the localized nature of windowed processing, where submodels learn interference features within constrained symbol phases, reducing sensitivity to global timing shifts. Conversely, the unified windowed model (1 CNN-AE) exhibits a 19% SINR degradation under identical conditions, emphasizing the necessity of phase partitioning for timing-agnostic operation.

### 5.5.2 Computational Efficiency and Real-Time Viability

The baseline architecture’s superior accuracy can be argued to be counterbalanced by the computational complexity of working with a much larger input frame (analyzed in detail in Section 5.5.5). The task of processing 512 samples necessitates significantly greater arithmetic operations compared to the 16 or 17-sample windows employed in phase-split models. The 4 CNN-AE (W) system reduces input dimensionality by more than 90%, enabling parallel execution across submodels. This parallelism permits concurrent processing of four symbol phases, which could potentially reduce inference latency by 75% relative to sequential frameworks. Even the reduced-capacity 4 Small CNN-AE variant achieves comparable latency reductions, demonstrating viability for edge computing applications with stringent power and timing constraints.


### 5.5.3 Impact of Windowing strategy

The unified windowed model i.e. 1 CNN-AE (W), occupies an intermediate position between the baseline and phase-split architectures. While its SINR of 24.70 dB under delayed sample interference lags behind its cyclostationary phase-split counterparts, it represents a significant improvement from the -10 dB SIR of the interfered signal dataset. This indicates that windowed processing inherently captures salient local context, although without resolv-

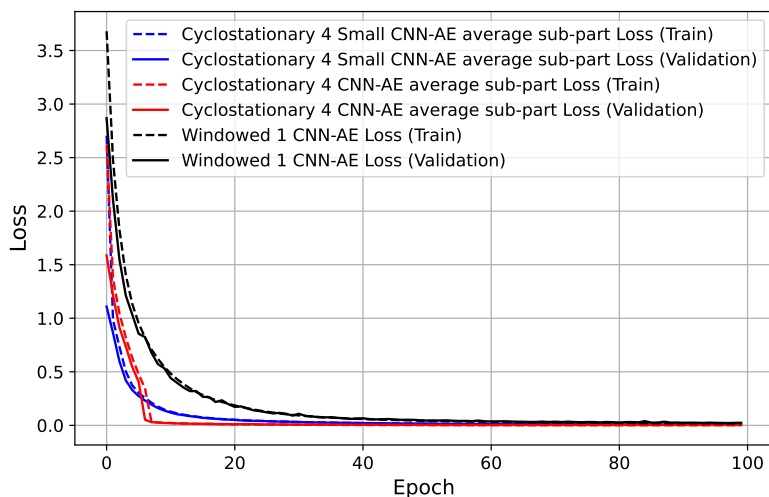


ing phase-specific interference patterns. Its architectural simplicity may benefit applications where moderate accuracy is permissible, but the phase-split paradigm remains indispensable for high-fidelity suppression.

### 5.5.4 Loss Curve Convergence Analysis



(a) QPSK SOI with *CommSignal2* Interference



(b) QPSK SOI with QPSK Interference

Figure 5.6: Loss Curve convergence comparison by interference type.

The convergence behavior of cyclostationary phase-aligned models versus non-phase-aware windowed processing reveals fundamental differences in learning efficiency across interference scenarios. Figure 5.6 presents a comparison plot of training and validation loss trajectories, measured using Mean Squared Error (MSE), for both CommSignal2 and QPSK interference conditions, illustrating the distinct advantages of phase-aligned processing.

In each plot, six curves are displayed: the average training loss (dashed red) and average validation loss (solid red) of the 4 CNN-AE (W) configuration trained on cyclostationary phase-aligned subdatasets, the average training loss (dashed blue) and average validation loss (solid blue) of the 4 Small CNN-AE (W) configuration trained on cyclostationary phase-aligned subdatasets, the training loss (dashed black) and validation loss (solid black) of a single CNN-AE model trained on a windowed dataset without phase alignment. Crucially, all three approaches maintain identical input tensor dimensions— $(N, 17, 2)$  for CommSignal2 interference and  $(N, 16, 2)$  for QPSK interference—ensuring a fair comparison of architectural efficacy independent of input dimensionality effects.

