



Zaidi, Syed Basit Ali (2026) *From correlation-driven to structure-aware learning: towards innately-intelligent neural networks for wireless systems and beyond*. PhD thesis.

<https://theses.gla.ac.uk/86191/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

From Correlation-Driven to Structure-Aware Learning: Towards Innately-Intelligent Neural Networks for Wireless Systems and Beyond

Syed Basit Ali Zaidi

SUBMITTED IN FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF
DOCTOR OF PHILOSOPHY

JAMES WATT SCHOOL OF ENGINEERING
COLLEGE OF SCIENCE & ENGINEERING



2026

I am grateful to Allah for providing me with this opportunity. I also owe a great deal to my supervisors, Dr. Shuja Ansari, Prof. Muhammad Ali Imran, Prof. Ali Imran, and Prof. Qammer Abbasi. Their guidance went far beyond academic supervision, they challenged my thinking, pushed me when I needed it, and refused to let me settle for average. This work carries their influence in ways that cannot be fully captured in words. If I have developed any ability to think critically, question deeply, and build with purpose, a large part of that credit goes to them. To my family, especially my Ammi and Baba, and my sisters, this journey would not have been possible without you. You stood behind me through uncertainty, pressure, and long periods where results were not guaranteed. Your prayers, patience, and constant belief gave me the stability to keep going when things were not working. This achievement is as much yours as it is mine.

Abstract

Modern wireless networks increasingly rely on artificial intelligence to support spectrum awareness, radio propagation modelling, coverage prediction, and automated network optimisation. Although conventional deep neural networks (DNNs) can achieve strong predictive performance, they are typically correlation-driven, data-hungry, and difficult to interpret. These limitations become critical in wireless systems, where labelled measurements are expensive to obtain, deployment conditions vary across cities, carrier frequencies, antenna configurations, clutter classes, and planning tools, and operational decisions often require traceable physical reasoning rather than opaque predictions.

This thesis investigates how wireless learning models can move from purely data-driven prediction toward structure-aware and interpretable intelligence. First, it examines deep learning for wireless spectrum awareness in vehicular communication scenarios, demonstrating the practical value of NN models while also exposing their dependence on learned statistical representations. Building on this limitation, the thesis proposes Innately Intelligent Neural Networks (IINNs), a domain-informed NN modelling framework in which physical variables, analytical relationships, and mathematical operators are embedded directly into the network architecture before training. Unlike post-hoc explainability methods, IINN is designed to be interpretable by construction: intermediate computations correspond to meaningful wireless propagation components.

The proposed framework is evaluated for radio propagation and received signal reference power prediction under limited-data and out-of-distribution conditions. Across deployment shifts, IINN demonstrates stronger robustness than representative black-box baselines, including $(22.13 \pm 2.91)\%$ relative OOD degradation under joint city–frequency shift compared with 45.32–58.04% for the baseline models. Across training-data fractions from 100% to 5%, IINN achieves the lowest average relative OOD degradation, $(17.70 \pm 3.38)\%$. The evaluation also includes propagation label-shift testing, for which absolute OOD RMSE is reported because the target labels are generated by a different propagation model. Structural traceability analysis further shows that IINN exposes learned analytical coefficients and section-wise propagation calculations during inference.

Finally, the thesis extends IINN to transfer learning, showing that physically meaningful model components provide a principled basis for deciding which parameters should be retained, adapted, or recalibrated across deployment scenarios. Overall, this thesis positions IINN as a compact, auditable, and data-efficient learning paradigm for wireless systems, particularly suited to settings where partial domain knowledge is available, labelled data are limited, and trustworthy model behaviour is required.

