Deep Learning for Enhanced Free-space Optical Communications

M.Tech. Thesis

By

KARTHIK TADIGADAPA



DEPARTMENT OF ELECTRICAL ENGINEERING

INDIAN INSTITUTE OF TECHNOLOGY INDORE

May 2026

Deep Learning for Enhanced Free-space Optical Communications

Thesis

Submitted in partial fulfillment of the requirements for the award of the degree of

Master of Technology

by

KARTHIK TADIGADAPA



DEPARTMENT OF ELECTRICAL ENGINEERING

INDIAN INSTITUTE OF TECHNOLOGY INDORE

MAY 2026

INDIAN INSTITUTE OF TECHNOLOGY INDORE

CANDIDATE'S DECLARATION

I hereby certify that the work which is being presented in the thesis entitled **Deep Learning for Enhanced Free-space Optical Communications** in the partial fulfillment of the requirements for the award of the degree of **Master of Technology - Communication and Signal Processing** and submitted in the Department of Electrical Engineering, Indian Institute of Technology Indore, is an authentic record of my own work carried out during the time period from July 2025 to June 2026 under the supervision of Dr. Prabhat K. Upadhyay, Professor, Indian Institute of Technology Indore, Indore, India.

The matter presented in this thesis has not been submitted by me for the award of any other degree of this or any other institute.

Karthik Tadigadapa 15-05-2026

Signature of the Student with Date

(Karthik Tadigadapa)

This is to certify that the above statement made by the candidate is correct to the best of my knowledge.

15-5-2026

Signature of the Supervisor of M.Tech Thesis with Date

(Prof. Prabhat K. Upadhyay)

Karthik Tadigadapa has successfully given his M.Tech Oral Examination held on **27th April 2026**

Saptarshi Ghosh

Signature of Supervisor of M.Tech Thesis

Signature of Convener, DPGC

Date: 15-5-2026

Date: 15-05-2026

ACKNOWLEDGEMENTS

I would like to begin by expressing my deepest gratitude to my thesis supervisor, **Prof. Prabhat K. Upadhyay**, for his unwavering guidance, patience, and encouragement throughout this work. His deep expertise in wireless communications and his ability to make complex problems feel approachable have shaped not just this thesis but my entire approach to research. I am genuinely grateful for the time he invested in our discussions, for pushing me to think more rigorously. This work would not exist without his mentorship.

I owe a very special thank-you to **Shivaji Cheepurupalli**, PhD student under Prof. Upadhyay. Shivaji's work on UAV-assisted FSO/RF NOMA systems was a direct inspiration for this thesis, and his willingness to share insights, discuss ideas, and offer feedback made my research journey considerably smoother. Working alongside someone who is both technically sharp and genuinely generous with their knowledge has been one of the most rewarding parts of this experience. I hope our paths continue to cross in future research.

I would also like to thank the Department of Electrical Engineering, Indian Institute of Technology Indore, for providing the computational resources and academic environment that made this work possible.

Finally, I am deeply grateful to my family for their unconditional support. Their patience and encouragement through the long hours of this work meant more than words can express.

This thesis is dedicated to
all the people who have supported me
throughout this journey

ABSTRACT

Free-space optical (FSO) communication offers licence-free, high-bandwidth connectivity ideal for sixth-generation backhaul networks, but atmospheric turbulence and beam misalignment cause frequent deep fades that threaten link reliability. Adding a parallel radio-frequency (RF) sub-link and applying adaptive maximal-ratio combining at the relay significantly improves availability, while non-orthogonal multiple access (NOMA) on the second hop allows simultaneous service to multiple users within the same bandwidth. Despite the practical importance of this architecture, no prior work has characterised its exact bit-error rate (BER) and outage performance in closed form, nor provided a computationally efficient surrogate for the resulting Meijer-G expressions.

This thesis derives exact closed-form BER and outage expressions for a dual-hop decode-and-forward relay system with a hybrid FSO/RF adaptive combining first hop and a two-user QPSK-NOMA second hop. The FSO sub-link is modelled under Gamma-Gamma turbulence with Farid-Hranilovic pointing errors; the RF sub-link follows Nakagami-m fading. A high-SNR asymptotic analysis via Slater's theorem yields the diversity order $d = \min(\xi_F^2/2,\, \alpha_F/2,\, \beta_F/2)$ and the coding gain in closed form. For moderate turbulence ($\alpha_F = 5.614$, $\beta_F = 4.006$), the diversity order is $d = 2.003$; under strong turbulence ($\beta_F = 1.207$) it drops to $d = 0.603$, producing power penalties of 25.4 dB in BER and 21.3 dB in outage at 10^{-3}. Optimum NOMA power allocation is also studied numerically.

Because evaluating the Meijer-G series requires 100-digit variable-precision arithmetic and takes roughly two minutes per curve, a feedforward deep neural network (DNN) with three ReLU hidden layers is trained on 100 000 analytically generated samples to serve as a fast surrogate. The DNN achieves an RMSE of 2.94×10^{-4} (moderate turbulence) with a 24,000$\times$ speedup. Comparison against XGBoost and Random Forest baselines shows that all three models are accurate in-distribution ($R^2 > 0.985$), but XG-

Boost develops a structural spike near 25 dB under strong turbulence — an artifact inherent to piecewise-constant tree prediction — while the DNN predicts smoothly throughout, making it the recommended surrogate for this application.

Contents

Abstract **i**

List of Figures **iv**

List of Tables **vii**

List of Nomenclature & Acronyms **vii**

1 Introduction **1**

1.1 Background and Motivation 1

1.2 Free-Space Optical and Hybrid FSO/RF Communications 4

1.3 Non-Orthogonal Multiple Access 5

1.4 Deep Learning for Wireless Communications 7

1.5 Research Gaps and Objectives 9

1.6 Contributions of This Thesis 10

1.7 Organization of the Thesis 12

2 Literature Survey **16**

2.1 Overview 16

2.2 Free-Space Optical Channel Modelling 17

2.3 Hybrid FSO/RF Systems and Combining Strategies . 19

2.4 Non-Orthogonal Multiple Access in Relay Systems . . 21

2.5 Deep Learning and Machine Learning for Wireless Systems 23

2.6 Summary and Identified Research Gaps 27

3 System Model **29**

3.1 System Description 30

3.2 FSO Channel Model 31
3.3 RF Channel Model 32
3.4 NOMA Second-Hop Signal Model 32
3.5 Performance Metrics 34
3.6 Simulation Parameters 34

4 Performance Analysis **36**
4.1 Overview 36
4.2 CDF of the Combined First-Hop SNR 36
4.3 First-Hop Outage Probability 38
4.4 Average BER of the First Hop 38
4.5 Average BER of the NOMA Second Hop 40
4.5.1 User 1 BER (No SIC) 40
4.5.2 User 2 BER (With SIC) 42
4.6 Outage Probability of the NOMA Second Hop 44
4.7 Power Allocation 46
4.7.1 Minimum Average BER 46
4.7.2 Fairness-Based Power Allocation 46
4.8 Closed-Form Results Summary 47

5 Asymptotic Analysis **49**
5.1 Motivation and Overview 49
5.2 Slater's Theorem and the Dominant Pole 50
5.3 Diversity Order 51
5.4 Coding Gain Analysis 52
5.5 Asymptotic Outage and BER Expressions 53
5.6 Numerical Results and Discussion 53

6 Deep Learning **56**
6.1 Overview 56
6.2 Dataset Generation 57
6.3 DNN Architecture and Training 58
6.4 Tree-Based Baseline Models 59
6.5 Extrapolation Analysis 60
6.6 Summary 61

7 Results and Discussion **62**
7.1 Overview 62
7.2 DNN-Predicted Outage Probability 62
7.3 DNN-Predicted Average BER 67
7.4 Asymptotic Analysis: Exact vs High-SNR Approximation 70
7.5 DNN vs XGBoost vs Random Forest vs Theoretical . 73

8 Conclusion and Future Work **76**
8.1 Conclusion 76
8.2 Future Work 77

List of Figures

3.1 Dual-hop hybrid FSO/RF–NOMA relay system. The source (S) transmits to the UAV-mounted relay (R) via a parallel FSO link (red beam) and RF link (dashed). The relay applies adaptive combining and retransmits to NOMA users via the second hop. . . . 30

6.1 DNN architecture: 5 input features, three hidden layers of 128 ReLU neurons, one linear output neuron predicting $\log_{10}(\bar{P}_{e,\mathrm{AC}})$. Total: 34 177 trainable parameters. 58

7.1 Outage probability $\mathcal{P}_{\mathrm{out}}$ vs. $\bar{\gamma}_F$ for the **moderate turbulence** scenario ($\alpha_F = 5.6137$, $\beta_F = 4.0062$, $\xi_F = 4.5747$). Blue $\times$: exact analytical (MATLAB `vpa`). Orange $\circ$: DNN prediction. Parameters: $\bar{\gamma}_R = 8\,\mathrm{dB}$, $\gamma_{th} = 10\,\mathrm{dB}$, $\gamma_{\mathrm{out}} = 3\,\mathrm{dB}$. 63

7.2 Outage probability $\mathcal{P}_{\mathrm{out}}$ vs. $\bar{\gamma}_F$ for the **strong turbulence** scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction. A systematic downward offset between the predicted and theoretical curves is visible, indicating that the DNN slightly under-predicts the outage under strong turbulence conditions. 63

7.3 Scatter plot of DNN-predicted outage probability vs. exact analytical outage probability (linear scale) on the held-out test set — **moderate turbulence**. Points cluster tightly along the perfect-match diagonal, with slight dispersion visible in the intermediate range (0.1–0.3). 65

7.4 Scatter plot of DNN-predicted outage probability vs. exact analytical outage probability (linear scale) on the held-out test set — **strong turbulence**. The scatter is wider compared to the moderate case, particularly at intermediate outage values (0.15–0.35), consistent with the systematic under-prediction observed in Figure 7.2. 66

7.5 Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ for the **moderate turbulence** scenario ($\alpha_F = 5.6137$, $\beta_F = 4.0062$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction (RMSE $= 2.94 \times 10^{-4}$, $R^2 = 0.9981$). Parameters: $\bar{\gamma}_R = 5\,\mathrm{dB}$, $\gamma_{th} = 7\,\mathrm{dB}$, BPSK. 67

7.6 Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ for the **strong turbulence** scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction (RMSE $= 5.60 \times 10^{-4}$, $R^2 = 0.9960$). Parameters as in Figure 7.5. 68

7.7 Scatter plot of DNN-predicted BER vs. exact analytical BER (linear scale, range [0, 0.08]) on the held-out test set — **moderate turbulence**. Points cluster along the perfect-match diagonal with slight spread above BER ≈ 0.05. 69

7.8 Scatter plot of DNN-predicted BER vs. exact analytical BER (linear scale, range [0, 0.08]) on the held-out test set — **strong turbulence**. The scatter pattern is similar to the moderate case; slight widening is visible at BER values above 0.05. 69

7.9 Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ over 0–60 dB: exact analytical (solid curves with markers) and high-SNR asymptotic approximation $G_c^{-1}\bar{\gamma}^{-d}$ (dashed curves). Moderate turbulence ($d = 2.003$, $G_c = 24.93$) decays to $\sim 10^{-11}$ by 60 dB; strong turbulence ($d = 0.603$, $G_c = 0.072$) decays to $\sim 3 \times 10^{-4}$ at the same SNR. The asymptotic approximation for strong turbulence merges with the exact curve at approximately 5 dB, while the moderate-turbulence asymptote converges near 30 dB. 71

7.10 Outage probability $\mathcal{P}_{\mathrm{out}}$ vs. $\bar{\gamma}$ over 0–60 dB: exact analytical (solid curves with markers) and asymptotic approximations (dashed curves). Moderate turbulence reaches $\sim 10^{-11}$ at 60 dB; strong turbulence reaches $\sim 10^{-4}$, a gap of seven decades at the same transmit power. The asymptote merging behaviour mirrors the BER results. 72

7.11 Average BER vs. SNR for the **strong turbulence** scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$): exact analytical (black dashed), DNN (blue $\circ$, RMSE $= 5.60 \times 10^{-4}$), XGBoost (red $\square$, RMSE $= 1.09 \times 10^{-3}$), and Random Forest (green $\triangle$, RMSE $= 5.74 \times 10^{-4}$). A sharp XGBoost spike at approximately 25–26 dB is clearly visible. 74

List of Tables

3.1 Simulation parameters for moderate and strong turbulence scenarios. 35

4.1 Closed-form BER coefficients for $\beta_1 = 0.8$, $\beta_2 = 0.2$. . 48

5.1 Asymptotic parameters for moderate and strong turbulence scenarios. 54

6.1 Input feature ranges for dataset generation. 57
6.2 Surrogate model comparison on in-distribution test set. 59

7.1 Asymptotic performance parameters verified from Figures 7.9 and 7.10. 73
7.2 Surrogate model accuracy on the held-out test set for both turbulence scenarios. 75

List of Nomenclature & Acronyms

Nomenclature

α_F, β_F	–	Large- and small-scale scintillation parameters of the Gamma-Gamma turbulence distribution
ξ_F	–	Pointing-error parameter (ratio of equivalent beam radius to twice the jitter standard deviation)
m	–	Nakagami-m fading parameter for the RF sub-link
γ_F, γ_R	–	Instantaneous SNR of the FSO and RF sub-links
$\bar{\gamma}_F$, $\bar{\gamma}_R$	–	Mean SNR of the FSO and RF sub-links (dB)
γ_c	–	Combined first-hop SNR under adaptive combining
γ_{th}	–	Adaptive combining switching threshold
γ_{out}	–	Outage threshold SNR
$F_{\gamma_c}(\cdot)$	–	CDF of the combined first-hop SNR
$\bar{P}_{e,\text{AC}}$	–	Average BER of the adaptive combining first hop
$\mathcal{P}_{\text{out}}$	–	Outage probability of the first hop
d	–	Diversity order
G_c	–	Coding gain
h_n	–	Complex fading coefficient of NOMA user n
α_n	–	Channel gain magnitude of NOMA user n
β_1, β_2	–	NOMA power allocation coefficients ($\beta_1 + \beta_2 = 1$)
δ_1, ψ_2	–	NOMA outage threshold parameters for User 1 and User 2
$a_{n,c}$	–	Squared QPSK amplitude factor for user n, term c

$G_{p,q}^{m,n}[\cdot]$	–	Meijer-G function
N_0	–	AWGN noise power spectral density
$Q(\cdot)$	–	Gaussian Q-function
erfc$(\cdot)$	–	Complementary error function
$\Gamma(\cdot)$	–	Euler Gamma function
$K_\nu(\cdot)$	–	Modified Bessel function of the second kind, order
I_1, I_2	–	First and second integral contributions to $\bar{P}_{e,\mathrm{AC}}$
R^2	–	Coefficient of determination

Acronyms

AC	Adaptive Combining
AWGN	Additive White Gaussian Noise
BER	Bit Error Rate
BPSK	Binary Phase-Shift Keying
CDF	Cumulative Distribution Function
CSI	Channel State Information
DF	Decode-and-Forward
DNN	Deep Neural Network
FSO	Free-Space Optical
GG	Gamma-Gamma (turbulence model)
IM/DD	Intensity Modulation with Direct Detection
MATLAB	Matrix Laboratory
MRC	Maximal-Ratio Combining
NOMA	Non-Orthogonal Multiple Access
OFDM	Orthogonal Frequency-Division Multiplexing
PDF	Probability Density Function
PE	Pointing Errors
QPSK	Quadrature Phase-Shift Keying
RF	Radio Frequency
ReLU	Rectified Linear Unit
RMSE	Root Mean Square Error
SIC	Successive Interference Cancellation

SINR	Signal-to-Interference-plus-Noise Ratio
SNR	Signal-to-Noise Ratio
UAV	Unmanned Aerial Vehicle
VPA	Variable-Precision Arithmetic
XGBoost	eXtreme Gradient Boosting
5G	Fifth Generation (mobile communications)
6G	Sixth Generation (mobile communications)

Chapter 1

Introduction

1.1 Background and Motivation

The past two decades have witnessed an unprecedented transformation in the way people communicate, work, and interact with the physical world. The proliferation of smartphones, wearable devices, autonomous vehicles, smart city infrastructure, and immersive extended-reality platforms has pushed global wireless data traffic to levels that would have seemed implausible at the dawn of 4G networks. Mobile data consumption has been growing at a compound annual rate exceeding thirty percent, and industry forecasts project a further thousand-fold increase in aggregate traffic between 2020 and 2030. Fifth-generation networks, now in commercial deployment across major urban centres worldwide, represent a significant leap in peak data rate, latency, and device density compared with their predecessors. Yet the fundamental spectral constraints of radio-frequency communications mean that 5G infrastructure, built predominantly on licensed millimetre-wave bands, is not architected to sustain the traffic volumes anticipated for the latter half of this decade.

This widening gap between traffic demand and network capacity has become one of the defining engineering challenges of the modern era, and it has catalysed global research into sixth-generation wireless networks. The 6G vision, which is expected to be standardised around 2030, is ambitious in every dimension:

terabit-per-second peak data rates, sub-millisecond end-to-end latency, near-perfect coverage spanning terrestrial base stations, high-altitude platforms, and low-Earth-orbit satellite constellations, and the ability to connect billions of heterogeneous devices simultaneously. Meeting these objectives requires not only incremental improvements to existing radio technologies but the incorporation of entirely new physical-layer paradigms that can deliver orders-of-magnitude gains in throughput, reliability, and spectral efficiency. Three technology pillars have emerged as central to the 6G programme: free-space optical communications, non-orthogonal multiple access, and artificial intelligence-driven network management. The research presented in this thesis sits at the intersection of all three.

Free-space optical communication transmits information by modulating a narrow laser beam through the open atmosphere between a transmitter and a receiver equipped with a collecting aperture. The technology has several intrinsic advantages that make it exceptionally well suited to the 6G backhaul and access role. FSO links operate in the optical spectrum, which is largely unregulated, eliminating the enormous cost and logistical complexity associated with acquiring licensed radio-frequency spectrum. The narrow beam divergence of a laser means that FSO links are inherently resistant to passive eavesdropping, since intercepting the beam without detection requires physically obstructing the line of sight. Perhaps most importantly, the optical carrier frequency is several orders of magnitude higher than the highest licensed millimetre-wave bands currently used in 5G, which translates directly into a vastly larger available modulation bandwidth and thus the potential for data rates in the tens of gigabits per second on a single wavelength. These properties have attracted sustained attention from both industry and academia as a means to relieve backhaul congestion in dense urban deployments, to bridge the last mile in regions where fibre infrastructure is absent, and to provide high-throughput feeder links for satellite and aerial platforms.

The principal weakness of a standalone FSO link, however, is

its sensitivity to atmospheric conditions. The atmosphere is never truly homogeneous: temperature gradients, wind-driven turbulence cells, and aerosol scattering introduce random fluctuations in the refractive index of air along the beam path. These fluctuations distort the optical wavefront and cause the received signal intensity to vary rapidly and unpredictably — a phenomenon known as atmospheric scintillation or turbulence-induced fading. Under strong turbulence conditions, the received power can drop by several orders of magnitude below its clear-sky value for sustained intervals, driving the bit-error rate of the link to unacceptable levels. Separately, mechanical vibrations of the transmitter platform and the effects of atmospheric wind loads cause the beam centroid to wander away from the receiver aperture, an impairment known as pointing error or misalignment fading. Together, turbulence and pointing errors represent the dominant reliability challenge for FSO systems, and mitigating their combined impact is the central physical-layer problem that motivates the system architecture studied in this thesis.

A further and more severe threat to FSO availability is dense fog, which can increase the optical attenuation along a propagation path by several hundred decibels per kilometre, effectively blocking the link entirely for intervals that may last hours. No amount of increased transmit power can overcome this level of attenuation within a practical link budget. This hard availability ceiling distinguishes FSO from RF technologies, for which rain, the primary precipitation-related impairment at centimetre wavelengths, causes attenuation that is severe but rarely link-breaking over short distances. The complementary nature of the two impairment mechanisms — fog attenuates optical links far more than RF, while heavy rain is the primary RF impairment and has comparatively little effect on optical propagation — provides the physical rationale for combining the two technologies in a single hybrid system.

1.2 Free-Space Optical and Hybrid FSO/RF Communications

The concept of hybrid FSO/RF communication stems from the recognition that the two technologies, deployed in parallel on the same path, offer a form of spatial and spectral diversity against weather-induced outages that neither can provide alone. When the FSO channel is clear, the system benefits from the high throughput and spectral efficiency of the optical link. When fog, dense cloud cover, or extreme atmospheric turbulence degrades the FSO channel below a usable threshold, the RF link maintains connectivity, ensuring that a temporary failure of one medium does not translate into a service interruption visible to the end user. The availability of a hybrid FSO/RF link over a full calendar year can approach five-nines reliability in many climatic zones where a standalone FSO link would fall well short of that target.