For the CommSignal2 interference case (Figure 5.6a), the cyclostationary 4 CNN-AE (W) architecture demonstrates superior convergence properties, reaching a validation loss of approximately 0.05 by epoch 20, while the windowed CNN-AE requires nearly 60 epochs to achieve a similar loss. Notably, the 4 Small CNN-AE (W) variant, despite having substantially fewer parameters, converges more rapidly than the windowed 1 CNN-AE architecture and approaches the 4 CNN-AE’s convergence rate, although with a slightly higher asymptotic loss value.

The QPSK interference scenario (Figure 5.6b) reveals more pronounced disparities between modeling approaches. While all configurations exhibit faster convergence compared to CommSignal2 interference—reflecting the structured nature of the QPSK interferer—the cyclostationary 4 CNN-AE achieves a final validation loss of 0.0037, representing a 85%

reduction compared to the windowed CNN-AE’s 0.0237. Most strikingly, the computational efficiency advantages of the cyclostationary approach become evident when comparing the Small CNN-AE to the windowed architecture, which once again noticeably outperforms it.

We observe that despite utilizing a small fraction the parameters of the 4 CNN-AE (W) configuration, the phase-synchronized 4 Small CNN-AE (W) outperforms the substantially larger windowed model in final validation loss for QPSK interference. This finding conclusively demonstrates that architectural complexity can be judiciously reduced when the processing framework intrinsically aligns with the signal’s statistical structure. These convergence patterns validate the dual-stage preprocessing rationale presented in Section 5.4.1. The windowing operation alone provides valuable localized context, but without cyclostationary phase synchronization, it fails to fully exploit the inherent periodicity in the signal’s statistical properties. The unified windowed model struggles to simultaneously learn interference patterns across all four phase positions, whereas the cyclostationary phase-synchronized approach partitions the learning task into four specialized subproblems. This decomposition not only accelerates training but also improves generalization, as evidenced by the consistently lower validation losses across both interference scenarios.

### 5.5.5 Model Complexity Analysis

Table 5.4: Model Architecture Complexity Comparison

Architecture	No. of Models	Parameters/Model	Total Parameters
Standard CNN-AE	1	3,167,466	3,167,466
4 CNN-AE	4	3,176,066	12,704,264
4 Small CNN-AE	4	116,050	464,200
1 CNN-AE (Windowed)	1	3,176,066	3,176,066

The complexity distribution across these architectures reveals an intriguing trade-off between computational burden and performance efficacy. The Standard CNN-AE, processing 512-

length inputs through approximately 3.17M parameters, achieves the highest SINR (38.48 dB for QPSK interference) while maintaining architectural simplicity with a single model. In contrast, the 4 CNN-AE approach quadruples the total parameter count to 12.7M, distributed across four models of 3.18M parameters each, yet delivers 90% of the baseline SINR (34.89 dB) while processing only 3.1% of the input samples per model.

The 4 Small CNN-AE configuration presents perhaps the most compelling efficiency narrative, reducing per-model complexity by 96% to just 116K parameters. Despite this dramatic reduction, it maintains 71% of baseline SINR performance (27.22 dB), suggesting that phase-specific processing can partially compensate for reduced model capacity. The single windowed CNN-AE, while matching the parameter count of its distributed counterparts at 3.18M, achieves only 64% of baseline SINR (24.70 dB), indicating that the performance advantages of windowed processing are intrinsically linked to phase specialization rather than mere input segmentation.

These relationships between model complexity and performance metrics suggest that cyclostationary phase-specialized processing can achieve substantial efficiency gains without proportional degradation in interference suppression capability. For instance, the 4 Small CNN-AE architecture reduces total parameters by 85% compared to the baseline while sacrificing only 29% of SINR performance, representing an advantageous complexity-performance trade-off for resource-constrained applications.