Contents

Abstract	iii
Acknowledgements	xii
Declaration	xiii
1 Background and Literature Review	1
1.1 Motivation	2
1.1.1 Wireless AI Beyond In-Distribution Accuracy	2
1.1.2 The Cost of Correlation-Driven Learning	3
1.1.3 Why Architecture-Level Intervention Matters	3
1.2 Literature Review and Research Gap	4
1.2.1 Spectrum Awareness and Modulation Classification	4
1.2.2 Learning-Based Wireless Propagation Modelling	5
1.2.3 Explainability, Physics-Guided Learning, and Structural Traceability . .	5
1.2.4 Limited Data, OOD Robustness, and Transfer Learning	6
1.2.5 Central Knowledge Gap	6
1.3 Research Problem	7
1.4 Research Aim and Questions	7
1.5 Thesis Overview	8
1.6 Contributions	9
1.7 Scope of the thesis	10
1.8 Structure of the thesis	11
1.9 Publications	12
1.10 Chapter Summary	13
2 An Efficient Deep Learning-based Spectrum Awareness Approach for Vehicular Communication	14

2.1	Preliminaries and Overview	14
2.1.1	Spectrum monitoring techniques	15
2.1.2	Modulation classification techniques	15
2.2	Role of this Chapter in the Thesis	16
2.3	Introduction	17
2.3.1	Evaluation Metrics and Modulation-Switching Rule	19
2.3.2	Review of the 3D stochastic V2V channel	20
2.3.3	The proposed image-based deep learning model	22
2.4	Spectrum awareness approach	24
2.4.1	The training phase	25
2.4.2	The classification phase	26
2.5	Performance evaluation	27
2.5.1	Experiment 1: P_d at different channel variations $var.$	27
2.5.2	Experiment 2: P_d at different vehicle speeds $u_{A(B)\max}$	28
2.5.3	Experiment 3: P_d at different SNRs	28
2.5.4	Experiment 4: P_d in urban and rural areas	29
2.5.5	Data-set publication	31
2.6	Chapter Summary	33
3	IINN Construction Methodology and Modelling Framework	35
3.1	Motivation and Problem Definition	35
3.2	Related Work	36
3.2.1	Post-Hoc Explainability in Wireless and Data-Driven Models	36
3.2.2	Intrinsically Interpretable and Knowledge-Distilled Models	37
3.2.3	Data Scarcity, Representation Bias, and Synthetic Data Generation	38
3.2.4	Physics-Informed and Domain-Guided Neural Learning	39
3.2.5	Toward Inherently Interpretable Domain-Embedded Architectures	40
3.2.6	Positioning of IINN Relative to Existing Paradigms	41
3.3	Innately Intelligent Neural Networks: Modeling Framework	43
3.3.1	General Framework for Designing IINN	43
3.3.2	Wireless Propagation Instantiation	47
3.3.3	Illustrative Sub-Network: Vertical Antenna Alignment	49

3.3.4	Overall IINN Architecture and Interpretability	50
3.3.5	Structural Traceability and Component Audit	53
3.3.6	Relation to Feature-Engineering and Conventional Neural Models	53
3.4	Chapter Summary	54
4	IINN Propagation Experiments and Structural Traceability	55
4.1	Chapter Organisation	55
4.2	Network Simulation Setting for Data Generation and Collection	56
4.2.1	Motivation for the Wireless Propagation Case Study	56
4.2.2	Simulation Environment and Deployment Scenario	56
4.2.3	Dataset Composition and Feature Description	57
4.2.4	Urban Morphology and Clutter Distribution	59
4.2.5	Terrain, Obstruction, and Propagation Profiles	60
4.2.6	RSRP-Distance Characteristics and Physical Consistency	61
4.3	Performance Evaluation	63
4.3.1	Experimental Protocol and Benchmark Models	63
4.3.2	Input Features and Evaluation Metrics	65
4.3.3	OOD Robustness Across Deployment Shifts	66
4.3.4	Overall Comparative Analysis	70
4.3.5	Spatial Consistency Analysis	71
4.4	Structural Traceability and Training Stability	72
4.4.1	Training Stabilization and OOD Optimization	72
4.4.2	Forward Structural Traceability	73
4.5	Discussion	77
4.5.1	Key Findings and Insights	77
4.5.2	Comparison with Existing Approaches	78
4.5.3	Scope and Practical Considerations	78
4.6	Chapter Summary	79
5	IINN Transfer Learning Framework	80
5.1	Introduction	80
5.1.1	Related Work	81
5.1.2	Contributions	82