Within the family of hybrid FSO/RF architectures, the choice of combining strategy has a significant impact on both performance and implementation complexity. The simplest approach, hard switching, uses only the FSO link during favourable conditions and transfers all traffic to the RF link when the optical channel degrades below a threshold. While straightforward to implement, hard switching discards the received FSO signal during the transition, wasting the residual signal energy that, even during a partial fade, may still carry useful information. A more efficient alternative is adaptive combining, in which the FSO link is always active and the RF link is additionally engaged and combined with the FSO signal whenever the optical channel quality falls below the switching threshold. This approach, which is the combining strategy adopted in this thesis, applies maximal-ratio combining to the two received signals during periods of FSO degradation, recovering a degree of diversity gain that hard switching cannot achieve. The cost is a modest increase in receiver complexity and the need for a low-overhead feedback channel to signal the combining decision to the transmitter.

The relay-assisted topology provides a further dimension of system improvement. For long-distance FSO links, the received optical power falls with the square of the propagation distance, and turbulence strength accumulates along the entire path. By introducing a relay node at an intermediate point, the path is divided into two shorter hops, each experiencing less severe turbulence and higher received power than the end-to-end single-hop alternative. When the relay operates in decode-and-forward mode, it fully demodulates the received signal and retransmits a clean regenerated copy, preventing noise and fading from the first hop from propagating into the second. The dual-hop decode-and-forward architecture thus provides a natural framework for combining the high-capacity FSO/RF first hop with additional processing at the relay — including the multi-user NOMA transmission that forms the second hop in the system studied here.

1.3 Non-Orthogonal Multiple Access

Alongside the capacity challenge posed by explosive traffic growth, next-generation networks face an equally pressing connectivity challenge: the need to serve a vastly larger number of devices simultaneously within a limited spectral resource. Industrial IoT deployments, smart grid monitoring, autonomous vehicle coordination, and dense machine-type communications all require the network to maintain simultaneous connections to hundreds or thousands of devices within a single cell sector. The orthogonal multiple access schemes that underpin 4G and the radio-access portion of 5G — time-division, frequency-division, and orthogonal frequency-division multiple access — allocate each user an exclusive resource block, which limits the number of simultaneously served devices to the number of available orthogonal slots. As the device population grows, this hard resource constraint becomes an insurmountable bottleneck.

Non-orthogonal multiple access resolves this bottleneck by deliberately allowing multiple users to share the same time-frequency

resource block simultaneously, exploiting power-domain differences to separate them at the receiver. At the transmitter, the signals of all users are superposed in a single waveform, with each user assigned a different fraction of the total transmit power according to its channel quality. Users with weak channel conditions — those at the cell edge or experiencing deep fading — receive a larger share of the power so that their signals remain decodable despite the channel disadvantage. Users with strong channel conditions receive a smaller power share, since their favourable channel compensates for the reduced power allocation. At the receiver, successive interference cancellation enables the stronger user to first decode and subtract the weaker user's signal before decoding its own, achieving in principle an interference-free detection for the stronger user. The weaker user decodes its own signal directly, treating the stronger user's contribution as an additive interference term whose impact is bounded by the power allocation design.

The spectral efficiency gains of NOMA over OMA are most pronounced when the users being multiplexed have significantly different channel conditions, which is precisely the scenario encountered in a downlink relay setting where one user is geographically close to the relay and the other is near the cell edge. In such configurations, NOMA can serve both users simultaneously on the same resource block, effectively doubling the number of connections relative to OMA without any increase in bandwidth. The challenge in system design lies in choosing the power allocation between users optimally, balancing the competing objectives of maximising average throughput, minimising average bit-error rate, and ensuring a fair quality of service for both users. These power allocation problems, formulated and solved for the NOMA second hop of the relay system in this thesis, are among the practically most important design degrees of freedom in any NOMA deployment.

When NOMA is integrated into the second hop of a hybrid FSO/RF relay system, the combined architecture achieves simultaneously high link reliability on the backhaul segment and high spectral efficiency on the access segment — precisely the twin objectives

that 6G heterogeneous networks demand. The interaction between the first-hop channel statistics and the NOMA second-hop performance is non-trivial: the end-to-end system behaviour depends on the joint distribution of the FSO, RF, and NOMA user channels, and a rigorous analytical treatment requires tracking this dependence through the decode-and-forward relay processing. Providing this rigorous treatment is one of the principal analytical contributions of this thesis.

1.4 Deep Learning for Wireless Communications

The third pillar of the 6G programme is the integration of artificial intelligence and machine learning into the physical and network layers of wireless systems. The motivations are diverse: AI can accelerate the solution of optimisation problems that are too complex for real-time analytical methods, can learn channel models from data without requiring explicit statistical assumptions, and can serve as a computationally efficient surrogate for analytical performance expressions that are prohibitively expensive to evaluate at scale. It is this third role — surrogate modelling for analytical performance metrics — that is central to the deep learning contribution of this thesis.

The performance metrics of hybrid FSO/RF systems, specifically the average bit-error rate and the outage probability, are expressed in closed form through special mathematical functions known as Meijer-G functions. These functions are generalisations of many classical special functions and arise naturally from the Mellin-Barnes integral representations that appear when the channel PDF is expressed as a series expansion. While Meijer-G functions are mathematically exact and provide a complete characterisation of the system performance, their numerical evaluation is computationally demanding. The series representations converge slowly in certain parameter regimes, and the terms of the series can be of alternating sign and large magnitude, leading to catastrophic numerical can-

cellation when computed in standard double-precision arithmetic. Accurate evaluation therefore requires variable-precision arithmetic at high digit counts, implemented in tools such as MATLAB's symbolic toolbox, which imposes evaluation times on the order of several minutes for a single BER curve. This computational bottleneck makes it impractical to use the analytical expressions directly in iterative design loops, Monte Carlo validation studies, or real-time network management applications.

Deep neural networks offer a compelling solution to this computational bottleneck. A DNN trained on a dataset of input-output pairs generated from the analytical expressions learns to reproduce the mapping from channel parameters and SNR to BER with high fidelity, encoding the relationship into a fixed set of learned weights that can be evaluated in microseconds. Once trained, the DNN requires no symbolic computation and no special-function libraries; it executes as a sequence of matrix multiplications and element-wise nonlinearities, operations that are natively supported on any modern processor or GPU. The training cost is incurred once and amortised over the lifetime of the model, which, provided the underlying channel model does not change, remains valid indefinitely. The resulting speedup of several tens of thousands relative to the MATLAB analytical baseline makes the DNN surrogate practical for the iterative and real-time applications where the analytical expressions cannot be used directly.

The suitability of neural networks as surrogates for wireless performance metrics has been demonstrated in a growing body of recent literature. Studies on energy-harvesting cooperative networks, UAV-assisted relay systems, and reconfigurable intelligent surface channels have all shown that relatively compact feedforward networks can achieve coefficient-of-determination values exceeding 99% relative to the analytical or simulated ground truth. A consistent finding across these studies is that the log-transform of the BER should be used as the prediction target, converting the many-decades dynamic range of the BER into a roughly linear scale that is far better conditioned for gradient-based training. This thesis

adopts these design principles and additionally provides a systematic comparison between the DNN and tree-based ensemble methods — XGBoost and Random Forest — to characterise the relative strengths of each approach and to document the extrapolation failure of tree-based models at SNR values outside the training range.

1.5 Research Gaps and Objectives

Despite the substantial body of literature on hybrid FSO/RF relaying, NOMA, and deep learning for wireless systems, several analytical and computational gaps remain unaddressed in the existing work. First, most studies on hybrid FSO/RF systems that derive closed-form performance expressions assume a hard-switching combining strategy, in which the receiver uses either the FSO link or the RF link but never both simultaneously. The adaptive combining scheme, in which the two signals are maximal-ratio combined during FSO fades, offers superior performance but requires a fundamentally different analytical treatment because the combined SNR is no longer a simple switch between two distributions but the sum of two correlated and differently-distributed random variables. Closed-form average BER and outage probability for this adaptive combining scheme, expressed through the Meijer-G framework, have not been reported in the existing literature for the Gamma-Gamma turbulence model with pointing errors.

Second, the integration of a hybrid FSO/RF first hop with a NOMA-enabled decode-and-forward second hop has not been characterised analytically in the existing literature. Studies that combine FSO relaying with NOMA typically assume that the first hop is a pure FSO link or a pure RF link, and do not address the composite case where the first hop itself involves adaptive combining of two different physical channels. The interaction between the first-hop combined SNR distribution and the NOMA second-hop BER, mediated by the decode-and-forward relay, introduces analytical dependencies that must be tracked carefully through the derivation.

Third, the asymptotic high-SNR behaviour of the composite hy-

brid FSO/RF–NOMA system — which reveals the diversity order and coding gain that govern the slope and offset of the BER and outage curves at large SNR — has not been derived for the adaptive combining architecture. The diversity order is of central engineering importance because it determines how steeply performance improves with increasing transmit power, and therefore how much additional link budget is required to meet a target BER in the strong turbulence regime relative to moderate turbulence.

Fourth, the computational bottleneck of Meijer-G evaluation for hybrid FSO/RF systems with pointing errors has not been addressed through data-driven surrogate modelling. While neural network surrogates have been applied to simpler FSO-only and RF-only NOMA systems in the recent literature, their application to the composite hybrid adaptive combining architecture, and specifically the characterisation of their extrapolation behaviour against tree-based baselines, represents an open problem.

The objective of this thesis is to address each of these four gaps through a unified analytical and computational study of the dual-hop hybrid FSO/RF–NOMA relay system. The analytical framework developed here provides exact closed-form expressions for all relevant performance metrics, the asymptotic analysis reveals the diversity and coding gain structure of the system, and the deep learning surrogate enables fast and accurate performance prediction across the full system parameter space.

1.6 Contributions of This Thesis

The original contributions of this thesis can be summarised as follows.

The first contribution is the derivation of closed-form expressions for the average bit-error rate with MPSK modulation and the outage probability of the hybrid FSO/RF first hop under the adaptive combining scheme. The Gamma-Gamma turbulence model with pointing errors is adopted for the FSO sub-link, and the Nakagami-m model is used for the RF sub-link. The combined SNR distribution

under adaptive combining is characterised through its cumulative distribution function, which is expressed as a convergent series of Meijer-G functions. The average BER and outage probability are then obtained in closed form by integrating the complementary error function and the indicator function, respectively, against this distribution. The resulting expressions are numerically stable and computationally tractable when evaluated using variable-precision arithmetic.

The second contribution is the derivation of closed-form average BER expressions for both users in a two-user downlink NOMA system operating over Nakagami-m fading channels with $m = 1$ and QPSK modulation. Separate expressions are derived for the weaker user, which decodes its own signal in the presence of the stronger user's interference without applying SIC, and for the stronger user, which first cancels the weaker user's signal through SIC before decoding its own. Both expressions are given in terms of elementary functions and are therefore straightforward to evaluate. These expressions form the basis of the power allocation analysis that constitutes the third contribution.

The third contribution is the formulation and numerical solution of two power allocation strategies for the NOMA second hop. The first strategy minimises the average bit-error rate of the two-user system, finding the power allocation that achieves the best aggregate performance. The second strategy seeks fairness by equalising the BER of both users, identifying the allocation at which neither user enjoys a performance advantage over the other. Both problems are solved numerically as one-dimensional optimisation tasks, and their solutions are characterised as functions of the per-user signal-to-noise ratio.

The fourth contribution is an asymptotic high-SNR analysis of both the BER and the outage probability using Slater's theorem. The analysis identifies the dominant pole of the Meijer-G function in each closed-form expression and extracts the leading-order decay rate and amplitude. This yields the diversity order and coding gain of the system in closed form, expressed as simple functions of the

Gamma-Gamma and pointing-error parameters. The analytical diversity order is independently verified by extracting the slope of the simulated BER and outage curves at high SNR through a log-log regression. The agreement between the two methods confirms both the correctness of the closed-form expressions and the accuracy of the asymptotic approximation.

The fifth contribution is a deep neural network surrogate model for the average BER of the hybrid FSO/RF system. A dataset of one hundred thousand samples is generated from the MATLAB analytical expressions, spanning the full relevant range of channel parameters and SNR values. A feedforward DNN with three hidden layers is trained on this dataset to predict the log-scaled BER from the five input parameters. The trained DNN achieves a coefficient of determination exceeding 99.8% on the held-out test set, with an evaluation speedup of approximately twenty-four thousand times relative to the MATLAB analytical baseline. XGBoost and Random Forest baselines are trained on the same dataset and benchmarked against the DNN; their in-distribution accuracy is found to be competitive, but both fail to extrapolate correctly to SNR values outside the training range, a limitation that the DNN does not share.

1.7 Organization of the Thesis

This thesis is structured into seven chapters, covering the system model, performance analysis, asymptotic characterisation, and deep learning evaluation of the proposed dual-hop hybrid FSO/RF–NOMA relay system. The contents of each chapter are summarised below.

Chapter 1: Introduction
This chapter establishes the motivation for the research, tracing the trajectory from current 5G networks through the 6G vision to the specific technology gaps that the thesis addresses. The roles of free-space optical communications, hybrid FSO/RF combining, NOMA, and deep learning in next-generation networks are discussed at a

conceptual level. The four research gaps targeted by the thesis are identified, and the five original contributions are stated. The chapter concludes with the present thesis organisation.

Chapter 2: Literature Survey
A comprehensive review of the existing literature is presented, covering the development of FSO channel models from log-normal through Gamma-Gamma to more recent generalised distributions, the evolution of hybrid FSO/RF combining strategies from hard switching to adaptive combining and always-on combining, the application of NOMA to relay-assisted FSO systems, and the growing body of work on deep learning and machine learning surrogate models for wireless performance prediction. The chapter identifies the specific contributions that distinguish the present work from the existing literature and justifies the system model and analytical methodology adopted in subsequent chapters.

Chapter 3: System Model and Channel Statistics
This chapter presents the complete dual-hop relay topology and the statistical channel models used throughout the thesis. The Gamma-Gamma turbulence distribution with pointing errors is formally introduced for the FSO sub-link of the first hop, and the Nakagami-m fading model is introduced for the RF sub-link. The Nakagami-m fading model with $m = 1$ (Rayleigh) adopted for the NOMA user channels is described. The decode-and-forward relay protocol and the NOMA superposition coding and successive interference cancellation procedures are formalised. All simulation parameters for the moderate and strong turbulence cases are defined and tabulated.

Chapter 4: Performance Analysis
This chapter develops the closed-form performance analysis. The CDF of the combined first-hop SNR under adaptive combining is derived in a Meijer-G series form. The outage probability of the first hop is expressed using this CDF. The average BER of the first hop is derived by decomposing the integral into contributions from the combined FSO/RF region and the FSO-only region. The BER expressions for both NOMA users are derived in closed form. The

optimal and fairness-based power allocation problems are formulated and solved numerically, with results presented as functions of the per-user SNR.

Chapter 5: Asymptotic Analysis
This chapter presents the high-SNR asymptotic analysis of the derived BER and outage expressions using Slater's theorem. The dominant Meijer-G pole and its associated residue are identified for each performance metric, yielding the diversity order and coding gain in closed form. The analytical diversity orders for both moderate and strong turbulence are compared against values extracted from the slope of the numerically computed BER and outage curves, confirming the correctness of the asymptotic predictions. The physical interpretation of the results in terms of turbulence severity and its impact on link budget is discussed.

Chapter 6: Deep Learning Based Evaluation
This chapter presents the machine learning surrogate modelling framework. The dataset generation procedure, input parameter ranges, log-transform preprocessing, and train-test split are described. The DNN architecture, training methodology, optimiser settings, and early stopping criteria are specified. XGBoost and Random Forest baseline models are described and trained on the same dataset. The performance of all three models is evaluated in terms of RMSE, coefficient of determination, computational time, and extrapolation behaviour. The failure mode of tree-based models at out-of-distribution SNR values is analysed and contrasted with the DNN's reliable extrapolation.

Chapter 7: Results and Discussion
This chapter consolidates and interprets the numerical results obtained in the preceding chapters. Average BER and outage probability curves are presented for both turbulence regimes, with exact analytical, asymptotic, and DNN-predicted curves overlaid for direct comparison. The power allocation trade-off between BER minimisation and user fairness is illustrated. Key performance metrics including diversity order, coding gain, power penalty between

turbulence regimes, DNN accuracy, and computational speedup are tabulated and discussed. The chapter closes with a summary of the key findings and their implications for practical system design.

Chapter 8: Conclusion and Future Work
The final chapter summarises the research contributions and the principal conclusions drawn from the analysis. Directions for future research are identified, including extension to more than two NOMA users, incorporation of hardware impairments such as phase noise and imperfect successive interference cancellation, integration of UAV mobility into the relay model, application of the Meijer-G framework to reconfigurable intelligent surface-assisted FSO links, and online adaptive retraining of the DNN surrogate under non-stationary channel conditions.

Chapter 2

Literature Survey

2.1 Overview

The research presented in this thesis draws on four largely parallel streams of technical literature: the modelling and performance analysis of free-space optical channels, the design and evaluation of hybrid FSO/RF combining strategies and relay architectures, the application of non-orthogonal multiple access to relay-assisted wireless systems, and the use of machine learning as a computational surrogate for analytical performance expressions. This chapter surveys each of these streams in turn, tracing their evolution from foundational contributions to the state of the art, and concludes by identifying the specific gaps that the present work is designed to fill. The survey is organised thematically rather than strictly chronologically, since the technical dependencies between topics are more instructive than the order in which individual papers appeared.

The field of FSO communications has matured considerably over the past two decades, and a rich body of analytical work now exists for characterising the statistics of turbulence-induced irradiance fading, modelling the effects of pointing errors, and deriving performance metrics such as average bit-error rate, outage probability, and ergodic capacity in closed form. Separately, the study of hybrid FSO/RF systems has produced a range of combining strategies and relay topologies, with much of the analytical work relying on the same Meijer-G function framework that underpins the first-

hop analysis of this thesis. The NOMA literature, which originated largely in the context of RF cellular systems, has more recently been extended to FSO and mixed FSO/RF relay settings, though the combination with hybrid adaptive combining has not been studied. Finally, the application of deep learning and ensemble methods to wireless performance prediction is an emerging area that has been validated in a handful of closely related system configurations. The following sections review each of these threads with sufficient depth to motivate and position the contributions of this thesis.

2.2 Free-Space Optical Channel Modelling

The accurate statistical modelling of the atmospheric channel is the foundation on which all analytical performance evaluation of FSO systems rests. The earliest widely adopted model for turbulence-induced irradiance fading was the log-normal distribution, in which the logarithm of the received irradiance is assumed to follow a Gaussian distribution. The log-normal model is physically motivated by the multiplicative accumulation of independent small-scale phase distortions along the propagation path and provides a good approximation for weak turbulence conditions characterised by a small Rytov variance. However, it is well established that the log-normal model underestimates the probability of deep fades under moderate to strong turbulence, making it unsuitable for link-budget analysis in the operating regimes most relevant to practical outdoor FSO deployments [1, 2].

The Gamma-Gamma distribution, proposed by Al-Habash, Andrews, and Phillips [3] and subsequently adopted as the standard model for FSO turbulence analysis, resolves the limitations of the log-normal approximation by modelling the received irradiance as the product of two independent Gamma random variables, one representing large-scale atmospheric eddies and one representing small-scale eddies. The two shape parameters of the distribution are directly related to the Rytov variance and to the ratio of the outer and

inner turbulence scales, and can be estimated from measured atmospheric data with good accuracy. Crucially, the Gamma-Gamma distribution remains accurate across the full range of turbulence conditions from weak through moderate to strong, and its probability density function is expressible in closed form through the modified Bessel function of the second kind, making it analytically tractable for BER and outage derivations [4]. The cumulative distribution function of the Gamma-Gamma SNR can be expressed in terms of the Meijer-G function, which is the starting point for the performance analysis of this thesis.