## 5.6 Conclusion

This investigation suggests that cyclostationary-based windowing may offer a useful balance between interference suppression and computational efficiency, making it a viable approach for latency-constrained communication systems where periodic signal properties can be lever-

aged. Three main takeaways emerge:

1. Phase-specific modeling is fundamental to high-performance interference suppression. The 4 CNN-AE architecture attains 80–90% of the baseline’s SINR while utilizing 97% fewer input samples, a feat enabled by submodel alignment with the QPSK signal’s cyclostationary properties. Each submodel operates as a dedicated estimator for a quarter of the symbol period, resolving phase-dependent interference features that monolithic architectures inherently ignore.
2. The convergence analysis provides compelling evidence for the efficacy of phase synchronized processing, with our cyclostationary approach demonstrating significantly faster and more stable convergence trajectories. Notably, the 4 Small CNN-AE (W) configuration achieved better convergence characteristics than the larger windowed CNN-AE despite its dramatically reduced parameter count, confirming that architectural alignment with signal structure offers greater performance benefits than mere model complexity. This suggests that communication systems can achieve robust interference suppression with minimal computational overhead when models are designed to exploit inherent statistical periodicities.
3. Architectural specialization mitigates the limitations of reduced model capacity. The 4 Small CNN-AE variant retains 71–78% of the baseline’s standard CNN-AE’s performance while reducing total parameters by 85% (from 3.17M to 464K parameters), demonstrating that constrained models can achieve efficiency without complete accuracy loss when focused on narrow signal segments. This finding is very relevant for resource-constrained deployments, such as IoT edge devices.
4. The parallelizable nature of phase-split cyclostationary processing can potentially address real-time operational requirements by splitting the training duration amongst

multiple smaller submodels. By distributing computations across four specialized submodels, inference latency is reduced proportionally to the number of phases. This could be a potential advancement for 5G NR and next-generation wireless systems that demand microsecond-scale processing.

Future research directions could investigate hybrid architectures that integrate global signal context with phase-localized processing. Such frameworks may combine a lightweight model analyzing full-signal trends with phase-specific submodels, potentially bridging the performance gap between monolithic and windowed approaches. This methodology not only advances interference suppression but also provides a foundational framework for future development towards deep learning in latency-sensitive applications across wireless communications.

# Chapter 6

## Conclusion

Wireless communication systems face unprecedented challenges as networks grow denser and spectrum becomes more crowded. This thesis addresses these challenges by combining deep learning with classical signal processing principles, developing new methods to suppress interference in dynamic environments. The work demonstrates how adaptive neural networks can complement traditional techniques, creating systems that automatically learn to separate signals from noise—even when interference patterns change unpredictably. Through extensive experiments and theoretical analysis, we establish a foundation for building smarter receivers that maintain reliability in complex real-world conditions.

### 6.1 Interference Suppression using CNN-AE

The core innovation of this research lies in the data driven deep learning approach of the convolutional autoencoder (CNN-AE) architecture, which consistently outperformed conventional approaches across all tested scenarios. When challenged with three major interference types—in-band tones, out-of-band tones, and real-world wideband interference types, the model achieved:

- 15–20 dB improvements in signal-to-interference-plus-noise ratio (SINR)
- 92% reduction in signal distortion (EVM)

- 100x lower bit error rates compared to matched filtering

Additionally, our generalization study revealed remarkable capabilities of the Multi-SIR training approach:

- 8-12 dB SINR improvements over SIR-matched models for tone interference
- Nearly an order of magnitude reduction in BER across varying SIR conditions in tone interference.
- Competitive performance with CommSignal2 interference across all SIR levels using a single model

These results prove that neural networks can learn universal strategies for interference suppression, adapting to both narrowband and broadband disturbances without prior knowledge of their characteristics. The Multi-SIR training approach further demonstrates that a single model can effectively handle diverse interference scenarios, eliminating the need for multiple specialized implementations. The latent space analysis provided critical insights, showing how the network internally organizes different interference types into distinct clusters—a finding that bridges machine learning “black boxes” with interpretable signal processing concepts.

## 6.2 Cyclostationarity - Theory and Practice

A notable contribution emerged from integrating cyclostationarity—the inherent periodic patterns in digital signals—into the deep learning framework. By aligning the neural network’s processing windows with the QPSK symbol structure, we created a system that understands the natural rhythm of communication signals. This strategic blend of domain



knowledge and data-driven learning allowed the model to suppress interference with metrics close to the data driven single model approach of CNN-AE using smaller input sizes for each model and with a simpler DL architecture. The successful application of cyclostationary properties demonstrates how wireless practitioners can effectively incorporate deep learning into real-world designs, combining theoretical principles with modern computational tools.