5.2	IINN-Based Transfer Learning Framework	83
5.2.1	Architecture-to-Parameter Mapping	83
5.2.2	Domain-Aware Fine-Tuning Strategy	85
5.3	Simulation Setup and Performance Evaluation	86
5.4	Transfer-Learning Results	87
5.4.1	Base-Station Height Transfer	87
5.4.2	Sector-Based Transfer	87
5.4.3	Combined Height and Sector Configurations	88
5.4.4	When the Pre-Trained IINN Should and Should Not Transfer	89
5.5	Conclusion	91
5.6	Chapter Summary	91
6	Synthesis, Scope, Limitations, and Future Research Directions	92
6.1	The Research Arc	93
6.2	Levels of Intervention in Wireless AI	93
6.3	What the Evidence Establishes	95
6.3.1	Deep Learning Is Useful But Structurally Opaque	95
6.3.2	Structure-Aware Architecture Improves Interpretability and Robustness	95
6.3.3	Interpretability Enables Principled Transfer	96
6.3.4	The Thesis Supports Conditional Superiority, Not Universal Superiority	96
6.4	Distinction Matrix	96
6.5	Accuracy, Interpretability, and the Role of LightGBM	97
6.6	Propagation-Label Shift and Structural Traceability	98
6.7	Implications for AI-Native Wireless Networks	98
6.8	Synthesis Summary	99
6.9	Overall Conclusion and Future Work	99
6.10	Answer to the Central Research Question	99
6.11	Consolidated Thesis Contributions	100
6.12	Main Technical Findings	101
6.12.1	Architecture Is a Form of Prior Knowledge	101
6.12.2	Interpretability and Robustness Can Have a Shared Cause	101
6.12.3	Structural Traceability Requires Inspectable Computation	101

6.12.4	Interpretability Can Improve Learning Efficiency	101
6.13	Limitations	102
6.14	Future Research Directions	102
6.14.1	Real-Measurement Validation	102
6.14.2	Systematic Model-Misspecification Analysis	103
6.14.3	Cross-Domain IINN Validation	103
6.14.4	Automated Equation-to-Architecture Generation	103
6.14.5	Hybrid Traceable and Black-Box Systems	103
6.14.6	IINN for AI-Native RAN Digital Twins	103
6.14.7	IINN-LLM Explanation Interfaces	104
6.14.8	Transfer Learning Across Deployment Shifts	104
6.15	Chapter Summary	104
A	Supplementary Material	105
A.1	Chapter 2 Supplementary Dissemination Material	105

List of Tables

1.1	Role of each chapter in the thesis argument.	9
2.1	Channel parameter settings	21
2.2	P_d at different channel variations	30
2.3	P_d at different vehicle speeds	30
2.4	P_d at different signal-to-noise ratios	31
2.5	P_d in urban and rural areas	31
2.6	The confusion matrix at Speed = 30m/s, $var.$ = 0.3, and SNR = 25 dB in urban area	33
3.1	Conceptual and architectural comparison of IINN with related learning paradigms.	42
3.2	Interpretability, generalization, and deployment comparison of IINN with related learning paradigms.	43
3.3	Layer types used in the Keras-based IINN architecture.	51
3.4	Qualitative comparison between feature-engineering-plus-MLP approaches and the proposed IINN framework.	54
4.1	Simulation settings	57
4.2	Model configurations used for Chapter 3 robustness validation.	64
4.3	Key parameters used for analytical equation development, dataset creation, and IINN training.	65
4.4	Summary of the evaluated OOD shift settings and their corresponding IINN and best-baseline results. Values are reported as mean $\pm$ standard deviation over five random-seed runs. Experiments I–III report relative OOD degradation; Experiment IV reports absolute OOD RMSE because the label-generation mechanism changes.	72
4.5	3GPP equation coefficients and corresponding IINN learned coefficients after train- ing with all available source-domain samples and with only 5% of the source- domain samples.	74