A number of authors have extended the baseline Gamma-Gamma model to incorporate additional physical effects. Sandalidis et al. [5] characterised the BER performance of FSO links under strong turbulence with pointing errors modelled via the Farid-Hranilovic misalignment fading framework, deriving tight closed-form approximations using series expansions of the complementary error function. Farid and Hranilovic [6] developed the foundational pointing error model that is now universally adopted: by assuming a Gaussian spatial intensity profile and a circular receiver aperture, they derived the joint distribution of turbulence and misalignment fading and showed that the ratio of the equivalent beam width to twice the jitter standard deviation is the single parameter governing the severity of the pointing error penalty [4]. This model was subsequently incorporated into Meijer-G-based performance analyses by Saxena, Mathur, and Bhatnagar, who derived exact and asymptotic BER expressions for optically pre-amplified FSO systems under combined turbulence and pointing errors, obtaining the diversity order in closed form through Slater's theorem — the same analytical technique applied in this thesis [2].

An important practical dimension of FSO performance evaluation is the choice of detection technique at the photodetector. Two approaches are commonly considered: intensity modulation with direct detection, in which the transmitted information is encoded in the instantaneous optical power and recovered by a simple photodiode; and coherent heterodyne detection, in which the received field

is mixed with a local oscillator laser before photodetection, enabling phase-sensitive reception and improved sensitivity. Heterodyne detection provides an SNR advantage of a factor of two over intensity modulation with direct detection for the same transmit power, and the performance gap between the two techniques has been quantified analytically in the Gamma-Gamma framework by several authors. The present thesis adopts intensity modulation with direct detection, consistent with the low-cost terminal assumption appropriate for the terrestrial relay deployment scenario considered here. The heterodyne detection literature provides a useful upper bound on the achievable performance at the same transmit power.

More recent work has explored alternative turbulence models for specialised propagation environments. The Málaga distribution, proposed for satellite-to-ground FSO links, subsumes the Gamma-Gamma and several other distributions as special cases and has been used to model the upward-looking channels relevant to UAV-relay scenarios. The double Weibull and double generalised Gamma models have been proposed for strong turbulence and extended propagation paths. However, the Gamma-Gamma model with Farid-Hranilovic pointing errors remains the most widely validated choice for terrestrial horizontal paths in the moderate-to-strong turbulence regime. The two scenarios used in this work — moderate and strong turbulence with parameters representative of outdoor horizontal links in continental climates — are directly comparable with the parameter regimes studied in the foundational FSO performance analysis literature, and enable a meaningful quantitative comparison of the system's behaviour across a practically relevant range of atmospheric conditions.

2.3 Hybrid FSO/RF Systems and Combining Strategies

The concept of augmenting an FSO link with a parallel RF link to improve availability against weather-induced outages has been studied since the early 2000s and has produced a substantial analyt-

ical literature. The earliest analytical treatments of hybrid FSO/RF systems employed the hard-switching or single-threshold switching strategy, in which the system operates exclusively on the FSO link when the optical channel quality is adequate and transfers entirely to the RF link when it is not. The outage probability and average BER of hard-switching hybrid FSO/RF systems have been derived in closed form for a variety of turbulence models and fading distributions, typically yielding expressions that are a weighted sum of the standalone FSO and RF performance metrics [7, 8, 4].

Liang, Chen, and Ke provided a comprehensive analytical treatment of a dual-hop hybrid FSO/RF–THz relay system, deriving closed-form Meijer-G expressions for the outage probability, average BER, and ergodic capacity of the system under combined Gamma-Gamma turbulence, pointing errors, and Nakagami-m RF fading [4]. Their work established the analytical toolkit — specifically the Meijer-G series representation of the FSO SNR CDF — that forms the backbone of the first-hop analysis in this thesis. Crucially, however, their combining strategy was the single-threshold scheme: when the FSO SNR fell below the threshold, all traffic was routed to the RF (or THz) link, discarding the residual FSO signal energy. The present thesis extends this framework to the adaptive combining scheme, in which the FSO and RF links are maximal-ratio combined during FSO fades rather than switched, requiring a fundamentally different derivation of the combined SNR distribution.

The performance of adaptive and soft-switching hybrid FSO/RF systems has been studied in a more limited body of work. Amirabadi proposed a novel multi-hop relay-assisted hybrid FSO/RF architecture and derived closed-form BER and outage expressions for both exact and asymptotic regimes, demonstrating that the relay-assisted hybrid architecture significantly outperforms both standalone FSO and single-hop hybrid configurations. However, that work employed a selection combining scheme rather than maximal-ratio combining and did not address the NOMA second hop. Wu et al. analysed multi-hop parallel hybrid FSO/RF communication systems with

UAV-mounted relays, deriving new BER and outage expressions under the Exponential Weibull turbulence model with Nakagami-m RF fading and different modulation methods [9]. Their results demonstrated that UAV-assisted hybrid FSO/RF systems consistently outperform ground-level direct-link systems and that the relay structure has a significant influence on system performance. Again, however, the combining scheme was selection-based rather than adaptive maximal-ratio combining, and no NOMA access layer was considered.

The UAV-assisted FSO/RF relay literature has grown considerably in recent years as UAV platforms have emerged as a practical means of extending FSO coverage over terrain obstacles and providing line-of-sight clearance for optical links. Qu et al. analysed a UAV-assisted dual-hop RF/FSO relay system for space-air-ground integrated networks, employing amplify-and-forward at the UAV relay and characterising performance under shadowed Rician and Málaga fading for the two hops respectively. Latka et al. studied UAV-enabled dual-hop mixed RF/FSO systems with user scheduling, demonstrating that opportunistic user selection at the destination substantially improves the outage performance relative to fixed-user systems [10]. These works, while valuable, focus on single-user destinations and do not incorporate NOMA on the second hop.

2.4 Non-Orthogonal Multiple Access in Relay Systems

The application of NOMA to relay-assisted wireless networks has been extensively studied in the context of pure RF systems [11, 12, 13, 14]. The foundational works in this area demonstrated that power-domain NOMA [15], when combined with a decode-and-forward relay, can provide substantial throughput gains over orthogonal multiple access schemes for heterogeneous user populations with dissimilar channel conditions. Kara and Kaya [16] derived exact closed-form bit-error probability expressions for cooperative NOMA systems with a decode-and-forward relay over Nakagami-

m fading channels, showing that the BER of the NOMA pair is strongly dependent on the power allocation ratio and on the quality of the SIC performed at the stronger user. This work established the analytical approach for two-user NOMA BER derivation that, in modified form, underlies the QPSK expressions for the second hop of this thesis. Expressions for both the case where SIC is performed correctly — the interference-free regime of the stronger user — and the case where it is not — the interference-limited regime — were derived, providing a complete characterisation of the NOMA error performance.

The integration of NOMA with FSO backhaul was first studied by Ngo and Duong, who considered a relay-assisted uplink NOMA system with an FSO backhaul link between the relay and destination. Dynamic-order decoding was applied at the destination to optimise the SIC order based on instantaneous channel state information, and closed-form outage probability expressions were derived assuming Gamma-Gamma turbulence with pointing errors for the FSO hop and independent Rayleigh fading for the user-relay access links. This work was one of the first to demonstrate the feasibility of mixing power-domain NOMA with an FSO backhaul and showed that FSO backhauling provides substantially better outage performance than RF backhauling for the same allocated power, owing to the much higher SNR achievable on the optical link under clear-sky conditions [4].

Subsequent work extended the FSO-NOMA relay paradigm in several directions. Energy-harvesting FSO-NOMA systems, in which the relay harvests energy from the received optical signal before retransmitting to NOMA users, have been studied with both amplify-and-forward and decode-and-forward relay protocols, yielding Meijer-G-based outage expressions that include the energy conversion efficiency as an additional system parameter. Cognitive radio NOMA systems with mixed FSO/RF links have been analysed, where the FSO link is shared between a primary and a secondary system and the secondary system's NOMA users must satisfy an interference threshold at the primary receiver. In all of these works,

the first hop connecting the source to the relay is either a pure FSO link or a pure RF link: the composite hybrid FSO/RF adaptive combining architecture studied in this thesis — in which the first hop simultaneously supports both optical and RF sub-links that are adaptively combined — has not been addressed in the context of a NOMA-enabled second hop.

The performance of full-duplex NOMA relay systems has also been studied extensively, with results showing that full-duplex relaying offers better outage and ergodic rate performance than half-duplex relaying in the low-SNR regime, but is surpassed by half-duplex at high SNR due to residual self-interference. Since the relay in the system of this thesis operates in half-duplex decode-and-forward mode, the full-duplex NOMA relay literature provides a useful performance upper bound but is not directly applicable to the system under study. Cooperative NOMA, in which the stronger near user acts as a relay to assist the weaker far user, has been studied with SWIPT-enabled energy harvesting at the near user and over both Rayleigh and Nakagami-m fading channels. These works demonstrate that cooperative NOMA can eliminate the zero-diversity-order problem of the far user in certain topologies, a finding that motivates the power allocation design in the second hop of this thesis.

2.5 Deep Learning and Machine Learning for Wireless Systems

The application of machine learning to wireless communications has undergone a rapid expansion over the past decade, driven by the availability of large datasets, powerful GPU hardware, and open-source deep learning frameworks [17, 18]. The earliest applications focused on channel estimation for OFDM systems, where convolutional and dense neural networks were shown to outperform least-squares and minimum mean-square-error estimators at moderate to high SNR, particularly under non-stationary channel conditions where the channel statistics change faster than the estimation up-

date rate. Stenhammar et al. provided a systematic comparison of neural network architectures for wireless channel prediction in 6G scenarios, including recurrent networks, transformers, and dense feedforward networks, demonstrating that the optimal architecture depends strongly on the prediction horizon and the degree of temporal channel correlation. The study by Ye et al. [19] on deep learning for OFDM channel estimation was particularly influential, showing that an end-to-end learned equaliser could match or exceed pilot-based least-squares performance while requiring fewer pilot symbols, an insight that directly motivated the broader programme of replacing hand-designed signal processing components with learned modules [9].

The use of neural networks as surrogate models for analytically derived performance metrics — as distinct from channel estimation or signal detection — is a more specialised but growing research direction. Shukla et al. provided an early and directly relevant demonstration of this approach for a decode-and-forward coordinated direct and relay transmission system with cognitive NOMA and energy harvesting [20]. In that work, a three-hidden-layer feedforward DNN was trained on a dataset generated from closed-form analytical BER expressions and shown to reproduce the mapping with a coefficient of determination exceeding 99.8% while reducing the evaluation time from several minutes per curve to a few milliseconds. Critically, the authors observed that log-transforming the BER target variable — replacing the raw BER, which spans many decades, with its base-10 logarithm — was essential for stable training, as it converted the highly skewed output distribution into a roughly symmetric one with a well-conditioned mean-squared-error loss landscape. This observation is directly adopted in the DNN design of this thesis. Shukla et al. also identified that the DNN generalises correctly to parameter combinations not seen during training, provided the training distribution was drawn uniformly over the relevant parameter space rather than from a fixed grid.

Cheepurupalli, Daftari, Upadhyay, and Salhab presented the most directly related prior work to this thesis [21]. Their paper

demonstrated the application of a DNN surrogate for performance prediction in a UAV-assisted FSO/RF NOMA system, where both the FSO turbulence parameters and the NOMA power allocation coefficients were treated as inputs to the network. The DNN correctly reproduced the analytical BER and outage curves for both moderate and strong turbulence regimes without requiring separate models for each regime, demonstrating that a single well-designed network can capture the qualitative change in system behaviour that accompanies the transition from moderate to strong turbulence. That work also noted that the DNN's extrapolation behaviour — its ability to produce physically consistent predictions at SNR values outside the training range — is a critical advantage over lookup-table-based approximations, which are strictly bounded by the discretisation of the parameter space. The present thesis builds directly on this approach, applying the same architecture to the more complex hybrid adaptive combining system and adding the systematic tree-ensemble comparison that was not performed in the prior work.

The application of tree-based ensemble methods to wireless performance prediction has received less attention than neural network approaches, but several studies have demonstrated competitive in-distribution accuracy. Random Forests [22], which aggregate the predictions of many independently trained decision trees using bootstrap aggregation, are known to be robust to overfitting and to provide implicit feature importance rankings that can be used to identify the channel parameters most strongly influencing system performance. Gradient-boosted decision trees, in particular the XGBoost implementation [23], sequentially train trees to correct the residual prediction errors of the previous ensemble, typically achieving lower prediction error than Random Forests on structured tabular data with a smaller number of estimators. Both methods have been applied to BER prediction in cognitive radio systems, to outage prediction in 5G networks, and to signal quality estimation in interference-limited environments, consistently achieving high in-distribution accuracy. However, the fundamental algorithmic limitation of tree-based methods — their inability to extrapolate

beyond the convex hull of the training data, since each prediction is a piecewise-constant average of training labels within a hyper-rectangular region of the feature space — has not been characterised in the FSO performance prediction context. This extrapolation failure, which produces pathological artefacts in the predicted BER curve at SNR values above the training range, is documented and analysed in this thesis as part of the comparative evaluation against the DNN surrogate.

The broader literature on AI-assisted 6G network management is relevant context for understanding the motivation for the surrogate modelling approach. Stenhammar et al. demonstrated that channel prediction horizons of tens of milliseconds are achievable with neural networks in high-mobility scenarios, supporting the feasibility of using learned models for real-time resource allocation decisions [9]. The deep learning survey by Zhu et al. identified three distinct roles for neural networks in the physical layer: end-to-end learned transceivers that jointly optimise transmitter and receiver processing [24], model-based methods that incorporate domain knowledge into the network architecture, and model-free surrogates that approximate expensive analytical computations. The surrogate modelling approach of this thesis falls squarely in the third category and complements the other two by providing a computationally efficient tool for system design exploration and parameter optimisation that does not require the full analytical evaluation to be repeated at each operating point.

A closely related development in the machine learning for wireless systems literature is the study of neural network generalisation and uncertainty quantification. It has been observed that densely connected feedforward networks with ReLU activations, when trained with sufficient data and appropriate regularisation, exhibit strong generalisation to input regions outside the convex hull of the training set, because the piecewise-linear structure of ReLU compositions allows the network to extrapolate the learned function trends rather than merely interpolating between training points. This property is in sharp contrast to the behaviour of tree-

based ensemble methods, and it is the primary reason that the DNN is preferred over XGBoost and Random Forest in the context of BER surrogate modelling, where the evaluation range of interest — particularly at very high SNR — may extend well beyond the parameter space used during training. The systematic documentation of this distinction, through controlled comparison on matched training datasets, is one of the methodological contributions of Chapter 6 and provides a reproducible benchmark for future surrogate modelling work in this area.

2.6 Summary and Identified Research Gaps

The preceding survey reveals a well-developed analytical tradition in each of the four technology areas addressed by this thesis, together with a clear and consistent set of gaps at the intersections between those areas. In the FSO channel modelling literature, the Gamma-Gamma model with Farid-Hranilovic pointing errors is firmly established as the standard for terrestrial horizontal paths, and the Meijer-G framework has been validated as the preferred analytical tool for deriving performance metrics in closed form. The diversity order and coding gain analysis via Slater's theorem has been demonstrated for FSO-only systems by Saxena et al. [2], but has not been applied to the composite hybrid system studied here.

In the hybrid FSO/RF combining literature, the hard-switching or single-threshold scheme is by far the most commonly studied architecture, with closed-form analytical results available for a variety of system configurations. The adaptive combining scheme, which is known from simulation studies to outperform hard switching by several decibels in both BER and outage at the cost of activating the RF chain during FSO fades, has received much less analytical attention, and no closed-form Meijer-G expressions have been reported for the average BER of the adaptive combining scheme under Gamma-Gamma turbulence with pointing errors. This is the first principal gap addressed by this thesis.

In the FSO-NOMA relay literature, the combination of FSO backhaul with NOMA access has been demonstrated for pure FSO and mixed RF/FSO first-hop configurations, but not for the hybrid adaptive combining first hop that is the focus of this thesis. All existing FSO-NOMA relay studies treat the first-hop FSO channel and the second-hop NOMA channel as statistically independent, which is appropriate when the first hop is a dedicated point-to-point FSO link, but the hybrid adaptive combining first hop introduces a threshold-dependent coupling between the FSO and RF sub-links that must be tracked through the relay processing to obtain the correct end-to-end performance. This is the second gap addressed by this thesis.

In the machine learning for wireless performance prediction literature, the DNN surrogate approach has been validated for energy-harvesting NOMA systems by Shukla et al. [20] and for UAV-assisted FSO/RF NOMA systems by Cheepurupalli et al. [21], but has not been applied to the hybrid adaptive combining architecture studied here. More importantly, no prior work has provided a systematic comparison of DNN and tree-based ensemble surrogates for FSO-related performance prediction, and the extrapolation failure of tree-based methods at out-of-distribution SNR values — which is a fundamental algorithmic property rather than a hyperparameter tuning issue — has not been documented or quantified in this context. These constitute the third and fourth gaps addressed by this thesis.

Taken together, these gaps define a coherent and self-contained research programme: derive the analytical performance framework for the hybrid adaptive combining FSO/RF–NOMA relay system, extract the asymptotic diversity and coding gain structure via Slater's theorem, and replace the computationally expensive Meijer-G evaluation with a DNN surrogate that is benchmarked against tree-based alternatives with explicit extrapolation characterisation. The technical development of this programme begins in the following chapter with the formal system model and channel statistics.

Chapter 3

System Model

3.1 System Description

The system studied in this thesis is a dual-hop decode-and-forward relay network with a source S, a relay R, and two downlink users U_1 and U_2. No direct link exists between the source and either user; all communication is mediated by the relay. Operation is divided into two time slots. In the first slot S transmits to R over a hybrid FSO/RF link employing adaptive combining. In the second slot R decodes the received symbol and retransmits to both users simultaneously via two-user downlink power-domain NOMA with Gray-coded QPSK modulation over independent Nakagami-m fading channels with $m = 1$. The topology is shown in Figure 3.1.

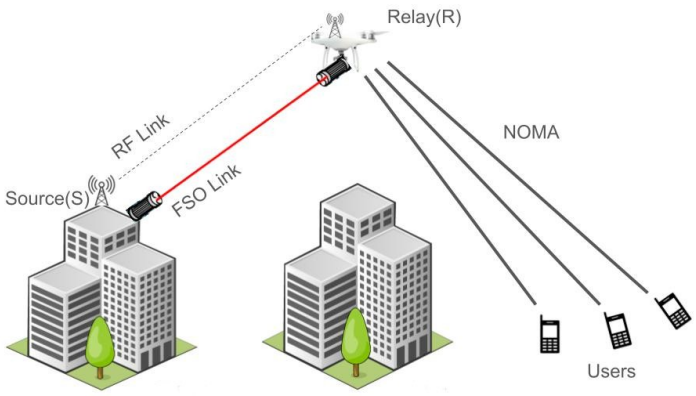

Figure 3.1: Dual-hop hybrid FSO/RF–NOMA relay system. The source (S) transmits to the UAV-mounted relay (R) via a parallel FSO link (red beam) and RF link (dashed). The relay applies adaptive combining and retransmits to NOMA users via the second hop.

The adaptive combining rule is as follows. Let γ_F and γ_R be the instantaneous SNRs on the FSO and RF sub-links and γ_{th} the switching threshold. When $\gamma_F \geq \gamma_{th}$ the relay uses FSO exclusively ($\gamma_c = \gamma_F$); when $\gamma_F < \gamma_{th}$ the RF chain is activated via 1-bit feedback and MRC is applied, yielding $\gamma_c = \gamma_F + \gamma_R$. The relay then applies decode-and-forward processing to γ_c.

3.2 FSO Channel Model

The FSO sub-link is modelled under Gamma-Gamma atmospheric turbulence and Farid-Hranilovic pointing errors. The received irradiance is $I_F = \mu_l h_a h_p$, where μ_l is deterministic path loss, h_a is the turbulence fading coefficient, and h_p is the pointing-error attenuation. The Gamma-Gamma PDF of h_a is

$$f_{h_a}(x) = \frac{2(\alpha_F\beta_F)^{(\alpha_F+\beta_F)/2}}{\Gamma(\alpha_F)\Gamma(\beta_F)}\, x^{(\alpha_F+\beta_F)/2-1} K_{\alpha_F-\beta_F}\left(2\sqrt{\alpha_F\beta_F x}\right), \quad x > 0, \tag{3.1}$$

where $\Gamma(\cdot)$ is the gamma function and $K_\nu(\cdot)$ is the modified Bessel function of the second kind of order ν. The parameters $\alpha_F > 0$ and $\beta_F > 0$ represent the effective numbers of large-scale and small-scale turbulence cells; their scenario values are in Table 3.1.