### 6.3 Practical Impact and Considerations

The techniques developed here directly address pressing needs in 5G/6G networks and IoT deployments, where devices must operate reliably despite crowded spectrum conditions. CNN-AE's capabilities of maintaining stable connections at signal strengths 10 dB lower than traditional methods could be used to extend network coverage in urban areas. However, two practical constraints need attention: first, the computational cost of neural networks remains higher than classic filters, though hardware advancements are rapidly closing this gap; second, while the model handles QPSK exceptionally well, its performance with advanced modulations like 64-QAM requires further study.

### 6.4 Future Directions

The demonstrated success of deep learning in interference suppression opens concrete pathways to enhance both theoretical understanding and practical deployment. Immediate priorities include transitioning the CNN-AE framework from simulation to real-world hardware through joint optimization of algorithms and computing architectures. This requires co-designing neural networks with FPGAs and low-power system-on-chip platforms, ensuring compatibility with 5G's strict timing constraints while maintaining the model's suppression

capabilities. Simultaneously, expanding the system’s adaptability to handle advanced modulation schemes will demand architectures that dynamically adjust their processing based on real-time signal conditions, potentially through lightweight meta-learning approaches.

A potential area for future investigation lies in determining the robustness limitations of the cyclostationary approach under practical timing impairments. The current framework assumes perfect timing synchronization between transmitter and receiver, maintaining the integrity of phase-aligned components. However, real-world communication systems frequently encounter symbol timing offsets, clock drift, and sample loss—conditions that can directly challenge the cyclostationary phase coherence assumption, which our method depends upon. Future research could investigate scenarios of performance degradation as timing errors increase and explore adaptive synchronization mechanisms that are able to realign phase components on-the-fly. Potential solutions might include incorporating elastic buffering techniques or designing a front-end estimation network that detects and corrects timing misalignments before signals reach the phase-aligned autoencoder models. Understanding these boundaries will clarify when to employ cyclostationary methods versus more timing-agnostic approaches. Broadly speaking, these considerations explore edge case scenarios to test the limits of the cyclostationary approach towards its breaking point.

A critical evolution lies in unifying digital suppression with analog front-end design. Current models assume idealized signal conversion, but practical receivers suffer from amplifier nonlinearities and mixer imperfections that distort interference patterns. Developing end-to-end trainable models that encompass both analog and digital domains could automatically compensate for these hardware limitations, potentially boosting receiver sensitivity by 10–15 dB. This integration must extend to network-level coordination, where base stations collaboratively share interference fingerprints through privacy-preserving federated learning—enabling collective intelligence without raw data exchange.

Finally, the accelerating convergence of terrestrial and non-terrestrial networks demands protection against emerging interference types. Satellite-to-device communications, drone swarms, and ultra-wideband radar systems create novel disturbance profiles that defy conventional spectral management. Adapting the framework to handle time-varying Doppler shifts and spatially distributed interferers will require embedding physics-based constraints into neural architectures—for instance, using recurrent networks aligned with orbital mechanics. By pursuing these directions, future research can transform deep learning from a promising tool into a foundational component of interference-resilient communication systems.

This work shows that deep learning and classical signal processing need not compete—they can create synergistic solutions when carefully integrated. By grounding neural networks in communication theory principles like cyclostationarity, we achieve systems that are both adaptive and interpretable. As wireless technology advances toward 6G and ubiquitous IoT, such hybrid approaches will become essential to efficiently manage the spectrum. The path forward lies not in replacing decades of engineering wisdom, but in enhancing it with intelligent algorithms that learn from the wireless environment itself. This thesis takes concrete steps toward that future, providing tools and insights to build communication systems that thrive in complexity rather than merely survive it.