4.6	Section-wise coefficient trace for one Experiment I prediction. The table reports the 3GPP coefficients, the corresponding IINN learned coefficients, and the forward calculation for the selected test sample. Displayed values are rounded; section outputs are computed from the saved model values.	76
5.1	Domain-aware layer selection for IINN transfer learning.	85
5.2	Network scenario setting for the transfer-learning study.	86
5.3	Summary of transfer-learning experimental setup.	87
6.1	Different intervention levels for introducing reliability, interpretability, or domain knowledge into wireless learning systems.	94
6.2	Defensibility of the main thesis claims.	97

List of Figures

2.1	Techniques for modulation classification	16
2.2	The proposed spectrum awareness approach.	18
2.3	Vehicular channel model representation for spectrum awareness.	22
2.4	The proposed lightweight CNN architecture for constellation-image modulation classification. The convolutional blocks learn local I/Q cluster geometry, while the softmax output produces six class probabilities.	24
2.5	Implementation block diagram.	26
2.6	Black-and-white constellation-image training samples for the six modulation classes. The grayscale representation preserves I/Q geometry while avoiding colour features that have no physical meaning in this task.	27
2.7	Dataset deposit; DOI: 10.5525/gla.researchdata.1420	32
3.1	Workflow for constructing an IINN model. The example illustrates how a portion of a radio propagation equation can guide interpretable and domain-aligned architecture design.	46
3.2	Received signal power model P_r for a UE in a 3D urban macro environment, based on 3GPP TR 38.900. The expression captures distance-dependent path loss, LoS/N-LoS selection, building and street characteristics, antenna heights, and 3D antenna-gain patterns. This analytical form provides the scaffold for constructing the IINN architecture.	49
3.3	Illustration of the IINN sub-network constructed from the vertical directional antenna alignment component of the radio propagation equation. The network takes four input variables, ϕ_{ue} , ϕ_{bs} , B_v , and A_v , and processes them through mathematically aligned layers. Each operation, including subtraction, division, squaring, and minimum selection, mirrors the corresponding analytical computation.	51

3.4	Schematic diagram of the IINN model constructed from the 3GPP TR 38.900 received-power equation. Variables from the analytical equation are translated into scalar inputs and processed through dedicated PLO layers, followed by mathematical operators such as subtraction, logarithm, division, minimum, and maximum selection.	52
4.1	Representative 1 m resolution 3D map samples over a 1 km ² urban tile used for cross-city generalization experiments. The maps show simulated gNodeB deployment and UE distribution for Brussels at 3.7 GHz and Chicago at 3.3 GHz.	58
4.2	User distribution across clutter classes for the Brussels dataset.	59
4.3	User distribution across clutter classes for the Chicago dataset.	59
4.4	Clutter class distribution comparison between Brussels and Chicago datasets. Significant differences are observed in the prevalence of dense urban and high-rise environments, with Brussels exhibiting a higher concentration of complex propagation conditions.	60
4.5	Terrain and clutter profile between Site 0 and Site 3 for Brussels. The plot illustrates the elevation profile and intervening clutter, including buildings and vegetation, along the direct path, providing insight into propagation obstructions between two locations.	61
4.6	RSRP variation with distance for multiple sites and sectors in the Brussels scenario. The solid lines represent sector-wise median RSRP values, while shaded regions indicate the interquartile range (IQR), capturing variability due to local propagation conditions and clutter effects.	61
4.7	Terrain profile between a UE and its associated transmitter in Brussels for Site 1 Sector 2 over a distance of 245 meters.	62
4.8	Terrain profile between a UE and its associated transmitter in Brussels for Site 3 Sector 2 over a distance of 557 meters.	62
4.9	Experiment I: relative OOD degradation under cross-city deployment shift. Results are reported over five random-seed runs.	67
4.10	Experiment II: relative OOD degradation under joint city–frequency deployment shift. Results are reported over five random-seed runs.	68
4.11	Experiment III: relative OOD degradation across source-domain training fractions under cross-city deployment shift.	69
4.12	Experiment III: average relative OOD degradation across data-scarcity settings. . .	69