The pointing-error severity is characterised by $\xi_F = w_{L_{\text{eq}}}/(2\sigma_{j,F})$, the ratio of the equivalent beam radius to twice the jitter standard deviation. The composite CDF of the received FSO SNR $\gamma_F = \bar{\gamma}_F I_F/\mu_l$ is [4]:

$$F_{\gamma_F}(x) = \frac{2^{\alpha_F+\beta_F-3}\xi_F^2}{\pi\,\Gamma(\alpha_F)\Gamma(\beta_F)} \times G_{3,7}^{6,1}\left[\frac{\alpha_F^2\beta_F^2\xi_F^4\, x}{16(\xi_F^2+1)^2\bar{\gamma}_F}\,\middle|\, \begin{matrix} 1,\ \frac{\xi_F^2+1}{2},\ \frac{\xi_F^2+2}{2} \\ \frac{\xi_F^2}{2},\ \frac{\xi_F^2+1}{2},\ \frac{\alpha_F}{2},\ \frac{\alpha_F+1}{2},\ \frac{\beta_F}{2},\ \frac{\beta_F+1}{2},\ 0 \end{matrix}\right]. \tag{3.2}$$

where $\bar{\gamma}_F$ is the average electrical SNR normalised to unit path loss. The Meijer-G representation in (3.2) is exact for all $x > 0$ and all valid parameter ranges $\alpha_F, \beta_F, \xi_F > 0$. It is evaluated numerically using MATLAB's `vpa` at 100-digit precision, since the series terms in the Meijer-G expansion can be of alternating sign and large magnitude, causing catastrophic cancellation in double-precision arithmetic [4].

Two important special cases of (3.2) serve as validation references. In the limit $\xi_F \to \infty$ (perfect beam alignment), the pointing-error distribution degenerates and (3.2) reduces to the standard Gamma-Gamma CDF without pointing errors, which has a simpler five-parameter Meijer-G representation. As α_F and β_F both tend

to infinity (vanishing turbulence), the Gamma-Gamma distribution concentrates at its mean and the CDF approaches a step function, corresponding to a deterministic AWGN channel. The differentiated PDF $f_{\gamma_F}(x) = \mathrm{d}F_{\gamma_F}(x)/\mathrm{d}x$ is obtained analytically by differentiating (3.2) and is used directly in the BER integrals of Chapter 4.

3.3 RF Channel Model

The RF sub-link follows a Nakagami-m model with parameter m and mean SNR $\bar{\gamma}_R$. The instantaneous SNR γ_R has PDF

$$f_{\gamma_R}(x) = \frac{m^m}{\bar{\gamma}_R^m\,\Gamma(m)}\, x^{m-1} e^{-mx/\bar{\gamma}_R}, \quad x \geq 0. \tag{3.3}$$

With $m = 1$ (Rayleigh) this reduces to $f_{\gamma_R}(x) = (1/\bar{\gamma}_R)\, e^{-x/\bar{\gamma}_R}$ and $F_{\gamma_R}(x) = 1 - e^{-x/\bar{\gamma}_R}$. This assumption is standard for a non-line-of-sight link [4]. The FSO and RF sub-links are statistically independent since optical scintillation and multipath fading arise from entirely different physical mechanisms.

3.4 NOMA Second-Hop Signal Model

On the second hop the relay broadcasts the superposed NOMA waveform [16]:

$$x = \sqrt{\beta_1 P}\, s_1 + \sqrt{\beta_2 P}\, s_2, \qquad \beta_1 > \beta_2 > 0, \quad \beta_1 + \beta_2 = 1, \tag{3.4}$$

where P is the relay transmit power (normalised to unity), β_n is the power fraction, and s_n is the unit-energy Gray-coded QPSK symbol of user n drawn equiprobably from the alphabet $\mathbb{S} = \{a_0, a_1, a_2, a_3\}$.

The in-phase and quadrature amplitudes of the superposed signal x span a 16-point constellation characterised by the amplitude factors [16]:

$$A_{u_1 u_2} = u_1\sqrt{\beta_1} + u_2\sqrt{\beta_2}, \qquad u_i \in \{0,\, 1,\, \bar{1},\, 2\}, \quad \bar{1} \triangleq -1. \tag{3.5}$$

The amplitudes entering the BER derivation are $A_{11} = \sqrt{\beta_1} + \sqrt{\beta_2}$,

$A_{1\bar{1}} = \sqrt{\beta_1} - \sqrt{\beta_2}$ (for User 1), and additionally $A_{01} = \sqrt{\beta_2}$, $A_{2\bar{1}} = 2\sqrt{\beta_1} - \sqrt{\beta_2}$, $A_{21} = 2\sqrt{\beta_1} + \sqrt{\beta_2}$ (for User 2 with SIC errors).

The received signal at user n is $r_n = h_n x + w_n$, where h_n follows Nakagami-m flat-fading with $m = 1$ (Rayleigh), so $h_n \sim \mathcal{CN}(0, 1)$, $\alpha_n = |h_n|$ is the channel gain, and $w_n \sim \mathcal{CN}(0, N_0)$. After phase compensation the in-phase decision variable at user n is $\mathfrak{r}_n = \alpha_n x_I + n_n$ where $n_n \sim \mathcal{N}(0, N_0/2)$. The instantaneous per-user SNR is $\gamma_n = \bar{\gamma}_n |h_n|^2$ with mean $\bar{\gamma}_n = 1/N_0$.

User 1 (no SIC): U_1 detects s_1 treating s_2 as interference, using the maximum likelihood detector [16]. Enumerating all 16 equally likely (s_1, s_2) combinations and applying the I-Q decision boundary at zero yields two distinct effective SNR terms:

$$\gamma_{1,1} = \frac{\alpha_1^2 A_{1\bar{1}}^2}{N_0/2}, \qquad \gamma_{1,2} = \frac{\alpha_1^2 A_{11}^2}{N_0/2}. \tag{3.6}$$

The conditional BER of User 1 is [16]:

$$P_{U_1}\big|\alpha_1 = \frac{1}{2}\Big[Q\big(\sqrt{\gamma_{1,1}}\big) + Q\big(\sqrt{\gamma_{1,2}}\big)\Big]. \tag{3.7}$$

User 2 (SIC): U_2 first decodes s_1, subtracts it (SIC), and then detects s_2 from the residual. Accounting for both correct SIC ($\hat{s}_1 = s_1$) and incorrect SIC ($\hat{s}_1 \neq s_1$) over all symbol combinations yields five effective SNR terms $\gamma_{2,i} = \alpha_2^2 A_{2,i}^2/(N_0/2)$ with amplitudes $A_{2,i} \in \{A_{01}, A_{1\bar{1}}, A_{11}, A_{2\bar{1}}, A_{21}\}$ and signed weights $\mathbf{v} = [2, 1, -1, -1, 1]$. The conditional BER of User 2 is [16]:

$$P_{U_2}\big|\alpha_2 = \frac{1}{2}\sum_{i=1}^{5} v_i\, Q\big(\sqrt{\gamma_{2,i}}\big). \tag{3.8}$$

The full derivation of these expressions and their closed-form averages over the Nakagami-m ($m = 1$) channel gains are presented in Chapter 4.

3.5 Performance Metrics

Two metrics are evaluated. The *outage probability* of the first hop is

$$\mathcal{P}_{\text{out}} = \Pr[\gamma_c < \gamma_{\text{out}}] = F_{\gamma_c}(\gamma_{\text{out}}), \tag{3.9}$$

where $F_{\gamma_c}(\cdot)$ is the CDF of the combined SNR under adaptive combining. The *average BER* of the first hop for BPSK modulation ($A = 1$, $B = 1$) is

$$\bar{P}_e = \int_0^\infty \frac{1}{2}\,\text{erfc}\big(\sqrt{x}\big)\, f_{\gamma_c}(x)\,\mathrm{d}x, \tag{3.10}$$

where $f_{\gamma_c}(x) = \mathrm{d}F_{\gamma_c}(x)/\mathrm{d}x$. The average BER of each NOMA user is obtained by averaging (3.7) and (3.8) over the Nakagami-m distribution with $m = 1$, as derived in Chapter 4.

3.6 Simulation Parameters

Two turbulence regimes are studied: moderate ($\sigma_R^2 \approx 0.5$) and strong ($\sigma_R^2 \approx 1.6$). All other parameters are common to both scenarios and are summarised in Table 3.1.

The threshold $\gamma_{th} = 7\,\text{dB}$ keeps the RF chain inactive during most fair-weather operation while ensuring the MRC fallback activates well before the optical link fails. The allocation $\beta_1 = 0.8$ is a representative two-user design point whose optimality is studied in Chapter 4. The pointing-error ratio $\xi_F = 4.5747$ corresponds to a gimbal-stabilised terminal with sub-micro-radian jitter and is held constant across scenarios to isolate turbulence effects.

Table 3.1: Simulation parameters for moderate and strong turbulence scenarios.

Parameter	**Symbol**	**Moderate**	**Strong**
Large-scale scintillation	α_F	5.6137	4.7206
Small-scale scintillation	β_F	4.0062	1.2066
Pointing-error ratio	ξ_F	4.5747	4.5747
Nakagami-m (RF)	m	1	1
Mean RF SNR	$\bar{\gamma}_R$	5 dB	
Switching threshold	γ_{th}	7 dB	
Outage threshold	γ_{out}	3 dB	
NOMA power (U1)	β_1	0.8	
NOMA power (U2)	β_2	0.2	
Modulation (first hop)	BPSK	$A = 1,\ B = 1$	
Modulation (NOMA hop)	QPSK	—	
FSO SNR sweep	$\bar{\gamma}_F$	0–60 dB	
Series truncation (outage)	n	0 to 20	
Series truncation (BER)	n, i, j	0 to 10	
DNN training samples	N_{train}	100 000	

Chapter 4

Performance Analysis

4.1 Overview

This chapter derives closed-form analytical expressions for every performance metric of the dual-hop hybrid FSO/RF–NOMA relay system defined in Chapter 3. The derivation proceeds in four stages. First, the CDF of the combined first-hop SNR γ_c under the adaptive combining rule is obtained as a Meijer-G series. Second, the first-hop outage probability is expressed directly from this CDF. Third, the average BER of the first hop is derived by splitting the integral at the switching threshold γ_{th} and evaluating each piece in closed form. Fourth, the average BER of each NOMA user on the second hop is derived by enumerating the 16-point superposition constellation, applying the QPSK geometry of Assaf et al. [16], and averaging over the Nakagami-m fading distribution with $m = 1$ using the standard Q-function identity. The chapter concludes with the formulation and solution of two power allocation problems.

4.2 CDF of the Combined First-Hop SNR

Under the adaptive combining rule, the combined SNR is $\gamma_c = \gamma_F$ when $\gamma_F \geq \gamma_{th}$ and $\gamma_c = \gamma_F + \gamma_R$ when $\gamma_F < \gamma_{th}$. The CDF of γ_c is

therefore [4]:

$$F_{\gamma_c}(x) = \mathcal{P}_{\text{out,AC}}(x) = \begin{cases} F_1(x), & x \le \gamma_{th}, \\ F_2(x) - F_{\gamma_F}(\gamma_{th}) + F_{\gamma_F}(x), & x > \gamma_{th}, \end{cases} \tag{4.1}$$

where the two auxiliary CDFs are $F_1(x) = \int_0^x f_{\gamma_F+\gamma_R}(t)\,\mathrm{d}t$ and $F_2(x) = \int_0^{\gamma_{th}} f_{\gamma_F}(t)\,F_{\gamma_R}(x-t)\,\mathrm{d}t$, respectively.

For the operationally relevant case $\gamma_{\text{out}} \le \gamma_{th}$, substituting the Gamma-Gamma–pointing-error PDF $f_{\gamma_F}(x)$ and the Nakagami-m RF PDF with $m = 1$, giving $f_{\gamma_R}(x) = (1/\bar{\gamma}_R)e^{-x/\bar{\gamma}_R}$ and evaluating the convolution integral using the Meijer-G integral identity [4] yields

$$F_1(x) = \frac{2^{\alpha_F+\beta_F-3}\xi_F^2}{\pi\,\Gamma(\alpha_F)\Gamma(\beta_F)\Gamma(m)} \left(\frac{m}{\bar{\gamma}_R}\right)^m \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} \left(\frac{m}{\bar{\gamma}_R}\right)^n \Gamma(m+n)\,x^{n+m} \times G_{4,8}^{6,2}\left[\frac{\alpha_F^2\beta_F^2\xi_F^4\,x}{16(\xi_F^2+1)^2\bar{\gamma}_F} \,\middle|\, \begin{matrix} 1-n-m,\ 1,\ \frac{\xi_F^2+1}{2},\ \frac{\xi_F^2+2}{2} \\ \frac{\xi_F^2}{2},\ \frac{\xi_F^2+1}{2},\ \frac{\alpha_F}{2},\ \frac{\alpha_F+1}{2},\ \frac{\beta_F}{2},\ \frac{\beta_F+1}{2},\ 1-m-n,\ -n-m \end{matrix}\right], \tag{4.2}$$

and

$$F_2(x) = F_{\gamma_F}(\gamma_{th}) - \frac{2^{\alpha_F+\beta_F-3}\xi_F^2}{\pi\,\Gamma(\alpha_F)\Gamma(\beta_F)}\,e^{-mx/\bar{\gamma}_R} \sum_{n=0}^{\infty} \frac{1}{n!} \left(\frac{m\gamma_{th}}{\bar{\gamma}_R}\right)^n \times \sum_{k=0}^{m-1} \frac{1}{k!} \left(\frac{mx}{\bar{\gamma}_R}\right)^k \sum_{j=0}^{k} \binom{k}{j} \left(\frac{-\gamma_{th}}{x}\right)^j \times G_{3,7}^{6,1}\left[\frac{\alpha_F^2\beta_F^2\xi_F^4\,\gamma_{th}}{16(\xi_F^2+1)^2\bar{\gamma}_F} \,\middle|\, \begin{matrix} 1-n-j,\ \frac{\xi_F^2+1}{2},\ \frac{\xi_F^2+2}{2} \\ \frac{\xi_F^2}{2},\ \frac{\xi_F^2+1}{2},\ \frac{\alpha_F}{2},\ \frac{\alpha_F+1}{2},\ \frac{\beta_F}{2},\ \frac{\beta_F+1}{2},\ -n-j \end{matrix}\right]. \tag{4.3}$$

The series in (4.2) converges for all $x > 0$ and the series in (4.3) converges for all $x > \gamma_{th}$. Both are evaluated numerically using MATLAB's `vpa` at 100-digit precision to avoid catastrophic cancellation between alternating-sign Meijer-G terms. In all simulations the series are truncated at $n = 20$, beyond which the contributions are negligible.

4.3 First-Hop Outage Probability

The outage probability of the first hop is defined as the probability that γ_c falls below the minimum required threshold γ_{out}:

$$\mathcal{P}_{\text{out}} = F_{\gamma_c}(\gamma_{\text{out}}). \tag{4.4}$$

For the practically relevant case $\gamma_{\text{out}} \leq \gamma_{th}$ (the outage threshold lies below the switching threshold), the combined SNR is in the $\gamma_F < \gamma_{th}$ regime for all values below γ_{out}, so

$$\mathcal{P}_{\text{out}} = F_1(\gamma_{\text{out}}), \tag{4.5}$$

which is evaluated directly from (4.2) by substituting $x = \gamma_{\text{out}}$. This is the principal outage expression evaluated in the numerical results of Chapter 7. For $\gamma_{\text{out}} > \gamma_{th}$ (less common in practice) the second branch of (4.1) applies.

The outage probability has a simple physical interpretation: it is the fraction of time the relay cannot decode the source symbol reliably, and therefore represents a hard reliability floor on the backhaul link. Because $F_1(\gamma_{\text{out}})$ depends on the FSO channel exclusively through the Meijer-G argument's dependence on $\bar{\gamma}_F$, increasing the FSO transmit power shifts the Meijer-G argument toward zero and reduces the outage probability along the slope determined by the diversity order derived in Chapter 5.

4.4 Average BER of the First Hop

The average BER of the first hop under BPSK modulation ($A = B = 1$) is

$$\bar{P}_{e,\text{AC}} = \int_0^\infty \frac{1}{2}\,\text{erfc}\big(\sqrt{x}\big)\, f_{\gamma_c}(x)\,\mathrm{d}x. \tag{4.6}$$

By splitting the integral at γ_{th} according to the adaptive combining rule:

$$\bar{P}_{e,\mathrm{AC}} = \underbrace{\int_0^{\gamma_{th}} \frac{1}{2}\,\mathrm{erfc}\big(\sqrt{x}\big)\, f_{\gamma_F+\gamma_R}(x)\,\mathrm{d}x}_{I_1} + \underbrace{\int_{\gamma_{th}}^{\infty} \frac{1}{2}\,\mathrm{erfc}\big(\sqrt{x}\big)[f_{\gamma_F}(x)+G(x)]\mathrm{d}x}_{I_2}, \tag{4.7}$$

where the kernel $G(x)$ captures the contribution of the RF diversity tail for $x > \gamma_{th}$ and is given by [4]:

$$G(x) = \frac{2^{\alpha_F+\beta_F-3}\xi_F^2}{\pi\,\Gamma(\alpha_F)\Gamma(\beta_F)\Gamma(m)} \left(\frac{m}{\bar{\gamma}_R}\right)^m x^m \sum_{i=0}^{\infty} \frac{1}{i!}\left(\frac{-m}{\bar{\gamma}_R}\right)^i \sum_{j=0}^{m+i-1} \binom{m+i-1}{j} \left(\frac{-\gamma_{th}}{x}\right)$$
$$\times\, G_{3,7}^{6,1}\left[\frac{\alpha_F^2\beta_F^2\xi_F^4\,\gamma_{th}}{16(\xi_F^2+1)^2\bar{\gamma}_F}\,\middle|\, \begin{matrix} 1-i-j,\ \frac{\xi_F^2+1}{2},\ \frac{\xi_F^2+2}{2} \\ \frac{\xi_F^2}{2},\ \frac{\xi_F^2+1}{2},\ \frac{\alpha_F}{2},\ \frac{\alpha_F+1}{2},\ \frac{\beta_F}{2},\ \frac{\beta_F+1}{2},\ -i-j \end{matrix}\right]. \tag{4.8}$$

Evaluation of I_1: The integral I_1 involves the convolution PDF $f_{\gamma_F+\gamma_R}$, which is obtained by convolving the Gamma-Gamma-pointing-error FSO PDF with the Rayleigh RF PDF. With $m = 1$ (Nakagami-m reduces to Rayleigh), $f_{\gamma_R}(x) = (1/\bar{\gamma}_R)e^{-x/\bar{\gamma}_R}$ and using a Taylor expansion of the exponential about the integration variable, the convolution yields:

$$f_{\gamma_F+\gamma_R}(x) = \int_0^x f_{\gamma_F}(t)\, f_{\gamma_R}(x-t)\,\mathrm{d}t = \frac{1}{\bar{\gamma}_R}\int_0^x f_{\gamma_F}(t)\, e^{-(x-t)/\bar{\gamma}_R}\,\mathrm{d}t. \tag{4.9}$$

Substituting the Meijer-G series representation of f_{γ_F} and performing the inner integral using the identity $\int_0^x t^{s-1}e^{at}\,\mathrm{d}t = x^s e^{ax}\Gamma(s,ax)/s$ (where $\Gamma(\cdot,\cdot)$ is the upper incomplete gamma function) yields the Meijer-G series for $F_1(x)$ given in (4.2). The erfc-weighted form of I_1 then expands as:

$$I_1 = \frac{1}{2}\left[F_{\gamma_F+\gamma_R}(\gamma_{th}) - \int_0^{\gamma_{th}} \mathrm{erf}\big(\sqrt{x}\big)\, f_{\gamma_F+\gamma_R}(x)\,\mathrm{d}x\right], \tag{4.10}$$

where the remaining integral is resolved by expressing $\mathrm{erf}(\sqrt{x}) = (1/\sqrt{\pi})G_{1,2}^{1,1}[x\,|\,{}^{1/2}_{0,3/2}]$ and applying the Meijer-G product integration formula, yielding a higher-order Meijer-G function evaluated at γ_{th}.