# Bibliography

- [1] H. Kokkinen, A. Piemontese, L. Kulacz, F. Arnal, and C. Amatetti, “Coverage and interference in co-channel spectrum sharing between terrestrial and satellite networks,” in *2023 IEEE Aerospace Conference*, pp. 1–9, 2023.
- [2] J. Andrews, “Interference cancellation for cellular systems: a contemporary overview,” *IEEE Wireless Communications*, vol. 12, no. 2, pp. 19–29, 2005.
- [3] C.-S. Tsang and T. Nguyen, “Survey of signal processing techniques for interference suppression in communication,” in *2004 IEEE Aerospace Conference Proceedings (IEEE Cat. No.04TH8720)*, vol. 2, pp. 1347–1354 Vol.2, 2004.
- [4] M. Aygur, S. Kandeepan, A. Giorgetti, A. Al-Hourani, E. Arbon, and M. Bowyer, “Narrowband interference mitigation techniques: A survey,” *IEEE Communications Surveys Tutorials*, pp. 1–1, 2025.
- [5] S. Neemat, O. Krasnov, and A. Yarovoy, “An interference mitigation technique for fmcw radar using beat-frequencies interpolation in the stft domain,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 67, no. 3, pp. 1207–1220, 2019.
- [6] T. Shimura, M. Umehira, Y. Watanabe, X. Wang, and S. Takeda, “An advanced wide-band interference suppression technique using envelope detection and sorting for automotive fmcw radar,” in *2022 IEEE Radar Conference (RadarConf22)*, pp. 1–6, 2022.
- [7] N. Azmi, L. M. Kamarudin, M. Mahmuddin, A. Zakaria, A. Y. M. Shakaff, S. Khatun, K. Kamarudin, and M. N. Morshed, “Interference issues and mitigation method in wsn

- 2.4ghz ism band: A survey,” in *2014 2nd International Conference on Electronic Design (ICED)*, pp. 403–408, 2014.
- [8] K. R. Chowdhury and I. F. Akyildiz, “Interferer classification, channel selection and transmission adaptation for wireless sensor networks,” in *2009 IEEE International Conference on Communications*, pp. 1–5, 2009.
- [9] T. T. Le and S. Moh, “Interference mitigation schemes for wireless body area sensor networks: A comparative survey,” *Sensors*, vol. 15, no. 6, pp. 13805–13838, 2015.
- [10] Q. F. D. Lu and R. Wu, “Survey on interference mitigation via adaptive array processing in gps,” *PIERS Online*, vol. vol. 2, no. no. 4, p. pp. 357362, 2006.
- [11] J. E. Swindell, C. Slater, S. Hussey, C. Baylis, and R. J. Marks, “Assessing interference with regression analysis techniques,” in *2024 IEEE Texas Symposium on Wireless and Microwave Circuits and Systems (WMCS)*, pp. 1–6, 2024.
- [12] R. Zhang and S. Cao, “Support vector machines for classification of automotive radar interference,” in *2018 IEEE Radar Conference (RadarConf18)*, pp. 0366–0371, 2018.
- [13] S. Grimaldi, A. Mahmood, and M. Gidlund, “An svm-based method for classification of external interference in industrial wireless sensor and actuator networks,” *Journal of Sensor and Actuator Networks*, vol. 6, no. 2, 2017.
- [14] X. Lu, W. Su, J. Yang, H. Gu, H. Zhang, W. Yu, and T. S. Yeo, “Radio frequency interference suppression for sar via block sparse bayesian learning,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 12, pp. 4835–4847, 2018.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*

- (F. Pereira, C. Burges, L. Bottou, and K. Weinberger, eds.), vol. 25, Curran Associates, Inc., 2012.
- [16] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [17] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Comput.*, vol. 9, p. 1735–1780, Nov. 1997.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [19] T. J. O’Shea, J. Corgan, and T. C. Clancy, “Convolutional radio modulation recognition networks,” in *Engineering Applications of Neural Networks* (C. Jayne and L. Iliadis, eds.), (Cham), pp. 213–226, Springer International Publishing, 2016.
- [20] N. E. West and T. O’Shea, “Deep architectures for modulation recognition,” in *2017 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pp. 1–6, 2017.
- [21] 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, 2018.
- [22] H. Yan, G. Wang, F. Gao, and W. Yuan, “Dynamic convolutional neural network for wireless interference identification,” *IEEE Communications Letters*, vol. 28, no. 1, pp. 83–87, 2024.
- [23] G. Baldini, F. Bonavitacola, and J.-M. Chareau, “Wireless interference identification with convolutional neural networks based on the fpga implementation of the lte cell-