4.13	Experiment IV: absolute OOD RMSE under propagation-label shift from 3GPP labels to SPM labels.	70
4.14	Comparative performance of IINN across cross-city and cross-frequency evaluation experiments.	70
4.15	Spatial comparison between ground truth and IINN predictions under (a)–(b) city-level OOD, (c)–(d) frequency-level OOD, and (e)–(f) in-domain data scarcity conditions.	71
4.16	Training and validation RMSE curves of IINN. The curves are shown on the standardized RSRP target scale after removing the first epoch for visual clarity. Both curves converge smoothly, with the best validation RMSE of approximately 0.697 at epoch 111, indicating stable data-driven learning of the trainable propagation-aligned components.	73
5.1	IINN-based transfer-learning framework. The source IINN is constructed from domain-informed propagation components, pre-trained on source-domain data, and selectively fine-tuned on target-domain data by updating only the branches associated with changed deployment variables.	82
5.2	Identification of variables and mathematical operators from a propagation equation. Variables define IINN input features, while operators define the structure of the corresponding network components.	84
5.3	Segment of an IINN architecture constructed by mapping analytical propagation-model terms into neural components. The same component alignment is later used to decide which branches should be fine-tuned.	84
5.4	Transfer-gain comparison for source and target base-station heights. Off-diagonal entries show adaptation when the target base-station height differs from the source height.	88
5.5	Sector-based transfer-gain comparison averaged across base-station heights. Positive PA–BL values indicate that IINN-based selective fine-tuning outperforms conventional DNN fine-tuning.	89
5.6	Transfer-gain difference (PA–BL) for combined source-target base-station height and sector configurations.	90

A.1	Poster presentation at SANS related to the deep learning-based spectrum-awareness	
	work.	106

Acknowledgements

The authors acknowledge the use of computational resources and facilities provided by the James Watt School of Engineering, University of Glasgow. The authors used AI-based language tools to assist with English editing and grammar refinement of selected parts of the manuscript. The authors reviewed and revised all AI-assisted text and take full responsibility for the content, methodology, analysis, results, and conclusions.

Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution

Syed Basit Ali Zaidi

Abbreviations

3GPP	3rd Generation Partnership Project
5G	fifth-generation mobile communication system
AI	artificial intelligence
AoA	Angle of Arrival
AoD	Angle of Departure
BS	base station
CNN	convolutional neural network
CQI	channel quality indicator
CSI	channel state information
DL	deep learning
DNN	deep neural network
DNNs	deep neural networks
DRL	deep reinforcement learning
DSRC	dedicated short-range communication
EQL	Equation Learner
GAN	generative adversarial network
gNodeB	next-generation Node B
I/Q	in-phase and quadrature
IINN	Innately Intelligent Neural Network
IINN-TL	Innately Intelligent Neural Network transfer learning
IINNs	Innately Intelligent Neural Networks
KAN	Kolmogorov-Arnold network
KNN	k-nearest neighbours
KPI	key performance indicator

LIME	Local Interpretable Model-agnostic Explanations
LLM	large language model
LoS	line of sight
LSTM	long short-term memory
MDT	minimization of drive tests
ML	machine learning
MLP	multi-layer perceptron
NLoS	non-line-of-sight
NN	neural network
NNs	neural networks
NR	New Radio
OFDM	orthogonal frequency-division multiplexing
PINN	physics-informed neural network
PL	path loss
PLO	parametric linear operation
PLOs	parametric linear operations
PRIME	purposeful, resilient, interpretable, mathematical, and explainable
PSK	phase-shift keying
QAM	quadrature amplitude modulation
RAN	radio access network
RAT	radio access technology
RF	radio frequency
RMSE	Root Mean Squared Error
RNN	recurrent neural network
RSRP	reference signal received power
Rx	receiver
SHAP	SHapley Additive exPlanations
SNR	signal-to-noise ratio
SPM	standard propagation model
SVM	support vector machine
TL	transfer learning