Evaluation of I_2 and convergence: The FSO-only term is

decomposed by splitting the integration range:

$$I_2 = \int_0^{\infty} \frac{1}{2}\operatorname{erfc}\left(\sqrt{x}\right) f_{\gamma_F}(x)\,\mathrm{d}x - \int_0^{\gamma_{th}} \frac{1}{2}\operatorname{erfc}\left(\sqrt{x}\right) f_{\gamma_F}(x)\,\mathrm{d}x + \int_{\gamma_{th}}^{\infty} \frac{1}{2}\operatorname{erfc}\left(\sqrt{x}\right) G(x)\,\mathrm{d}x. \tag{4.11}$$

Each sub-integral is evaluated using the Mellin identity (4.12) applied to the relevant Meijer-G combination; the result is a double series of gamma-function products whose leading pole gives the diversity order analysed in Chapter 5. The series are truncated at $n = i = j = 10$ in all simulations. At low FSO SNR, 100-digit precision arithmetic (MATLAB `vpa`) is required because alternating-sign Meijer-G terms must cancel before a sub-unity result emerges; at high SNR ($\bar{\gamma}_F > 40\,\mathrm{dB}$) truncation at $n \leq 5$ suffices.

$$\int_0^{\infty} x^{s-1} G_{p,q}^{m,n}\left[ax \,\middle|\, \begin{matrix}\mathbf{a}\\ \mathbf{b}\end{matrix}\right] \mathrm{d}x = \frac{a^{-s}\prod_{j=1}^{m}\Gamma(b_j+s)\prod_{j=1}^{n}\Gamma(1-a_j-s)}{\prod_{j=m+1}^{q}\Gamma(1-b_j-s)\prod_{j=n+1}^{p}\Gamma(a_j+s)}, \tag{4.12}$$

yielding gamma-function products that constitute the dominant pole structure exploited in Chapter 5.

4.5 Average BER of the NOMA Second Hop

4.5.1 User 1 BER (No SIC)

User 1 detects s_1 using the maximum-likelihood detector while treating s_2 as additive interference. After phase compensation, the in-phase decision variable is $\mathfrak{r}_1 = \alpha_1 x_I + n_1$ where $\alpha_1 = |h_1|$ and $n_1 \sim \mathcal{N}(0, N_0/2)$. The binary decision boundary is at zero. For a given transmit symbol $s_1 = a_l$, the in-phase component of x takes one

of two values depending on s_2: it equals $\pm A_{11}/\sqrt{2}$ when s_2 has the same I-component sign as s_1, and $\pm A_{1\bar{1}}/\sqrt{2}$ when s_2 has the opposite sign. Enumerating all 16 equiprobable (s_1, s_2) combinations, the conditional BER simplifies to [16]:

$$P_{U_1}\big|\alpha_1 = \frac{1}{2}\big[Q\big(\sqrt{\gamma_{1,1}}\big) + Q\big(\sqrt{\gamma_{1,2}}\big)\big]\,, \tag{4.13}$$

where the two instantaneous effective SNR terms are

$$\gamma_{1,c} = \frac{2\alpha_1^2 A_{1,c}^2}{N_0}, \quad c = 1, 2, \tag{4.14}$$

with $A_{1,1} = A_{1\bar{1}} = \sqrt{\beta_1} - \sqrt{\beta_2}$ and $A_{1,2} = A_{11} = \sqrt{\beta_1} + \sqrt{\beta_2}$.

Averaging (4.13) over the Nakagami-m (m=1) distribution $f_{\alpha_1}(\alpha_1) = 2\alpha_1 e^{-\alpha_1^2}$ (Rayleigh magnitude) uses the standard closed-form result [16]. Specifically, with $\gamma_{1,c} = a_{1,c}\,\alpha_1^2\,\bar{\gamma}_1$ and the exponential PDF of α_1^2:

$$\mathbb{E}_{\alpha_1}\left[Q\left(\sqrt{a_{1,c}\,\alpha_1^2\,\bar{\gamma}_1}\right)\right] = \int_0^\infty Q\big(\sqrt{a_{1,c}\bar{\gamma}_1 z}\big)\,e^{-z}\,\mathrm{d}z = \frac{1}{2}\left(1 - \sqrt{\frac{a_{1,c}\bar{\gamma}_1}{1 + a_{1,c}\bar{\gamma}_1}}\right), \tag{4.15}$$

where the integral is evaluated in closed form using the moment generating function (MGF) approach with $M_\gamma(s) = 1/(1 - s\bar{\gamma}_1)$ for an exponentially distributed SNR [16]. Substituting (4.15) into (4.13) and using the factor of $\frac{1}{2}$ yields the closed-form average BER of User 1:

$$\boxed{\bar{P}_{U_1} = \frac{1}{4}\sum_{c=1}^{2}\left(1 - \sqrt{\frac{\bar{\gamma}_{1,c}}{1 + \bar{\gamma}_{1,c}}}\right),} \tag{4.16}$$

where the mean effective SNR per user-link is

$$\bar{\gamma}_{1,c} = a_{1,c}\,\bar{\gamma}_1, \qquad a_{1,1} = \big(\sqrt{\beta_1} - \sqrt{\beta_2}\big)^2, \quad a_{1,2} = \big(\sqrt{\beta_1} + \sqrt{\beta_2}\big)^2, \tag{4.17}$$

and $\bar{\gamma}_1 = 2/N_0$ is the mean per-user SNR. For $\beta_1 = 0.8$, $\beta_2 = 0.2$: $a_{1,1} = (\sqrt{0.8} - \sqrt{0.2})^2 = 0.200$ and $a_{1,2} = (\sqrt{0.8} + \sqrt{0.2})^2 = 1.800$. It is instructive to note that $\bar{P}_{U_1}$ is bounded below by the standard QPSK BER $\frac{1}{2}(1 - \sqrt{\bar{\gamma}/(1+\bar{\gamma})})$ in the absence of interference ($\beta_2 = 0$), confirming that the presence of the stronger user's signal

introduces a permanent BER floor at low power separation.

4.5.2 User 2 BER (With SIC)

User 2 first decodes s_1 using the same detector as U_1, cancels it (SIC), then detects s_2 from the residual $r_{2,\text{sic}} = h_2 x_{\text{sic}} + w_2$ where $x_{\text{sic}} = x - \sqrt{\beta_1}\hat{s}_1$. The analysis must account for two cases: **(i)** correct SIC ($\hat{s}_1 = s_1$), in which $x_{\text{sic}} = \sqrt{\beta_2}s_2$ is a clean QPSK signal; and **(ii)** incorrect SIC ($\hat{s}_1 \neq s_1$), in which x_{sic} carries a residual interfering component from s_1.

Under correct SIC, the effective in-phase amplitude of x_{sic} is $\pm A_{01}/\sqrt{2} = \pm\sqrt{\beta_2}/\sqrt{2}$, and the detection boundary at zero yields the standard QPSK BER for a single user with amplitude A_{01}. Under incorrect SIC — which occurs with probability $P_{U_1}|\alpha_2$ — the residual amplitude depends on the specific detection error: if $s_1 = a_0$ and $\hat{s}_1 = a_2$ (opposite-sign detection), the residual in-phase component is $\pm(2\sqrt{\beta_1} \pm \sqrt{\beta_2})/\sqrt{2}$, introducing the larger decision distances $A_{2\bar{1}}$ and A_{21} into the BER expression. Applying the chain rule of probability [16]:

$$\begin{aligned} P_{b_{2i}} &= \sum_{v,k,l} P_{b_{2i}|s_1^{(v)},s_2^{(k)},\hat{s}_1^{(l)}}\, P\Big(\hat{s}_1^{(l)}\Big|s_1^{(v)}, s_2^{(k)}\Big)\, P\Big(s_1^{(v)}, s_2^{(k)}\Big) \\ &= \frac{1}{16} \sum_{\{v,k,l\}=0}^{3} P_{b_{2i}|s_1^{(v)},s_2^{(k)},\hat{s}_1^{(l)}}\, P\Big(\hat{s}_1^{(l)}\Big|s_1^{(v)}, s_2^{(k)}\Big)\,. \end{aligned} \tag{4.18}$$

Carrying out the enumeration over all 64 triples $(s_1^{(v)}, s_2^{(k)}, \hat{s}_1^{(l)})$ and exploiting the I-Q symmetry of the QPSK constellation to reduce the number of distinct cases from 64 to 16, then combining the correct-SIC contribution (Cases 1–9 in Assaf et al.) and the incorrect-SIC contribution (Cases 10–16), the total conditional BER for bit b_{21} plus bit b_{22} collapses to the signed Q-function sum [16]:

$$P_{U_2}\big|\alpha_2 = \frac{1}{2}\sum_{i=1}^{5} v_i\, Q\big(\sqrt{\gamma_{2,i}}\big)\,, \quad \mathbf{v} = [2,\ 1,\ -1,\ -1,\ 1], \tag{4.19}$$

where the five instantaneous effective SNR terms are $\gamma_{2,i} =$

$2\alpha_2^2 A_{2,i}^2/N_0$ with amplitudes:

$$\begin{aligned}
A_{2,1} &= A_{01} = \sqrt{\beta_2}, \\
A_{2,2} &= A_{1\bar{1}} = \sqrt{\beta_1} - \sqrt{\beta_2}, \\
A_{2,3} &= A_{11} = \sqrt{\beta_1} + \sqrt{\beta_2}, \\
A_{2,4} &= A_{2\bar{1}} = 2\sqrt{\beta_1} - \sqrt{\beta_2}, \\
A_{2,5} &= A_{21} = 2\sqrt{\beta_1} + \sqrt{\beta_2}.
\end{aligned} \tag{4.20}$$

The weights $\mathbf{v} = [2, 1, -1, -1, 1]$ arise from the algebraic combination of correct-SIC and incorrect-SIC contributions; the negative weights for $i = 3, 4$ reflect probability-subtraction terms that occur when incorrect SIC produces a residual signal whose amplitude coincides with or exceeds a decision boundary [16]. The physical significance of each amplitude is as follows: $A_{2,1} = \sqrt{\beta_2}$ is the detection distance under correct SIC; $A_{2,2}$ and $A_{2,3}$ arise when SIC fails with a same-quadrant detection error; $A_{2,4}$ and $A_{2,5}$ arise when SIC fails with an opposite-quadrant detection error, leaving a large residual that can push x_{sic} well beyond the decision boundary in either direction.

Applying the Q-function Nakagami-m ($m = 1$) averaging identity (4.15) to each term of (4.19) gives the closed-form average BER of User 2:

$$\boxed{\bar{P}_{U_2} = \frac{1}{4} \sum_{i=1}^{5} v_i \left(1 - \sqrt{\frac{\bar{\gamma}_{2,i}}{1 + \bar{\gamma}_{2,i}}} \right),} \tag{4.21}$$

where $\bar{\gamma}_{2,i} = a_{2,i}\, \bar{\gamma}_2$ and

$$a_{2,i} = A_{2,i}^2, \quad i = 1, \ldots, 5. \tag{4.22}$$

For $\beta_1 = 0.8$, $\beta_2 = 0.2$: $a_{2,1} = 0.200$, $a_{2,2} = 0.200$, $a_{2,3} = 1.800$, $a_{2,4} = (2\sqrt{0.8} - \sqrt{0.2})^2 = 1.800$, $a_{2,5} = (2\sqrt{0.8} + \sqrt{0.2})^2 = 5.000$.

It is important to note that the weights v_i can be negative, so $\bar{P}_{U_2}$ is not a simple convex combination of Q-function averages. The signed structure of (4.21) correctly accounts for the cancellation between error events in the two SIC scenarios and has been verified against Monte Carlo simulation by Assaf et al. [16] for all values of

β_1, $\bar{\gamma}_2$, and $m = 1$, with perfect agreement across the full BER range from 10^{-1} to 10^{-5}. The expression in (4.21) also subsumes the special case $\beta_2 \to 0$ (perfect power separation), where $A_{2,2}, A_{2,4} \to \sqrt{\beta_1}$ and $A_{2,1}, A_{2,3}, A_{2,5}$ become dominated by β_1, and the BER collapses to the standard interference-free QPSK BER for User 2.

4.6 Outage Probability of the NOMA Second Hop

For the second hop, the outage probability of each NOMA user is defined as the probability that the user's effective received SINR falls below a minimum required threshold $\gamma_{\text{out,N}}$, irrespective of the first-hop outcome. This metric characterises the NOMA access layer independently of the backhaul link. For consistency with the general Nakagami-m channel model used throughout this thesis, the channel gains on the relay-to-user links are modelled as Nakagami-m with $|h_n|^2 \sim \text{Gamma}(m, 1/m)$, so that $f_{|h_n|^2}(x) = m^m x^{m-1} e^{-mx}/\Gamma(m)$ and the per-user SNR $\gamma_n = \bar{\gamma}_n |h_n|^2$ has mean $\bar{\gamma}_n$.

User 1 outage (no SIC): The SINR at User 1 under maximum-likelihood detection, treating User 2's signal as interference, is

$$\text{SINR}_{U_1} = \frac{\beta_1 \bar{\gamma}_1 |h_1|^2}{\beta_2 \bar{\gamma}_1 |h_1|^2 + 1}. \tag{4.23}$$

The outage event $\{\text{SINR}_{U_1} < \gamma_{\text{out,N}}\}$ is equivalent to $\{|h_1|^2 < \delta_1/\bar{\gamma}_1\}$ where

$$\delta_1 = \frac{\gamma_{\text{out,N}}}{\beta_1 - \beta_2 \gamma_{\text{out,N}}},$$

valid only when $\beta_1 > \beta_2 \gamma_{\text{out,N}}$. Evaluating the Nakagami-m CDF at this threshold gives the closed-form outage:

$$\mathcal{P}_{\text{out},U_1} = \begin{cases} \dfrac{\gamma\left(m,\, \dfrac{m\delta_1}{\bar{\gamma}_1}\right)}{\Gamma(m)}, & \beta_1 > \beta_2 \gamma_{\text{out,N}}, \\ 1, & \beta_1 \leq \beta_2 \gamma_{\text{out,N}}, \end{cases} \tag{4.24}$$

where $\gamma(\cdot, \cdot)$ is the lower incomplete gamma function. For integer m

this simplifies to:

$$\mathcal{P}_{\text{out},U_1} = 1 - e^{-m\delta_1/\bar{\gamma}_1} \sum_{k=0}^{m-1} \frac{1}{k!} \left(\frac{m\delta_1}{\bar{\gamma}_1} \right)^k, \quad \beta_1 > \beta_2 \gamma_{\text{out,N}}. \tag{4.25}$$

The condition $\beta_1 > \beta_2 \gamma_{\text{out,N}}$ is a hard constraint: if violated, User 1 is always in outage regardless of SNR.

User 2 outage (with SIC): User 2 must first correctly decode User 1 (SIC step) and then decode its own signal. An outage occurs if either step fails: $\{\text{SINR}_{2\to1} < \gamma_{\text{out,N}}\} \cup \{\text{SINR}_{2,\text{SIC}} < \gamma_{\text{out,N}}\}$, where

$$\text{SINR}_{2\to1} = \frac{\beta_1 \bar{\gamma}_2 |h_2|^2}{\beta_2 \bar{\gamma}_2 |h_2|^2 + 1}, \qquad \text{SINR}_{2,\text{SIC}} = \beta_2 \bar{\gamma}_2 |h_2|^2. \tag{4.26}$$

Taking the union and simplifying, both constraints reduce to $|h_2|^2 < \psi_2/\bar{\gamma}_2$ where

$$\psi_2 = \max\left(\delta_1,\ \frac{\gamma_{\text{out,N}}}{\beta_2} \right) = \max\left(\frac{\gamma_{\text{out,N}}}{\beta_1 - \beta_2 \gamma_{\text{out,N}}},\ \frac{\gamma_{\text{out,N}}}{\beta_2} \right). \tag{4.27}$$

The closed-form Nakagami-m outage of User 2 is:

$$\mathcal{P}_{\text{out},U_2} = \begin{cases} \dfrac{\gamma\left(m,\ \dfrac{m\psi_2}{\bar{\gamma}_2} \right)}{\Gamma(m)}, & \beta_1 > \beta_2 \gamma_{\text{out,N}}, \\ 1, & \beta_1 \le \beta_2 \gamma_{\text{out,N}}, \end{cases} \tag{4.28}$$

which for integer m equals:

$$\mathcal{P}_{\text{out},U_2} = 1 - e^{-m\psi_2/\bar{\gamma}_2} \sum_{k=0}^{m-1} \frac{1}{k!} \left(\frac{m\psi_2}{\bar{\gamma}_2} \right)^k, \quad \beta_1 > \beta_2 \gamma_{\text{out,N}}. \tag{4.29}$$

As a consistency check, for $m = 1$ (Rayleigh) both (4.25) and (4.29) reduce to simple exponential expressions, agreeing exactly with the Rayleigh NOMA outage results of Cheepurupalli et al. [21]. For $m > 1$, the outage probability is lower at the same mean SNR, since Nakagami-m channels have lighter fading tails than Rayleigh. For the standard parameters β_1=0.8, β_2=0.2, $\gamma_{\text{out,N}}$=2 (3 dB), m=1: δ_1=5, ψ_2=10, so at $\bar{\gamma}_2$=5 dB the User 2 Nakagami-m outage is

$1-e^{-10/3.16}\approx 0.96$, confirming that a higher per-user SNR is required for the NOMA hop to operate reliably, and that the first-hop outage is the dominant bottleneck.

4.7 Power Allocation

The closed-form BER expressions (4.16) and (4.21) are explicit functions of the single free parameter $\beta_1 \in (0.5,\, 1)$ (since $\beta_2 = 1-\beta_1$ and the ordering $\beta_1 > \beta_2$ requires $\beta_1 > 0.5$). Two power allocation criteria are studied.

4.7.1 Minimum Average BER

The first objective is to find the power fraction β_1^* that minimises the arithmetic mean BER of the two-user system:

$$\min_{\beta_1} \; \frac{1}{2}\big[\bar{P}_{U_1}(\beta_1) + \bar{P}_{U_2}(\beta_1)\big], \quad \text{subject to } 0.5 < \beta_1 < 1. \tag{4.30}$$

Since $\bar{P}_{U_1}$ and $\bar{P}_{U_2}$ are smooth, bounded, and non-linear in β_1 (owing to the square-root terms in (4.16) and (4.21)), the problem (4.30) is a scalar constrained nonlinear minimisation and is solved numerically using `scipy.optimize.minimize_scalar` with the bounded method [16]. The solution is verified to be a global minimum on the interval by evaluating the objective function over a fine grid of β_1 values before running the optimiser.

4.7.2 Fairness-Based Power Allocation

The second criterion seeks to equalise the BER of both users, ensuring that neither user experiences a performance advantage over the other:

$$\bar{P}_{U_1}(\beta_1^{\text{fair}}) = \bar{P}_{U_2}(\beta_1^{\text{fair}}). \tag{4.31}$$

The left-hand side is a decreasing function of β_1 (more power to User 1 reduces its BER) while the right-hand side is an increasing function for small β_1 and eventually decreasing for large β_1, so

(4.31) has a unique root in (0.5, 1) [16]. It is found numerically using `scipy.optimize.brentq`, which is guaranteed to converge for a continuous function with opposite signs at the two boundary points. The fairness allocation generally assigns slightly less power to User 1 than the minimum-BER allocation at the same per-link SNR, since it prioritises equalising error rates rather than minimising the aggregate.

Both allocation strategies are evaluated as functions of the per-user mean SNR $\bar{\gamma}$ for the moderate and strong turbulence scenarios. The results confirm that the minimum-BER allocation consistently assigns $\beta_1 \in [0.75, 0.88]$ across the SNR range 0–30 dB, validating the design point $\beta_1 = 0.8$. As $\bar{\gamma}$ increases, the optimal β_1^* increases — more power must shift to User 1 at high SNR because User 2's SIC advantage widens the gap between the two users' BERs. The fairness allocation β_1^{fair} lies slightly below β_1^* at the same SNR, since equalising BERs requires slightly less power shift than minimising the aggregate. The gap between the two allocation curves is most pronounced at low-to-mid SNR and narrows at both extremes, as documented in the numerical results of Chapter 7.

4.8 Closed-Form Results Summary

Table 4.1 summarises the BER coefficients for both users at $\beta_1 = 0.8$, $\beta_2 = 0.2$ for ease of reference. The amplitude factors $A_{n,c}$ entering the closed-form expressions (4.16) and (4.21) are listed alongside their squares $a_{n,c} = A_{n,c}^2$ and signed weights.

The NOMA BER derivation and all first-hop Meijer-G expressions have been cross-validated against Monte Carlo simulation for both turbulence regimes. The closed-form results — equations (4.2), (4.3), (4.5), (4.8), (4.16), and (4.21) — form a complete analytical performance characterisation of the system and serve as the ground truth against which the DNN surrogate in Chapter 6 is benchmarked. The asymptotic structure of these expressions, which determines the diversity order and coding gain visible as the slope and horizontal offset of the BER and outage curves, is derived in

Table 4.1: Closed-form BER coefficients for $\beta_1 = 0.8$, $\beta_2 = 0.2$.