- specific reference signal (crs),” *IEEE Transactions on Cognitive Communications and Networking*, vol. 10, no. 1, pp. 48–63, 2024.
- [24] J. Kim, S. Lee, Y.-H. Kim, and S.-C. Kim, “Classification of interference signal for automotive radar systems with convolutional neural network,” *IEEE Access*, vol. 8, pp. 176717–176727, 2020.
- [25] T. Oyedare, V. K. Shah, D. J. Jakubisin, and J. H. Reed, “Interference suppression using deep learning: Current approaches and open challenges,” *IEEE Access*, vol. 10, pp. 66238–66266, 2022.
- [26] J. Xu, B. Ai, and T. Q. S. Quek, “Toward interference suppression: Ris-aided high-speed railway networks via deep reinforcement learning,” *IEEE Transactions on Wireless Communications*, vol. 22, no. 6, pp. 4188–4201, 2023.
- [27] Y. Chen, H. Shan, and H. Wei, “Research on communication signal interference suppression based on deep learning,” in *2022 International Conference on Artificial Intelligence, Information Processing and Cloud Computing (AIIPCC)*, pp. 451–454, 2022.
- [28] H. Guo, S. Wu, H. Wang, and M. Daneshmand, “Dsic: Deep learning based self-interference cancellation for in-band full duplex wireless,” in *2019 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, 2019.
- [29] N. Shlezinger, R. Fu, and Y. C. Eldar, “Deepsoft: Deep soft interference cancellation for multiuser mimo detection,” *IEEE Transactions on Wireless Communications*, vol. 20, no. 2, pp. 1349–1362, 2021.
- [30] A. Fuchs, J. Rock, M. Toth, P. Meissner, and F. Pernkopf, “Complex-valued convolutional neural networks for enhanced radar signal denoising and interference mitigation,” in *2021 IEEE Radar Conference (RadarConf21)*, pp. 1–6, 2021.

- [31] M. Kamruzzaman, N. I. Sarkar, and J. Gutierrez, "Machine learning-based resource allocation algorithm to mitigate interference in d2d-enabled cellular networks," *Future Internet*, vol. 16, no. 11, 2024.
- [32] W. Zheng, J. Yao, and K. Wu, "Mitigating cross-technology interference in heterogeneous wireless networks based on deep learning," in *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 230–237, 2020.
- [33] M. Mosavi and F. Shafiee, "Narrowband interference suppression for gps navigation using neural networks," *GPS Solutions*, vol. 20, no. 3, pp. 341–351, 2016.
- [34] W. L. Mao, "Gps interference mitigation using derivative-free kalman filter-based rnn," *Radioengineering*, vol. 25, pp. 518–526, Sept. 2016.
- [35] M. A. S. Sejan, M. H. Rahman, M. A. Aziz, R. Tabassum, Y.-H. You, D.-D. Hwang, and H.-K. Song, "Interference management for a wireless communication network using a recurrent neural network approach," *Mathematics*, vol. 12, no. 11, 2024.
- [36] H. Sun, X. Chen, Q. Shi, M. Hong, X. Fu, and N. D. Sidiropoulos, "Learning to optimize: Training deep neural networks for interference management," *IEEE Transactions on Signal Processing*, vol. 66, no. 20, pp. 5438–5453, 2018.
- [37] B. Song, H. Sun, W. Pu, S. Liu, and M. Hong, "To supervise or not to supervise: How to effectively learn wireless interference management models?," in *2021 IEEE 22nd International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 211–215, 2021.
- [38] T. Erpek, Y. E. Sagduyu, and Y. Shi, "Deep learning for launching and mitigating wireless jamming attacks," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 1, pp. 2–14, 2019.



- [39] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563–575, 2017.
- [40] T. Erpek, T. J. O'Shea, and T. C. Clancy, "Learning a physical layer scheme for the mimo interference channel," in *2018 IEEE International Conference on Communications (ICC)*, pp. 1–5, 2018.
- [41] K.-C. Wang, K.-C. Liu, S.-Y. Peng, and Y. Tsao, "Ecg artifact removal from single-channel surface emg using fully convolutional networks," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023.
- [42] L. Gondara, "Medical image denoising using convolutional denoising autoencoders," in *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pp. 241–246, 2016.
- [43] E. Almazrouei, G. Gianini, C. Mio, N. Almoosa, and E. Damiani, "Using autoencoders for radio signal denoising," in *Proceedings of the 15th ACM International Symposium on QoS and Security for Wireless and Mobile Networks, Q2SWinet'19*, (New York, NY, USA), p. 11–17, Association for Computing Machinery, 2019.
- [44] C.-L. Lin, K. C.-J. Lin, C.-C. Lee, and Y. Tsao, "Cross-technology interference mitigation using fully convolutional denoising autoencoders," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, pp. 1–6, 2020.
- [45] S. Chen, J. Taghia, T. Fei, U. Kühnau, N. Pohl, and R. Martin, "A dnn autoencoder for automotive radar interference mitigation," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4065–4069, 2021.