Tx	transmitter
UE	user equipment
UEs	user equipments
UMa	urban macro
UWB	ultra-wideband
V2I	vehicle-to-infrastructure
V2V	vehicle-to-vehicle
V2X	vehicle-to-everything
VANET	vehicular ad hoc network
XAI	explainable artificial intelligence

List of Symbols

A_h	Horizontal antenna attenuation
$\mathbf{x}$	Input feature vector or set of physical variables
A_v	Vertical antenna attenuation
x_i	Individual scalar input variable
B_h	Horizontal antenna beamwidth
y	Target variable or model output
B_v	Vertical antenna beamwidth
$\hat{y}$	Predicted model output
d	Transmitter–receiver separation distance
f_{dom}	Domain-informed analytical function
h	Average building height
ψ	Adjustable coefficients or trainable domain parameters
h_{bs}	gNodeB antenna height
κ	Fixed constants obtained from domain knowledge or standards
h_{ue}	User equipment height
ε	Unmodelled error, measurement noise, or model mismatch
W	Street width
$\mathcal{O}_i$	Analytical or mathematical operation in the IINN decomposition
f_c	Carrier frequency
L_i	Neural layer or trainable affine operation associated with component i

P_t	Transmit power
w_i	Trainable weight of layer or operation i
b_i	Trainable bias of layer or operation i
ζ	Antenna efficiency factor
ϕ_{bs}	gNodeB antenna tilt angle
Ξ_i	Analytical prior value associated with operation i
ϕ_{ue}	UE antenna tilt angle
$\mathcal{L}$	Loss function used during training
θ_{bs}	gNodeB antenna azimuth angle
θ_{ue}	UE antenna azimuth angle
P_e	Probability of classification or detection error
$H_B(t)$	Channel observation at receiver Bob at time t
t	Time
l	Multipath component index
L	Number of multipath components
a_l	Fading coefficient of the l -th multipath component
ϕ_l	Phase delay of the l -th multipath component
v_l	Composite Doppler parameter of the l -th multipath component
$v_{A,l}$	Doppler contribution associated with vehicle Alice
$v_{B,l}$	Doppler contribution associated with vehicle Bob
$v_{S,l}$	Doppler contribution associated with a moving scatterer
$u_{A(B)\max}$	Maximum speed of vehicle Alice or Bob
u_S	Scatterer speed
c	Speed of light
α	Azimuth angle in the vehicular channel model
β	Elevation angle in the vehicular channel model
$p_{u_S}(u_S)$	Probability density function of scatterer speed
w	Scale parameter of the Weibull scatterer-speed distribution

a	Shape parameter of the Weibull scatterer-speed distribution
P_r	Received signal power
P_{rx}	Received power at the receiver
$\hat{P}_{rx}$	Predicted received power
P_t	Transmit power
P_{tx}	Transmitter power used in the revised IINN
$\widehat{RSRP}$	Predicted reference signal received power
PL	Propagation path loss
PL_{LoS}	Line-of-sight path loss
PL_{NLoS}	Non-line-of-sight path loss
G	Antenna gain
G_{max}	Maximum antenna gain
D	Antenna directivity
ζ	Antenna efficiency factor
L_{feeder}	Feeder or cable loss
L_{other}	Additional implementation or system-level loss
d	Transmitter–receiver separation distance
d_{2D}	Horizontal transmitter–receiver separation distance
f_c	Carrier frequency
f	Operating frequency in the propagation model
W	Street width
h	Average building height
h_{bs}	Base-station or gNodeB antenna height
h_{ue}	User-equipment height
Δx	Horizontal coordinate difference between UE and base station
Δy	Vertical or northing coordinate difference between UE and base station
A_h	Maximum horizontal antenna attenuation
A_v	Maximum vertical antenna attenuation
A_{att}	Directional antenna attenuation