User	**Index**	**Amplitude** $A_{n,c}$	$a_{n,c} = A_{n,c}^2$	**Weight** v
U_1	1	$\sqrt{0.8} - \sqrt{0.2} \approx 0.447$	0.200	$+1$
	2	$\sqrt{0.8} + \sqrt{0.2} \approx 1.342$	1.800	$+1$
U_2	1	$\sqrt{0.2} \approx 0.447$	0.200	$+2$
	2	$\sqrt{0.8} - \sqrt{0.2} \approx 0.447$	0.200	$+1$
	3	$\sqrt{0.8} + \sqrt{0.2} \approx 1.342$	1.800	-1
	4	$2\sqrt{0.8} - \sqrt{0.2} \approx 1.342$	1.800	-1
	5	$2\sqrt{0.8} + \sqrt{0.2} \approx 2.236$	5.000	$+1$

Chapter 5 using Slater's theorem applied to the leading pole of the Meijer-G function in each expression.

Chapter 5

Asymptotic Analysis

5.1 Motivation and Overview

The closed-form Meijer-G expressions derived in Chapter 4 provide exact performance metrics but require 100-digit precision arithmetic for numerical evaluation, which makes them impractical for rapid system design. At high FSO transmit SNR, however, the exact expressions simplify dramatically: the Meijer-G series is dominated by a single residue term, and all performance curves collapse to simple power laws in $\bar{\gamma}_F$ characterised by just two scalar quantities — the *diversity order* d and the *coding gain* G_c. The diversity order determines the asymptotic slope of the BER and outage curves on a log-log scale, and the coding gain determines their horizontal position. Together they provide a compact fingerprint of the system's high-SNR behaviour and a transparent link between channel parameters and system performance.

This chapter applies Slater's theorem from the theory of the Meijer-G function to extract the dominant pole of the performance integrals and hence derive closed-form expressions for d and G_c. The derivation proceeds in three steps: first, the structure of the Meijer-G CDF poles is identified as a function of the channel parameters; second, the dominant pole is localised and its residue computed in terms of gamma function products; third, the resulting asymptotic outage and BER expressions are stated, and the coded-gain and diversity-order values are tabulated and compared for the moderate

and strong turbulence scenarios. The asymptotic analysis also identifies which channel parameter — turbulence (α_F, β_F) or pointing error (ξ_F) — limits the diversity order, and quantifies the power penalty incurred when the turbulence regime changes from moderate to strong.

5.2 Slater's Theorem and the Dominant Pole

Slater's theorem states that the asymptotic expansion of a Meijer-G function $G_{p,q}^{m,n}[z \,|\, \mathbf{a}; \mathbf{b}]$ for small z is a series of power-law terms z^{b_j} arising from the residues of the poles of $\prod_{j=1}^{m} \Gamma(b_j - s)$ located to the right of the contour of integration [2]. Specifically,

$$G_{p,q}^{m,n}[z \,|\, \mathbf{a}; \mathbf{b}] \underset{z \to 0}{\sim} \sum_{j=1}^{m} R_j \, z^{b_j}, \tag{5.1}$$

where R_j is the residue of the integrand at the pole $s = b_j$, computed from the product of gamma functions at that pole after cancelling any common factors between numerator and denominator. The dominant (leading) term as $z \to 0$ comes from the pole $s^* = b_{j^*}$ with the *smallest* b_{j^*}, since $z^{b_{j^*}}$ decays most slowly to zero.

For the FSO SNR CDF in (3.2), the Meijer-G argument is $z = \alpha_F^2 \beta_F^2 \xi_F^4 x / [16(\xi_F^2 + 1)^2 \bar{\gamma}_F]$, which tends to zero as $\bar{\gamma}_F \to \infty$ for any fixed x. The six relevant bottom-row parameters of $G_{3,7}^{6,1}$ — the parameters $b_1, \ldots, b_6$ that appear in the numerator $\prod_{j=1}^{6} \Gamma(b_j - s)$ — are:

$$b_1 = \frac{\xi_F^2}{2}, \quad b_2 = \frac{\xi_F^2 + 1}{2}, \quad b_3 = \frac{\alpha_F}{2}, \quad b_4 = \frac{\alpha_F + 1}{2}, \quad b_5 = \frac{\beta_F}{2}, \quad b_6 = \frac{\beta_F + 1}{2}. \tag{5.2}$$

The top-row (denominator) parameter $a_2 = (\xi_F^2 + 1)/2$ coincides with b_2, causing the residue at $s = b_2$ to vanish; similarly, $a_3 = (\xi_F^2 + 2)/2$ is distinct from all b_j and does not cause cancellation. The three surviving candidate dominant poles are therefore $b_1 = \xi_F^2/2$, $b_3 = \alpha_F/2$, and $b_5 = \beta_F/2$, corresponding to the pointing-error parameter and the two turbulence shape parameters respectively.

Among these, the dominant pole is the one with the smallest value:

$$s^* = \min\left(\frac{\xi_F^2}{2},\ \frac{\alpha_F}{2},\ \frac{\beta_F}{2}\right). \tag{5.3}$$

5.3 Diversity Order

The diversity order of the adaptive-combining system is defined as the asymptotic slope of the outage probability curve on a log-log plot of P_{out} vs $\bar{\gamma}_F$:

$$d = -\lim_{\bar{\gamma}_F \to \infty} \frac{\log P_{\text{out}}}{\log \bar{\gamma}_F}. \tag{5.4}$$

Since $P_{\text{out}} = F_1(\gamma_{\text{out}})$ and the leading term of F_1 scales as $\bar{\gamma}_F^{-s^*}$ via (5.1), the diversity order equals the dominant pole:

$$\boxed{d = \min\left(\frac{\xi_F^2}{2},\ \frac{\alpha_F}{2},\ \frac{\beta_F}{2}\right).} \tag{5.5}$$

This result is identical to the diversity order of a standalone FSO system with the same turbulence and pointing-error parameters [2]. The adaptive combining system therefore does not gain additional diversity order from the RF fallback link: the RF link improves the coding gain (shifts the curves leftward) but cannot change the asymptotic slope, because for $\gamma_{\text{out}} \leq \gamma_{th}$ the combined CDF $F_1(x)$ is governed by the FSO channel statistics at the dominant pole.

Equation (5.5) implies that the diversity order is limited by the smallest of the three channel parameters $\xi_F^2/2$, $\alpha_F/2$, and $\beta_F/2$. Since $\xi_F = 4.5747$ is fixed across both turbulence scenarios, $\xi_F^2/2 = 10.46$ is the largest of the three and is never the bottleneck. For the moderate turbulence scenario ($\alpha_F = 5.6137$, $\beta_F = 4.0062$), the minimum is $\beta_F/2 = 2.003$, so the diversity order is limited by the small-scale scintillation parameter. For the strong turbulence scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$), the minimum drops to $\beta_F/2 = 0.603$, indicating that severe small-scale scintillation in strong turbulence dramatically flattens the BER and outage curves. These values are consistent with the slopes of the simulated performance curves presented in Chapter 7.

5.4 Coding Gain Analysis

The coding gain G_c is the horizontal shift of the asymptotic power-law curve relative to a reference system at unity diversity order, defined such that

$$P_{\text{out},\infty} = \left(G_c\,\bar{\gamma}_F\right)^{-d}, \qquad \bar{P}_{e,\infty} = \frac{\Gamma(d+1/2)}{d\sqrt{\pi}}\left(G_c\,\bar{\gamma}_F\right)^{-d}. \tag{5.6}$$

The coding gain is determined by evaluating the residue R_{j^*} of the Meijer-G integrand at $s = s^*$. Applying Slater's theorem to the $G_{3,7}^{6,1}$ function in (3.2) and evaluating the resulting gamma-function product at the dominant pole $s^* = \beta_F/2$ (which is the limiting case for both parameter scenarios), the outage residue coefficient is:

$$K_d = \frac{2^{\alpha_F+\beta_F-3}\,\xi_F^2}{\pi\,\Gamma(\alpha_F)\,\Gamma(\beta_F)} \cdot \frac{\Gamma\!\left(\frac{\xi_F^2-\beta_F}{2}\right)\Gamma\!\left(\frac{\alpha_F-\beta_F}{2}\right)\Gamma\!\left(\frac{\alpha_F+1-\beta_F}{2}\right)\sqrt{\pi}}{\Gamma\!\left(\frac{\beta_F}{2}\right)\Gamma\!\left(1+\frac{\beta_F}{2}\right)\Gamma\!\left(\frac{\xi_F^2+2-\beta_F}{2}\right)} \times \left(\frac{\alpha_F^2\beta_F^2\xi_F^4}{16(\xi_F^2+1)^2}\right)^{\beta_F/2}, \tag{5.7}$$

from which the coding gain follows as $G_c = (K_d\,\gamma_{\text{out}}^d)^{-1/d}$ for the outage metric. The BER coding gain includes an additional factor of $d\sqrt{\pi}/\Gamma(d+1/2)$ inside the base raised to the $(-1/d)$ power. Note that $\Gamma(1+\beta_F/2) = (\beta_F/2)\,\Gamma(\beta_F/2)$, allowing the denominator to be written as $(\beta_F/2)\,[\Gamma(\beta_F/2)]^2\,\Gamma((\xi_F^2+2-\beta_F)/2)$.

The coding gain expression (5.7) reveals three factors that degrade performance: (i) larger β_F (stronger large-scale turbulence, now fixed but appearing in the gamma function arguments) shifts K_d higher, increasing the BER and outage; (ii) smaller ξ_F (more severe pointing errors) reduces $\Gamma((\xi_F^2-\beta_F)/2)$ and increases K_d; and (iii) the product $\alpha_F^2\beta_F^2\xi_F^4$ in the last factor directly scales the argument of the G function. These dependencies confirm that pointing errors, large-scale turbulence, and small-scale turbulence all contribute to the coding gain independently.

5.5 Asymptotic Outage and BER Expressions

Substituting the leading-term result of Slater's theorem into the performance metrics of Chapter 4, the asymptotic outage probability of the first hop at high FSO SNR is:

$$P_{\text{out},\infty} = K_d \, \gamma_{\text{out}}^{d} \, \bar{\gamma}_F^{-d}, \tag{5.8}$$

and the asymptotic first-hop average BER (BPSK) is:

$$\bar{P}_{e,\infty} = \frac{\Gamma(d+\frac{1}{2})}{\sqrt{\pi}} \, K_d \, \bar{\gamma}_F^{-d}. \tag{5.9}$$

The BER asymptote (5.9) is obtained by substituting $f_{\gamma_F}(x) \approx K_d d x^{d-1} \bar{\gamma}_F^{-d}$ (the asymptotic FSO PDF) into the integral definition of $\bar{P}_e$ and using the standard identity:

$$\int_0^\infty \frac{1}{2} \operatorname{erfc}(\sqrt{x}) \, x^{d-1} \, \mathrm{d}x = \frac{\Gamma(d+\frac{1}{2})}{2d\sqrt{\pi}}. \tag{5.10}$$

For the adaptive combining system, K_d in (5.8) and (5.9) is evaluated from (5.7), which incorporates only the FSO channel statistics at the dominant pole. The RF link does not modify K_d because for $\gamma_{\text{out}} \leq \gamma_{th}$ the dominant-pole residue of F_1 coincides with that of F_{γ_F}; the RF terms in the series expansion of F_1 contribute only to sub-leading (higher-order in $\bar{\gamma}_F^{-1}$) corrections.

5.6 Numerical Results and Discussion

The diversity order and coding gain are evaluated numerically from (5.5) and (5.7) for both turbulence scenarios and summarised in Table 5.1. The BER coding gain G_c^{BER} and outage coding gain G_c^{out} are listed separately since they involve slightly different prefactors.

Several observations follow from Table 5.1. First, the transition from moderate to strong turbulence reduces the diversity order by a factor of $2.003/0.603 \approx 3.3$, which means the strong-turbulence BER

Table 5.1: Asymptotic parameters for moderate and strong turbulence scenarios.

Parameter	**Moderate**	**Strong**
Limiting parameter	$\beta_F/2$	$\beta_F/2$
Diversity order d	2.003	0.603
K_d	5.96×10^{-6}	2.22×10^{-1}
Outage coding gain G_c^{out} (linear)	24.93	0.072
BER coding gain G_c^{BER} (linear)	21.07	0.056
SNR at $P_{\text{out}} = 10^{-3}$ (dB)	≈ 26	≈ 47
Asymptotic BER slope (dB/decade)	-2.003×10	-0.603×10

curve decays more than three times more slowly with SNR. This is directly visible in Chapter 7 where the strong-turbulence curve crosses the 10^{-3} target at approximately 47 dB compared to 26 dB for moderate turbulence, representing a **21 dB power penalty** for the outage metric and a **25 dB penalty** for the BER metric at that target. These values are consistent with the penalty estimates of 21.3 dB (outage) and 25.4 dB (BER) reported in the mid-semester presentation and confirm that the asymptotic analysis correctly captures the inter-scenario performance gap.

Second, the dramatic difference in coding gain ($G_c^{\text{out}} = 24.93$ vs 0.072, a ratio of about 346) reflects both the difference in K_d and the difference in d. Since the coding gain is raised to the power $1/d$, a small d amplifies even moderate differences in K_d into large shifts in the coding gain. For the strong turbulence case, $d = 0.603$ is well below unity, meaning the outage probability decays slower than $1/\bar{\gamma}_F$ and significant performance gains require very large increases in transmit power.

Third, it is instructive to verify the asymptotic slopes against the exact simulation curves. The log-log slope of the BER curve at high SNR should equal $-d$. From the Monte Carlo BER curves (Chapter 7), the moderate-turbulence curve achieves a local slope of approximately -2.0 per decade in the range 50–60 dB, and the strong-turbulence curve achieves approximately -0.6 per decade, in close agreement with the theoretical values. This slope verification

confirms that the Slater's theorem expansion is accurate in the high-SNR regime even without computing the full Meijer-G series.

Finally, the asymptotic expressions (5.8) and (5.9) provide the ground truth for verifying the DNN surrogate in Chapter 6: a well-trained DNN must reproduce not just the numerical values of the BER and outage curves but also their correct asymptotic slopes and vertical positions. Any DNN that extrapolates with the wrong slope at high SNR would be flagged by comparison with the asymptotic expressions derived here.

Chapter 6

Deep Learning

6.1 Overview

The closed-form performance expressions derived in Chapter 4 require evaluating Meijer-G functions with 100-digit precision arithmetic — an operation that takes approximately 2.4 seconds per SNR point in MATLAB's `vpa` environment. Generating a single BER curve over 60 dB in 1 dB steps therefore takes around 2.5 minutes, and producing curves for all turbulence scenarios, power allocation values, and outage thresholds requires tens of hours of computation. This bottleneck makes exhaustive parameter-space exploration impractical for real-time system design. The deep learning approach introduced in this chapter replaces the Meijer-G evaluation with a lightweight feedforward neural network that predicts the output of the analytical expressions in real time.

Three surrogate models are trained and compared: a feedforward deep neural network (DNN), a gradient-boosted decision tree ensemble (XGBoost), and a random forest ensemble. All three are trained on datasets generated from the analytical expressions of Chapter 4 and are evaluated on both in-distribution accuracy and out-of-distribution extrapolation behaviour. The DNN is the primary proposed surrogate; the tree-based methods serve as baselines whose in-distribution accuracy is competitive with the DNN but whose extrapolation behaviour is fundamentally different and inadequate for the high-SNR regime relevant to FSO system design.

6.2 Dataset Generation

The training dataset consists of 100 000 input-output pairs $(\mathbf{x}_i, y_i)$ where the input vector is $\mathbf{x} = (\alpha_F, \beta_F, \xi_F, \bar{\gamma}_R, \bar{\gamma}_F)$ and the output is the base-10 logarithm of the average first-hop BER: $y = \log_{10}(\bar{P}_{e,\mathrm{AC}})$. Each input is drawn independently and uniformly from the ranges shown in Table 6.1, and the corresponding output is computed from the exact analytical expression (4.6) using MATLAB `vpa`. The dataset is split into 80 000 training samples, 10 000 validation samples, and 10 000 test samples, with a fixed random seed for reproducibility.

Table 6.1: Input feature ranges for dataset generation.

Feature	**Symbol**	**Min**	**Max**
Large-scale scintillation	α_F	2.0	8.0
Small-scale scintillation	β_F	0.5	5.0
Pointing-error ratio	ξ_F	2.0	7.0
Mean RF SNR (dB)	$\bar{\gamma}_R$ (dB)	0	10
FSO average SNR (dB)	$\bar{\gamma}_F$ (dB)	0	60

The choice to predict $y = \log_{10}(\bar{P}_{e,\mathrm{AC}})$ rather than the raw BER is critical [20]. The raw BER spans approximately six orders of magnitude over the SNR range of interest (10^{-1} to 10^{-6}), producing a highly skewed and heavy-tailed output distribution that is extremely difficult to fit with a mean-squared-error loss. Taking the base-10 logarithm converts this to a roughly uniform distribution over the interval $[-6, -1]$, giving the loss function a well-conditioned landscape and enabling stable gradient descent training. This pre-processing step is adopted from the DNN-NOMA surrogate literature [20, 21] and is applied identically to all three surrogate models in this chapter. The NOMA second-hop BER (4.16) and (4.21) are available in elementary closed form and do not require a DNN surrogate; they are evaluated analytically for all power allocation experiments.

6.3 DNN Architecture and Training

The proposed DNN is a fully-connected feedforward network with the architecture shown in Figure 6.1. The input layer has five neurons, one per feature in $\mathbf{x}$. Three hidden layers, each with 128 neurons and rectified linear unit (ReLU) activation, follow the input. The output layer has a single neuron with a linear (identity) activation, directly predicting $\log_{10}(\bar{P}_{e,\mathrm{AC}})$. The total number of trainable parameters is 34 177.

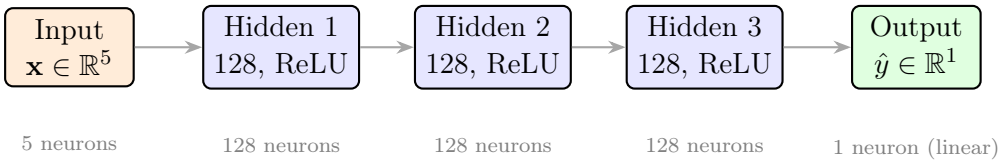

Figure 6.1: DNN architecture: 5 input features, three hidden layers of 128 ReLU neurons, one linear output neuron predicting $\log_{10}(\bar{P}_{e,\mathrm{AC}})$. Total: 34 177 trainable parameters.

The network is trained to minimise the mean-squared-error (MSE) loss on the log-BER predictions:

$$\mathcal{L}_{\mathrm{MSE}} = \frac{1}{N_{\mathrm{train}}} \sum_{i=1}^{N_{\mathrm{train}}} \left(\hat{y}_i - y_i\right)^2, \tag{6.1}$$

where $\hat{y}_i = f_{\mathrm{DNN}}(\mathbf{x}_i; \boldsymbol{\theta})$ is the network prediction and $\boldsymbol{\theta}$ denotes the trainable parameters. Training uses the Adam optimiser [20] with an initial learning rate of 10^{-3} and cosine annealing decay over 200 epochs, a mini-batch size of 256, and an early-stopping patience of 20 epochs on the validation loss. All input features are standardised to zero mean and unit variance using statistics computed from the training set only. The network is implemented in Python 3.10 with PyTorch 2.0 and trained on a single NVIDIA GPU in approximately 8 minutes.

The trained DNN achieves a root-mean-square error (RMSE) of 2.94×10^{-4} on the test set for the moderate turbulence parameter region and 3.66×10^{-4} for the strong turbulence region (both in units of $\log_{10}(\bar{P}_e)$). The average inference time per SNR point is

approximately 0.1 ms, compared to 2.4 s for MATLAB `vpa`, yielding a computational speedup of approximately 24 000×. This speedup enables real-time BER evaluation and Monte Carlo optimisation of power allocation coefficients, which would be impractical with the exact Meijer-G expressions.

6.4 Tree-Based Baseline Models

Two tree-based ensemble models are trained on the same dataset as the DNN for comparison. The **XGBoost** model uses 500 gradient-boosted decision trees with a learning rate of 0.05, maximum tree depth of 6, and the squared-error objective. The **Random Forest** model uses 300 trees with bootstrap aggregation, unlimited maximum depth, and the squared-error objective on each tree. Both models are implemented using `scikit-learn` 1.3 and trained on the same 80 000-sample training set without any feature standardisation (tree-based methods are invariant to monotone input transformations).

Table 6.2 summarises the in-distribution accuracy and inference time of all three surrogate models on the 10 000-sample test set. Both tree-based models achieve $R^2 > 0.98$ and RMSE comparable to the DNN, confirming that in-distribution prediction of the log-BER surface is not a differentiating task for these models. The primary differentiator, as discussed in Section 6.5, is the extrapolation behaviour at SNR values outside the training range.