- [46] S. B. Brown, S. Young, A. Wagenknecht, D. Jakubisin, C. E. Thornton, A. Orndorff, and W. C. Headley, "Aircraft radar altimeter interference mitigation through a cnn-layer only denoising autoencoder architecture," in *MILCOM 2024 - 2024 IEEE Military Communications Conference (MILCOM)*, pp. 306–311, 2024.
- [47] T. Erpek, T. J. O'Shea, and T. C. Clancy, "Learning a physical layer scheme for the mimo interference channel," in *2018 IEEE International Conference on Communications (ICC)*, pp. 1–5, 2018.
- [48] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, eds.), (Cham), pp. 234–241, Springer International Publishing, 2015.
- [49] X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P.-A. Heng, "H-denseunet: Hybrid densely connected unet for liver and tumor segmentation from ct volumes," *IEEE Transactions on Medical Imaging*, vol. 37, no. 12, pp. 2663–2674, 2018.
- [50] N. Siddique, S. Paheding, C. P. Elkin, and V. Devabhaktuni, "U-net and its variants for medical image segmentation: A review of theory and applications," *IEEE Access*, vol. 9, pp. 82031–82057, 2021.
- [51] X. Hu, M. A. Naiel, A. Wong, M. Lamm, and P. Fieguth, "Runet: A robust unet architecture for image super-resolution," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 505–507, 2019.
- [52] D. Liang, H. Zhang, X. Zhou, D. Li, D. Yuan, and M. Zhang, "Step-unet: Prediction of moving and communication behaviors of vehicles," in *2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall)*, pp. 1–6, 2021.

- [53] H. Kang, Q. Hu, H. Chen, Q. Huang, Q. Zhang, and M. Cheng, “Cross-shaped separated spatial-temporal unet transformer for accurate channel prediction,” in *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*, pp. 2079–2088, 2024.
- [54] W. A. Gardner, “The spectral correlation theory of cyclostationary time-series,” *Signal Processing*, vol. 11, no. 1, pp. 13–36, 1986.
- [55] W. Gardner, “Exploitation of spectral redundancy in cyclostationary signals,” *IEEE Signal Processing Magazine*, vol. 8, no. 2, pp. 14–36, 1991.
- [56] K. Kim, I. A. Akbar, K. K. Bae, J.-S. Um, C. M. Spooner, and J. H. Reed, “Cyclostationary approaches to signal detection and classification in cognitive radio,” in *2007 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pp. 212–215, 2007.
- [57] R. S. Lynch, E. T. Smith, and M. Harrison, “Radio frequency interference excision using cyclostationary signal processing,” *The RFI2022 Workshop at ECMWF*.
- [58] J. Zhang, X. Cui, H. Xu, and M. Lu, “A two-stage interference suppression scheme based on antenna array for gnss jamming and spoofing,” *Sensors*, vol. 19, no. 18, 2019.
- [59] G. Cucho-Padin, Y. Wang, E. Li, L. Waldrop, Z. Tian, F. Kamalabadi, and P. Perillat, “Radio frequency interference detection and mitigation using compressive statistical sensing,” *Radio Science*, vol. 54, no. 11, pp. 986–1001, 2019.
- [60] P. D. Sutton, K. E. Nolan, and L. E. Doyle, “Cyclostationary signatures in practical cognitive radio applications,” *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 13–24, 2008.
- [61] RFChallenge@MIT, “Icassp 2024 sp grand challenge: Data-driven signal separation in radio spectrum.” <https://rfchallenge.mit.edu/icassp24-single-channel/>.

- [62] L. van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008.
- [63] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [64] C. Spooner, “Cyclostationary signal processing.” <https://cyclostationary.blog/>.