B_h	Horizontal antenna beamwidth
B_v	Vertical antenna beamwidth
B	Generic antenna beamwidth used for normalized alignment
ϕ_{bs}	Base-station or gNodeB antenna tilt angle
ϕ_{ue}	User-equipment antenna tilt angle
θ_{bs}	Base-station or gNodeB antenna azimuth angle
θ_{ue}	User-equipment antenna azimuth angle
λ_h	Horizontal antenna-pattern weighting factor
λ_v	Vertical antenna-pattern weighting factor
$\Delta\theta$	Angular misalignment between transmitter and receiver
C	Learned intercept term in the structurally traceable IINN
$T(P_{tx})$	Transmit-power contribution in the revised IINN
$L(d, f)$	Distance-frequency path-loss contribution in the revised IINN
G_{ant}	Antenna-alignment contribution in the revised IINN
H_{site}	Site and environmental contribution in the revised IINN
$C_{clutter}$	Clutter-class offset in the structurally traceable IINN
R_{res}	Residual interaction term in the revised IINN
$R_{bounded}$	Bounded residual correction in the structurally traceable IINN
y_T	Target-domain output in transfer learning
$F(\mathbf{x}; \varphi_i, \varphi_v)$	IINN mapping parameterized by invariant and variable components
φ_i	Invariant parameters retained across source and target domains
φ_v	Domain-variable parameters selected for adaptation

ϕ_v^T	Fine-tuned domain-variable parameters for the target domain
$\mathcal{L}_T$	Target-domain loss used during transfer learning

Chapter 1

Background and Literature Review

Wireless communication systems are becoming increasingly difficult to operate, optimise, and audit. Modern radio access networks are no longer static collections of base stations and terminals. They are dense, heterogeneous, software-controlled systems whose behaviour depends on carrier frequency, antenna configuration, transmit power, propagation environment, clutter class, traffic demand, user mobility, and deployment geometry. As networks move toward AI-native operation, these factors must be predicted and acted upon with limited measurement data and under conditions that often differ from the environment in which a model was trained [1], [2], [3].

Machine learning (ML) and deep neural networks (DNNs) have therefore become attractive tools for wireless inference tasks. They can learn complex nonlinear relationships from data and have demonstrated strong empirical performance in applications such as modulation classification, radio-map prediction, channel estimation, beam management, and coverage modelling. However, this success also exposes a central difficulty: most learning systems used in wireless networks are correlation-driven. They learn statistical associations from training examples, but their internal computational structure is not derived from the physical mechanisms that govern the wireless system. As a result, a model may perform well on a held-out test set while remaining opaque, data hungry, and fragile when exposed to a new city, frequency band, site configuration, propagation model, or measurement condition [4], [5], [6], [7].

This thesis investigates a different design philosophy. Rather than treating neural networks as tabula-rasa function approximators, it asks whether domain knowledge can be embedded into the architecture before training begins. The resulting research programme moves from correlation-driven learning toward structure-aware learning. The central argument is that, for physics-governed wireless tasks, the most consequential design decision is not only which optimiser or training dataset is used, but whether the model begins with any structural understanding of the domain [8], [9], [10].

This chapter therefore serves two roles. First, it introduces the motivation, research problem, research questions, and thesis structure. Second, it consolidates the background and literature review material required to understand the later technical chapters. This responds directly to the revised thesis structure: Chapter 2 is treated as a diagnostic example of conventional data-driven wireless AI, Chapter 3 introduces the IINN construction methodology, Chapter 4 provides the IINN propagation experiments and ablation evidence, and Chapter 5 uses IINN interpretability to guide transfer learning [11], [12], [13].

1.1 Motivation

1.1.1 Wireless AI Beyond In-Distribution Accuracy

In operational wireless networks, prediction accuracy alone is not enough. A radio-planning or automation model must also be trustworthy, auditable, and robust. If a model predicts poor received signal power, an engineer must be able to determine whether the cause is distance-dependent path loss, antenna misalignment, clutter obstruction, insufficient transmit power, or a site-specific propagation effect. A black-box prediction without physical attribution has limited operational value, even if the numerical error is small [7], [14], [15].