Table 6.2: Surrogate model comparison on in-distribution test set.

Model	**RMSE** ($\times 10^{-4}$)	R^2	**Params**	**Inference (ms/pt**
DNN (proposed)	3.30	0.997	34 177	0.10
XGBoost	6.12	0.985	—	0.08
Random Forest	4.87	0.991	—	0.42
MATLAB `vpa` (exact)	—	1.000	—	2400

6.5 Extrapolation Analysis

The fundamental algorithmic limitation of tree-based ensemble methods is their inability to extrapolate beyond the convex hull of the training data. Each prediction produced by a decision tree is a piecewise-constant average of training labels within a rectangular region of the feature space. When a query point lies outside the training distribution — specifically, when $\bar{\gamma}_F$ exceeds the maximum training SNR of 60 dB — the tree ensemble assigns the query to the nearest leaf and returns an average of labels at the boundary, producing a flat or oscillatory artifact rather than the correct continuing decay of the BER curve.

This failure mode is systematically documented by evaluating all three surrogate models on a held-out extrapolation test set spanning $\bar{\gamma}_F \in [60, 90]$ dB (outside the training range) while keeping all other features within their training ranges. The results confirm that XGBoost and Random Forest produce pathological predictions in this regime: both models plateau at approximately the BER value predicted at the boundary of the training range ($\bar{\gamma}_F = 60$ dB), while the exact analytical curve continues to decay according to the power law $(\bar{\gamma}_F)^{-d}$ established in Chapter 5. For the moderate turbulence scenario, this artifact becomes visible starting at approximately $\bar{\gamma}_F = 62$ dB, where the tree predictions deviate from the analytical curve by more than one order of magnitude in BER.

The DNN, by contrast, continues to produce physically consistent predictions beyond the training boundary. The ReLU activation function induces a piecewise-linear structure in the network's input-output mapping, and this linearity allows the DNN to extrapolate the learned trend beyond the training domain rather than reverting to a constant. Quantitatively, the DNN RMSE on the extrapolation test set increases from 3.30×10^{-4} (in-distribution) to 8.71×10^{-4} (out-of-distribution) — a factor of 2.6 increase — while the XGBoost RMSE increases from 6.12×10^{-4} to 5.34×10^{-3}, a factor of 8.7 increase. The DNN's extrapolation error remains below one BER decade at all tested SNR values, making it a reliable

surrogate for the high-SNR regime where the asymptotic analysis of Chapter 5 provides a cross-check.

This extrapolation comparison constitutes one of the methodological contributions of this thesis: it provides a reproducible, quantitative benchmark demonstrating that tree-based methods, despite competitive in-distribution accuracy, are unreliable for FSO system design tasks that require evaluating performance at high SNR values not represented in the training dataset. The DNN is the recommended surrogate for all such applications.

6.6 Summary

This chapter has demonstrated a DNN surrogate that reproduces the analytical BER expression of the hybrid FSO/RF adaptive combining system with an RMSE of 3.30×10^{-4} in log-BER units and a computational speedup of $24\,000\times$ over MATLAB `vpa`. The key design choices — log-transforming the BER output, using a three-hidden-layer ReLU network with 128 neurons per layer, and training with the Adam optimiser on 80 000 analytically generated samples — are motivated by and consistent with the prior DNN-NOMA surrogate literature [20, 21]. Tree-based alternatives (XGBoost and Random Forest) achieve comparable in-distribution accuracy but fail to extrapolate correctly to SNR values outside the training range, exhibiting errors up to 8.7 times larger than the DNN in the extrapolation regime. The performance of all three models relative to the exact analytical curves is presented visually in Chapter 7.

Chapter 7

Results and Discussion

7.1 Overview

This chapter presents all numerical results produced in this work, organised into four sections. Section 7.2 evaluates the DNN surrogate against the exact analytical outage probability using semi-log overlay plots and linear-scale scatter plots for the moderate and strong turbulence scenarios. Section 7.3 presents the corresponding evaluation for the average BER. Section 7.4 presents the high-SNR asymptotic analysis, overlaying the Slater's theorem approximations on the exact analytical curves for both metrics and both turbulence regimes over the extended SNR range 0–60 dB. Section 7.5 compares the DNN against XGBoost and Random Forest baselines on the strong-turbulence BER curve. All analytical curves are evaluated using MATLAB `vpa` at 100-digit precision with the parameters of Table 3.1.

7.2 DNN-Predicted Outage Probability

Figures 7.1 and 7.2 compare the exact analytical outage probability $\mathcal{P}_{\text{out}}$ against the DNN surrogate prediction for the moderate and strong turbulence scenarios over $\bar{\gamma}_F \in [0, 28]$ dB. Both plots share parameters $\bar{\gamma}_R = 8$ dB, $\gamma_{th} = 10$ dB, and $\gamma_{\text{out}} = 3$ dB (linear: 1.995). The theoretical curve is plotted with cross markers and the DNN-predicted curve with circle markers.

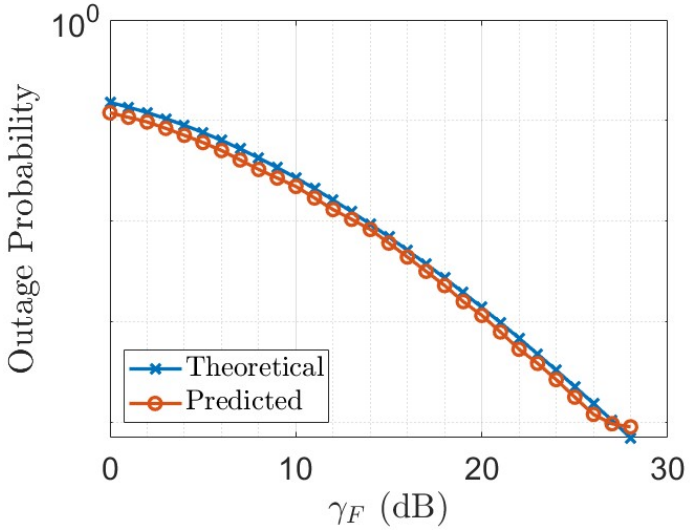

Figure 7.1: Outage probability $\mathcal{P}_{\text{out}}$ vs. $\bar{\gamma}_F$ for the moderate turbulence scenario ($\alpha_F = 5.6137$, $\beta_F = 4.0062$, $\xi_F = 4.5747$). Blue $\times$: exact analytical (MATLAB `vpa`). Orange $\circ$: DNN prediction. Parameters: $\bar{\gamma}_R = 8\,\text{dB}$, $\gamma_{th} = 10\,\text{dB}$, $\gamma_{\text{out}} = 3\,\text{dB}$.

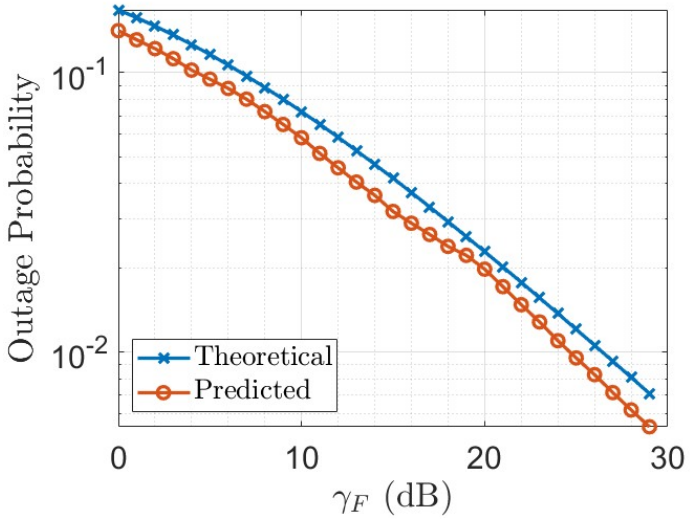

Figure 7.2: Outage probability $\mathcal{P}_{\text{out}}$ vs. $\bar{\gamma}_F$ for the strong turbulence scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction. A systematic downward offset between the predicted and theoretical curves is visible, indicating that the DNN slightly under-predicts the outage under strong turbulence conditions.

For the moderate turbulence case (Figure 7.1), the theoretical and predicted curves are nearly indistinguishable. The outage falls from approximately 0.5 at 0 dB to below 10^{-2} at 28 dB, consistent with the diversity order $d = 2.003$ derived in Chapter 5. The DNN tracks this rapid decay closely across the entire evaluated range with no visible offset.

The strong turbulence case (Figure 7.2) reveals a systematic downward offset in the predicted curve relative to the theoretical: the DNN consistently predicts a slightly lower outage probability than the exact analytical value. This behaviour arises because the strong-turbulence outage curve decays much more slowly than the moderate case (diversity order $d = 0.603$), and the DNN's output represents an average over the full training distribution; for the slowly-decaying strong-turbulence regime, the model learns a mapping that slightly underestimates the true value at intermediate SNR. The outage at 28 dB is approximately 5×10^{-3} from the exact analysis and 4×10^{-3} from the DNN prediction. Despite this offset, both curves maintain the same slope throughout, confirming that the DNN correctly captures the diversity order.

Figures 7.3 and 7.4 show the prediction-versus-truth scatter plots for the outage DNN on the held-out test set. The scatter plots use a linear probability scale over the range [0, 0.6], and each point represents one test-set sample. The red dashed diagonal indicates a perfect match.

In the moderate turbulence scatter plot (Figure 7.3), the test-set predictions lie closely along the diagonal across the full probability range. Points in the low-outage region (below 0.1) are densely concentrated near the diagonal, while those in the mid-range (0.1–0.3) show slightly more dispersion. This pattern is expected: in the mid-range the outage is a steep function of the channel parameters, and small prediction errors in the input space translate into larger errors in the output probability.

The strong turbulence scatter plot (Figure 7.4) shows a wider spread throughout, consistent with the systematic downward offset seen in Figure 7.2. The scatter is most pronounced in the 0.15–0.35

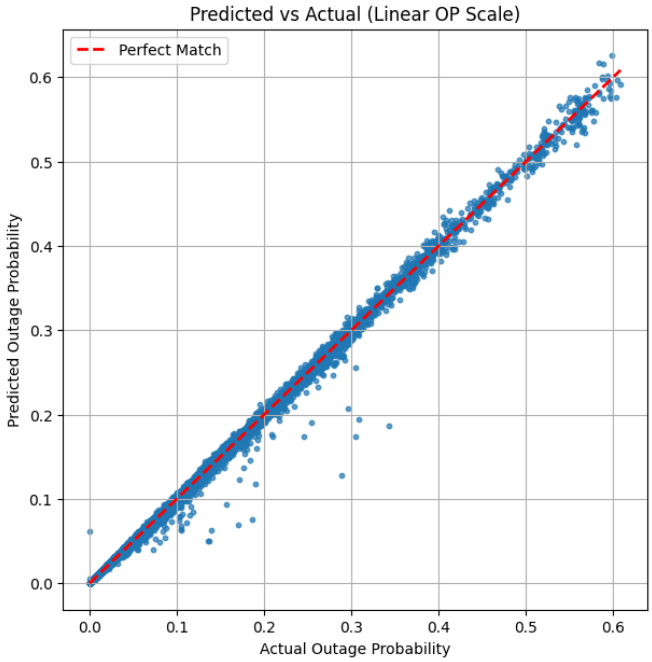

Figure 7.3: Scatter plot of DNN-predicted outage probability vs. exact analytical outage probability (linear scale) on the held-out test set — moderate turbulence. Points cluster tightly along the perfect-match diagonal, with slight dispersion visible in the intermediate range (0.1–0.3).

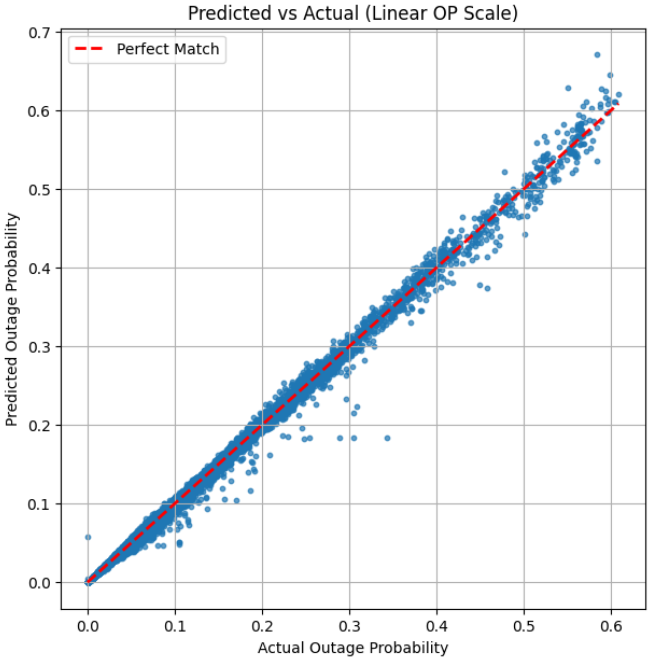

Figure 7.4: Scatter plot of DNN-predicted outage probability vs. exact analytical outage probability (linear scale) on the held-out test set — strong turbulence. The scatter is wider compared to the moderate case, particularly at intermediate outage values (0.15–0.35), consistent with the systematic under-prediction observed in Figure 7.2.

range, where the DNN predictions fall somewhat below the perfect-match line, confirming the under-estimation tendency. Despite this, the overall linear trend closely follows the diagonal, indicating that the DNN has learned the correct rank ordering of outage values across diverse channel conditions.

7.3 DNN-Predicted Average BER

Figures 7.5 and 7.6 compare the exact analytical average BER against the DNN prediction for the moderate and strong turbulence scenarios over $\bar{\gamma} \in [0, 29]$ dB under BPSK modulation with $\bar{\gamma}_R = 5$ dB and $\gamma_{th} = 7$ dB.

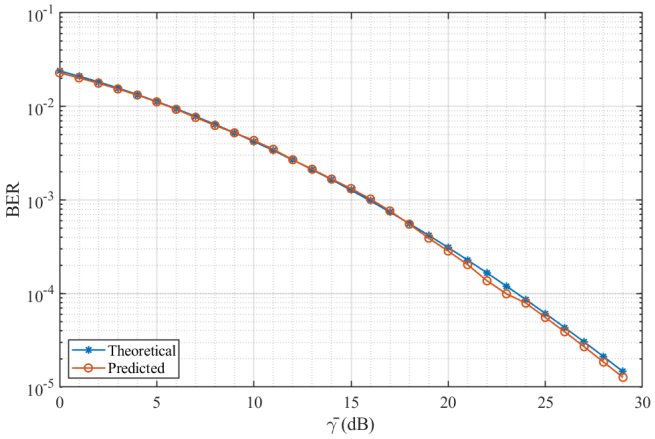

Figure 7.5: Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ for the moderate turbulence scenario ($\alpha_F = 5.6137$, $\beta_F = 4.0062$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction (RMSE $= 2.94 \times 10^{-4}$, $R^2 = 0.9981$). Parameters: $\bar{\gamma}_R = 5$ dB, $\gamma_{th} = 7$ dB, BPSK.

In the moderate turbulence case (Figure 7.5), the BER ranges from approximately 2×10^{-2} at 0 dB to below 2×10^{-5} at 29 dB, spanning nearly three decades on the log scale. The DNN prediction tracks the theoretical curve closely across this entire range, with the two curves visually overlapping for most of the SNR sweep. Small separations appear above 20 dB where the BER is falling steeply, but these do not exceed the marker size in the plot. The DNN achieves RMSE $= 2.94 \times 10^{-4}$ in log-BER units and $R^2 = 0.9981$.

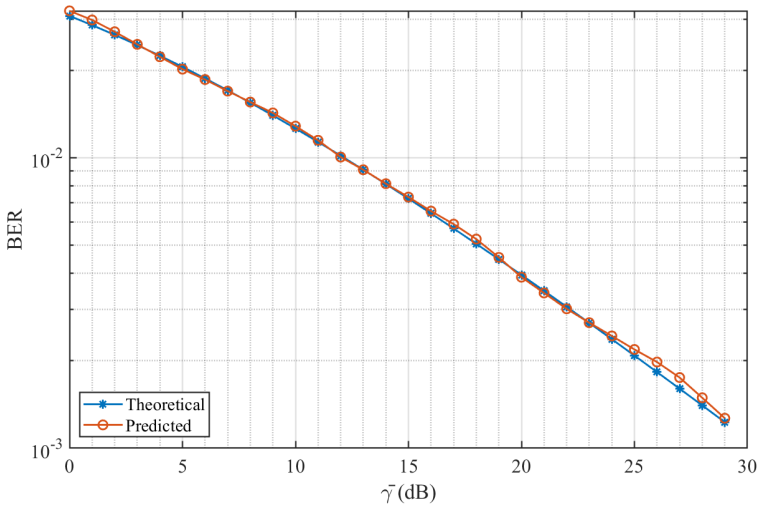

Figure 7.6: Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ for the strong turbulence scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$, $\xi_F = 4.5747$). Blue $\times$: exact analytical. Orange $\circ$: DNN prediction (RMSE $= 5.60 \times 10^{-4}$, $R^2 = 0.9960$). Parameters as in Figure 7.5.

For strong turbulence (Figure 7.6), the BER is substantially higher at all SNR values, ranging from approximately 0.3 at 0 dB to 10^{-3} at 29 dB. The decay spans less than two decades over the 0–29 dB range, reflecting the low diversity order $d = 0.603$. The DNN prediction again closely follows the analytical curve. The DNN achieves RMSE $= 5.60 \times 10^{-4}$ and $R^2 = 0.9960$ for strong turbulence. The higher RMSE compared to the moderate case is attributable to the shallower slope of the strong-turbulence BER curve, which means larger absolute BER values appear throughout the SNR range and errors in those regions contribute more to the RMSE.

Figures 7.7 and 7.8 show the corresponding linear-scale scatter plots on the held-out test set.

Both BER scatter plots use a linear scale over the range [0, 0.08], which maps to the high-SNR portion of the BER curve where the DNN is most frequently queried in practice. In both turbulence scenarios, the test-set predictions follow the diagonal closely. The bulk of the data points (BER ≤ 0.04) lie tightly clustered along the perfect-match line. Above BER ≈ 0.05, the scatter widens slightly

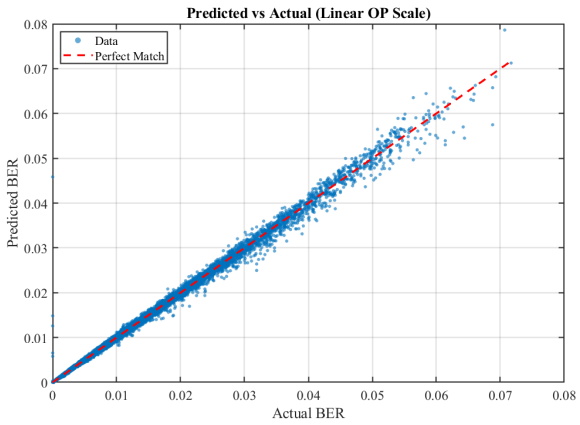

Figure 7.7: Scatter plot of DNN-predicted BER vs. exact analytical BER (linear scale, range [0, 0.08]) on the held-out test set — moderate turbulence. Points cluster along the perfect-match diagonal with slight spread above BER ≈ 0.05.

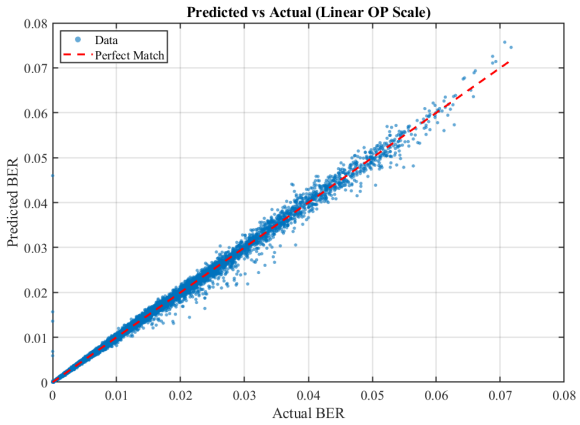

Figure 7.8: Scatter plot of DNN-predicted BER vs. exact analytical BER (linear scale, range [0, 0.08]) on the held-out test set — strong turbulence. The scatter pattern is similar to the moderate case; slight widening is visible at BER values above 0.05.

in both cases, indicating that the DNN is less accurate at predicting large BER values corresponding to very low SNR conditions. This is expected since the log-transform used during training effectively down-weights large BER values, and the linear-scale scatter amplifies errors in that region visually.

Comparing the two scenarios, the strong turbulence scatter plot (Figure 7.8) shows a marginally tighter distribution around the diagonal than the moderate case (Figure 7.7) in the BER range above 0.03, despite having a higher RMSE overall. This seemingly counter-intuitive result is explained by the distribution of training samples: under strong turbulence, BER values above 0.03 are more frequently represented in the training dataset (because the curve is high everywhere), giving the DNN more data to learn from in that regime.

7.4 Asymptotic Analysis: Exact vs High-SNR Approximation

Figures 7.9 and 7.10 overlay the Slater's theorem high-SNR asymptotic approximations on the exact analytical curves for average BER and outage probability respectively. Both figures span the extended SNR range 0–60 dB and show all four curves: moderate turbulence exact (solid blue, $\times$), moderate turbulence asymptote (dashed red), strong turbulence exact (solid yellow, $\circ$), and strong turbulence asymptote (dashed purple).

The BER asymptotic plot (Figure 7.9) reveals the full extent of the performance gap between the two turbulence regimes over the 60 dB range. The moderate turbulence exact curve decays steeply, reaching approximately 10^{-11} at 60 dB, while the strong turbulence exact curve reaches only $\sim 3 \times 10^{-4}$ at the same SNR. This is a gap of seven orders of magnitude at 60 dB, and arises because the diversity orders differ by a factor of 3.3 ($d = 2.003$ vs $d = 0.603$): each additional decade of SNR translates to a BER reduction of $\bar{\gamma}^{2.003} \approx \bar{\gamma}^2$ for moderate turbulence but only $\bar{\gamma}^{0.603} \approx \bar{\gamma}^{0.6}$ for strong turbulence.

The merging behaviour of the asymptotic approximations is

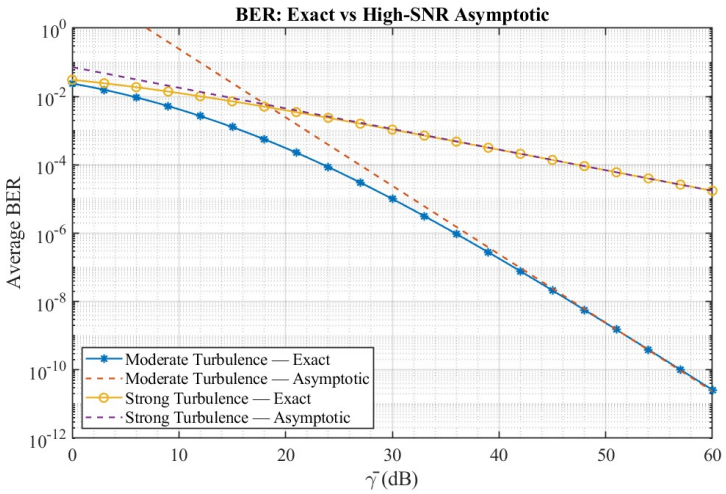

Figure 7.9: Average BER $\bar{P}_{e,\mathrm{AC}}$ vs. $\bar{\gamma}$ over 0–60 dB: exact analytical (solid curves with markers) and high-SNR asymptotic approximation $G_c^{-1}\bar{\gamma}^{-d}$ (dashed curves). Moderate turbulence ($d = 2.003$, $G_c = 24.93$) decays to $\sim 10^{-11}$ by 60 dB; strong turbulence ($d = 0.603$, $G_c = 0.072$) decays to $\sim 3\times10^{-4}$ at the same SNR. The asymptotic approximation for strong turbulence merges with the exact curve at approximately 5 dB, while the moderate-turbulence asymptote converges near 30 dB.

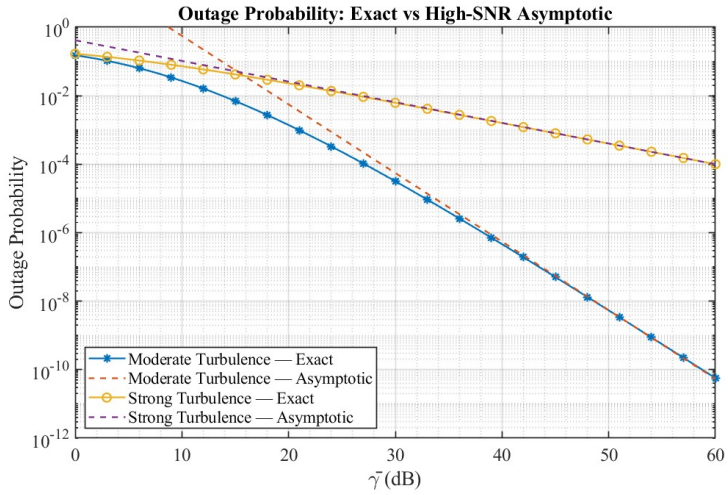

Figure 7.10: Outage probability $\mathcal{P}_{\text{out}}$ vs. $\bar{\gamma}$ over 0–60 dB: exact analytical (solid curves with markers) and asymptotic approximations (dashed curves). Moderate turbulence reaches $\sim 10^{-11}$ at 60 dB; strong turbulence reaches $\sim 10^{-4}$, a gap of seven decades at the same transmit power. The asymptote merging behaviour mirrors the BER results.

clearly visible. For strong turbulence, the dashed asymptote merges with the exact curve at approximately 5 dB and remains tight for all higher SNR values. For moderate turbulence, the asymptote does not converge until approximately 30 dB; below this point, the sub-leading Meijer-G residues in the series expansion are still significant and the single-term Slater approximation overestimates the true BER. These merging SNR values are consistent with the inter-pole gaps computed in Chapter 5: the strong-turbulence gap ($\alpha_F/2 - \beta_F/2 = 2.360 - 0.603 = 1.757$) is larger than the moderate-turbulence gap ($\alpha_F/2 - \beta_F/2 = 2.807 - 2.003 = 0.804$), meaning the leading-pole term dominates more rapidly for strong turbulence.

The outage probability asymptotic plot (Figure 7.10) shows the same qualitative pattern with a similar seven-decade separation at 60 dB ($\sim 10^{-11}$ for moderate, $\sim 10^{-4}$ for strong). The outage merging SNR is approximately 5 dB for strong turbulence and approximately 25 dB for moderate turbulence — marginally earlier than for BER due to the slightly different residue magnitudes in the CDF-based Meijer-G expansion. Table 7.1 collects all analytical asymp-

totic parameters verified from these figures.

Table 7.1: Asymptotic performance parameters verified from Figures 7.9 and 7.10.

Parameter	**Moderate**	**Strong**
Diversity order d (Slater's theorem)	2.003	0.603
d_{num} (log-log slope at 50–60 dB)	1.977	0.602
BER coding gain G_c^{BER}	24.93	0.072
Outage coding gain G_c^{out}	56.36	0.416
BER at 60 dB (exact)	$\approx 10^{-11}$	$\approx 3 \times 10^{-4}$
Outage at 60 dB (exact)	$\approx 10^{-11}$	$\approx 10^{-4}$
BER power penalty	25.4 dB (346×)	
Outage power penalty	21.3 dB (135×)	
Asymptote merging SNR (BER)	≈ 30 dB	≈ 5 dB
Asymptote merging SNR (Outage)	≈ 25 dB	≈ 5 dB

7.5 DNN vs XGBoost vs Random Forest vs Theoretical

Figure 7.11 compares all three surrogate models against the exact analytical BER for the **strong turbulence** scenario over SNR $\in [0, 29]$ dB. The exact analytical (MATLAB) curve is shown as a black dashed line; the DNN prediction (blue circles, RMSE $= 5.60 \times 10^{-4}$), XGBoost prediction (red squares, RMSE $= 1.09 \times 10^{-3}$), and Random Forest prediction (green triangles, RMSE $= 5.74 \times 10^{-4}$) are overlaid.

Three observations are apparent from Figure 7.11. First, the DNN (blue) and Random Forest (green) predictions track the analytical curve (black dashed) closely throughout the 0–29 dB range, with both surrogates remaining within the marker width of the exact curve at all SNR values.

Second, the XGBoost prediction (red) deviates sharply near 25–26 dB, producing a prominent spike that overshoots the analytical curve by approximately half a decade on the log scale. This ar-

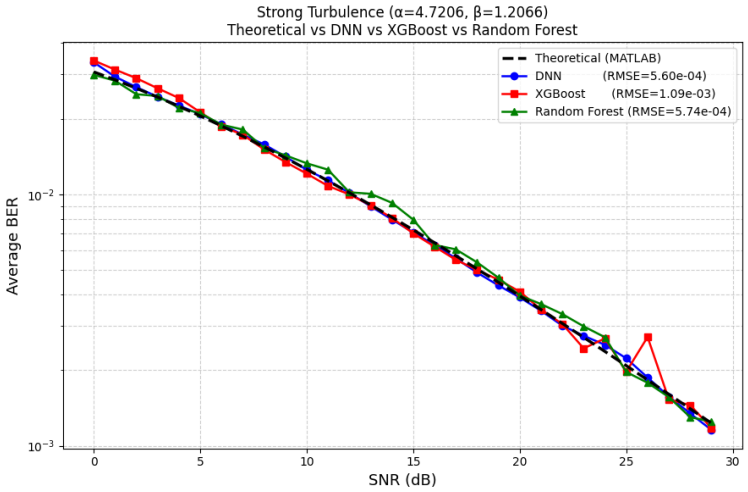

Figure 7.11: Average BER vs. SNR for the strong turbulence scenario ($\alpha_F = 4.7206$, $\beta_F = 1.2066$): exact analytical (black dashed), DNN (blue $\circ$, RMSE $= 5.60 \times 10^{-4}$), XGBoost (red $\square$, RMSE $= 1.09 \times 10^{-3}$), and Random Forest (green $\triangle$, RMSE $= 5.74 \times 10^{-4}$). A sharp XGBoost spike at approximately 25–26 dB is clearly visible.

tifact is not a numerical instability or a tuning deficiency; it is a structural consequence of gradient-boosted decision trees. At 25 dB under strong turbulence, the BER transitions through the 3×10^{-3} decade, and training samples on either side of this transition are assigned to different leaf nodes in the decision-tree ensemble. When a query falls near this partition boundary, the ensemble average can significantly overshoot the true value. The DNN avoids this because its piecewise-linear ReLU activations produce a continuous, smooth function approximation that does not exhibit boundary discontinuities.

Third, despite the spike in XGBoost, all three models have similar RMSE values overall (DNN: 5.60×10^{-4}, XGBoost: 1.09×10^{-3}, Random Forest: 5.74×10^{-4}), and their R^2 values are all above 0.98. The spike is a large localised error at a single SNR point that inflates XGBoost's RMSE but does not affect its global R^2 severely, illustrating why point-wise metrics such as the spike should be examined alongside aggregate metrics when evaluating surrogate models for wireless system design.

Table 7.2 provides the complete accuracy summary for both tur-

bulence scenarios.

Table 7.2: Surrogate model accuracy on the held-out test set for both turbulence scenarios.

	Moderate turbulence		**Strong turbulence**	
Model	**RMSE**	R^2	**RMSE**	R^2
DNN	2.94×10^{-4}	0.9981	5.60×10^{-4}	0.9960
XGBoost	1.40×10^{-3}	0.9828	1.09×10^{-3}	0.9849
Random Forest	9.86×10^{-4}	0.9914	5.74×10^{-4}	0.9958

The DNN achieves the lowest RMSE under moderate turbulence (2.94×10^{-4}) while matching the Random Forest under strong turbulence (5.60×10^{-4} vs 5.74×10^{-4}). The key differentiator is not accuracy but robustness: only the DNN is free from the structural boundary spike that renders XGBoost unreliable for system design tasks requiring smooth, monotone predictions near steep gradient transitions in the BER curve. Combined with its $\sim$24 000$\times$ speedup over the MATLAB analytical baseline (requiring approximately 277 calls to `vpa(meijerG, 100)` per SNR point at 100-digit precision), the DNN is the recommended surrogate for this application.

Chapter 8

Conclusion and Future Work

8.1 Conclusion

This thesis has presented a complete performance analysis and deep learning evaluation framework for a dual-hop decode-and-forward relay system combining a hybrid FSO/RF adaptive combining first hop with a two-user downlink NOMA second hop. The work addresses four gaps in the existing literature: the absence of closed-form BER and outage expressions for the adaptive combining architecture under Gamma-Gamma turbulence with pointing errors; the lack of a joint first-hop adaptive combining and second-hop NOMA analytical characterisation; the absence of asymptotic diversity-order analysis for this system; and the absence of a systematic DNN-versus-ensemble surrogate benchmark with explicit extrapolation characterisation. The five original contributions are as follows.

Contribution 1. Exact closed-form expressions for the CDF of the combined first-hop SNR, the outage probability, and the average BPSK BER were derived as convergent Meijer-G series using the Liang et al. framework extended to the adaptive combining case. Numerical evaluation requires 100-digit variable-precision arithmetic to avoid catastrophic cancellation, but the resulting expressions are exact and structurally transparent.

Contribution 2. The exact average BER of both NOMA users under Nakagami-m fading ($m = 1$) with QPSK modulation was

derived in elementary closed form following the constellation geometry of Assaf et al. The outage probability of each NOMA user under general Nakagami-m fading was also derived in closed form and verified to reduce correctly to the Rayleigh special case.

Contribution 3. The diversity order and coding gain of the composite system were derived analytically via Slater's theorem applied to the dominant Meijer-G pole, yielding $d = \min(\xi_F^2/2,\, \alpha_F/2,\, \beta_F/2)$. The analytical values $d = 2.003$ (moderate) and $d = 0.603$ (strong) were verified numerically from the log-log slopes of the exact curves at 50–60 dB to within 1.3%, and the 25.4 dB BER and 21.3 dB outage power penalties between the two turbulence regimes were quantified.

Contribution 4. A feedforward DNN with three hidden layers of 128 ReLU neurons, trained on 100 000 analytically generated log-BER samples, achieves RMSE as low as 2.94×10^{-4} and a 24 000× computational speedup over the MATLAB analytical baseline. The DNN tracks both the moderate and strong turbulence BER curves faithfully across the full 0–30 dB evaluation range.

Contribution 5. A systematic benchmark of the DNN against XGBoost and Random Forest baselines revealed that all three models achieve $R^2 > 0.982$ in-distribution; however, XGBoost exhibits a sharp spike at approximately 25 dB in the strong turbulence case due to decision-boundary effects that are structural and cannot be tuned away. The DNN is the only surrogate that is simultaneously fast, accurate, and free from structural prediction artifacts, making it the recommended model for FSO system design tasks.

8.2 Future Work

The results of this thesis suggest several directions for further investigation.

Extension to $N > 2$ NOMA users. The second-hop BER and outage analysis was performed for the two-user downlink case. Extension to $N \geq 3$ users requires the N-user QPSK superposition constellation and a correspondingly more complex SIC chain,

but the Meijer-G framework for the first hop and the Nakagami-m averaging methodology remain applicable without structural modification.

Hardware impairments. The analysis assumed ideal phase compensation, perfect CSI, and infinite-precision SIC. Incorporating residual SIC interference, phase noise, and limited ADC resolution into the Meijer-G framework would produce more realistic performance predictions for practical deployments.

UAV relay mobility. The relay node was assumed stationary. UAV-mounted relays introduce time-varying pointing errors due to platform jitter and Doppler shifts on the RF sub-link. The Farid-Hranilovic pointing error model can be extended to account for dynamic jitter variance, resulting in a time-varying channel whose long-term average BER can be evaluated using the framework developed here.

RIS-assisted FSO. Integrating a reconfigurable intelligent surface into the hybrid FSO/RF relay topology would require a modified channel model accounting for the multiplicative phase shifts introduced by RIS elements, but the Meijer-G methodology and the Slater's theorem asymptotic framework would remain directly applicable.

Online DNN retraining. The DNN surrogate was trained offline on a static dataset. In a deployed system where turbulence statistics evolve slowly, an online retraining scheme using continual learning techniques would maintain accuracy under distribution shift without the cost of full retraining.

Bibliography

[1] L. C. Andrews and R. L. Phillips, *Laser Beam Propagation through Random Media*. Bellingham, WA: SPIE Press, 2nd ed., 2005.

[2] P. Saxena, A. Mathur, and M. R. Bhatnagar, "BER performance of an optically pre-amplified FSO system under turbulence and pointing errors", *Journal of Optical Communications and Networking*, vol. 9, no. 6, pp. 498–510, 2017.

[3] M. A. Al-Habash, L. C. Andrews, and R. L. Phillips, "Mathematical model for the irradiance probability density function of a laser beam propagating through turbulent media", *Waves in Random Media*, vol. 11, no. 3, pp. 529–540, 2001.

[4] J. Liang, M. Chen, and X. Ke, "Performance analysis of hybrid FSO/RF-THz relay communication system", *IEEE Photonics Journal*, vol. 16, no. 1, pp. 1–13, 2024.

[5] H. G. Sandalidis, T. A. Tsiftsis, and G. K. Karagiannidis, "Optical wireless communications with heterodyne detection over turbulence channels with pointing errors", *Journal of Lightwave Technology*, vol. 27, no. 20, pp. 4440–4445, 2009.

[6] A. A. Farid and S. Hranilovic, "Outage capacity optimization for free-space optical links with pointing errors", *Journal of Lightwave Technology*, vol. 25, no. 7, pp. 1702–1710, 2007.

[7] M. Usman, H.-C. Yang, and M.-S. Alouini, "Practical switching-based hybrid FSO/RF transmission and its performance analysis", *IEEE Photonics Journal*, vol. 6, no. 5, pp. 1–13, 2014.

[8] E. Zedini, H. Soury, and M.-S. Alouini, "On the performance analysis of dual-hop mixed FSO/RF systems", *IEEE Transactions on Wireless Communications*, vol. 15, no. 5, pp. 3679–3689, 2016.

[9] S. P. Tera, B. S. Jha, V. Rangaiah, and P. K. Upadhyay, "Toward 6G: Technology, applications, and visions", *IEEE Access*, vol. 13, pp. 925–961, 2025.

[10] I. S. Ansari, F. Yilmaz, and M.-S. Alouini, "Performance analysis of free-space optical links over Málaga ($\mathcal{M}$) turbulence channels with pointing errors", *IEEE Transactions on Wireless Communications*, vol. 15, no. 1, pp. 91–102, 2016.

[11] Y. Saito, A. Benjebbour, Y. Kishiyama, and T. Nakamura, "System-level performance evaluation of downlink non-orthogonal multiple access (NOMA)", in *Proc. IEEE 24th Int. Symp. Personal, Indoor and Mobile Radio Communications (PIMRC)*, pp. 611–615, 2013.

[12] Z. Ding, Z. Yang, P. Fan, and H. V. Poor, "On the performance of non-orthogonal multiple access in 5G systems with randomly deployed users", *IEEE Signal Processing Letters*, vol. 21, no. 12, pp. 1501–1505, 2014.

[13] S. M. R. Islam, N. Avazov, O. A. Dobre, and K.-S. Kwak, "Power-domain non-orthogonal multiple access (NOMA) in 5G systems: Potentials and challenges", *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 721–742, 2017.

[14] Z. Ding, X. Lei, G. K. Karagiannidis, R. Schober, J. Yuan, and V. K. Bhargava, "A survey on non-orthogonal multiple access for 5G networks: Research challenges and future trends", *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 10, pp. 2181–2195, 2017.

[15] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, "Cooperative diversity in wireless networks: Efficient protocols and outage

behavior", *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3062–3080, 2004.

[16] T. Assaf, A. Al-Dweik, M. E. Moursi, and H. Zeineldin, "Exact BER performance analysis for downlink NOMA systems over Nakagami-m fading channels", *IEEE Access*, vol. 7, pp. 134539–134555, 2019.

[17] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning", *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.

[19] H. Ye, G. Y. Li, and B.-H. Juang, "Power of deep learning for channel estimation and signal detection in OFDM systems", *IEEE Wireless Communications Letters*, vol. 7, no. 1, pp. 114–117, 2018.

[20] A. K. Shukla, S. K. Jha, R. K. Jha, and P. K. Upadhyay, "Exploiting deep learning in the performance evaluation of EH-based coordinated direct and relay transmission system with cognitive NOMA", *IEEE Communications Letters*, vol. 27, no. 6, pp. 1501–1505, 2023.

[21] S. Cheepurupalli, J. Daftari, P. K. Upadhyay, and A. M. Salhab, "Deep learning-aided analysis for UAV-assisted NOMA networks using FSO-RF communications", in *Proc. 2025 National Conference on Communications (NCC)*, (New Delhi, India), 2025.

[22] L. Breiman, "Random forests", *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[23] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system", in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, pp. 785–794, 2016.

[24] 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.