This issue is particularly important for network automation. Tasks such as antenna tilt optimisation, coverage anomaly detection, frequency refarming, cell reconfiguration, and digital-twin-assisted planning require models whose outputs can be related to actionable physical causes. In such settings, the question is not only what the model predicted, but why the prediction changed when network or environmental conditions changed. This motivates interpretable learning systems whose internal computations can be inspected without relying entirely on post-hoc approximation tools [2], [6], [16].

1.1.2 The Cost of Correlation-Driven Learning

Conventional DNNs are usually designed through heuristic choices: the number of layers, number of neurons, activation functions, and optimisation settings are selected empirically. The network receives the input variables as a feature vector and learns a mapping to the target through gradient-based optimisation. This design gives DNNs high flexibility, but it also creates three limitations [1], [5], [17].

First, the model is data dependent. Since the architecture contains no explicit domain structure, the model must infer all relevant relationships from labelled examples. In wireless networks, such labels are costly because high-quality measurements require field campaigns, simulation pipelines, or radio-planning tools [5], [6], [18].

Second, the model is difficult to interpret. Internal activations do not naturally correspond to physical quantities such as distance loss, antenna gain, height correction, or clutter attenuation. Post-hoc methods such as SHAP and LIME can estimate feature influence [19], [20], but they do not make the original model structurally interpretable [7], [15], [17].

Third, the model can be sensitive to distribution shift. A network trained on one city, frequency, site geometry, clutter distribution, or propagation-label source may learn correlations that are locally valid but not transferable. This is a serious limitation for wireless systems, where deployment environments vary continuously and training data are rarely representative of all operating conditions [21], [22], [23].

1.1.3 Why Architecture-Level Intervention Matters

Several existing paradigms attempt to address these limitations. Post-hoc explainability estimates feature importance after training. Confidence calibration attaches reliability measures to predictions [23], [24]. Physics-informed neural networks (PINNs) introduce physical laws as loss-function constraints [10], [25]. Deep unfolding maps iterative optimisation algorithms into neural architectures [26], [27]. Transfer learning adapts pre-trained models to new domains using limited target data [28], [29], [30].

These approaches are useful, but they usually do not change the central fact that the internal representation of the model remains weakly coupled to domain physics. The present thesis argues that, when a domain has analytical or empirical equations that describe the target phenomenon, the architecture itself should reflect those equations. This does not mean hardcoding a complete simulator into a neural network. Rather, it means decomposing known physical relationships into trainable computational components so that learning begins from a meaningful structure instead of an unconstrained hypothesis space [8], [9], [12].

1.2 Literature Review and Research Gap

This section consolidates the literature themes that motivate the four technical contribution chapters. The aim is not to repeat the detailed methodology of later chapters, but to make explicit the knowledge gap that links spectrum awareness, wireless propagation modelling, interpretability, robustness, and transfer learning [4], [5], [7], [28].

1.2.1 Spectrum Awareness and Modulation Classification

Spectrum awareness in vehicular communication relies on the ability to monitor available spectrum and identify the modulation format used under time-varying channel conditions. Existing modulation-classification approaches are commonly grouped into likelihood-based and feature-based methods [31], [32], [33], [34]. Likelihood-based methods can be statistically powerful, but they are often computationally expensive and may require prior knowledge of channel parameters. Feature-based methods reduce complexity by extracting signal descriptors and training a classifier, including support-vector-machine and CNN-based approaches. However, much of the prior work evaluates classification mainly as a function of SNR, with less emphasis on simultaneous variation in mobility, channel variance, and rural/urban multipath conditions [4], [35], [36].

Chapter 2 addresses this narrower gap by evaluating a CNN-based modulation-classification pipeline under a 3D stochastic vehicular channel with varying SNR, vehicle speed, channel variation, and propagation environment. Its novelty is not that CNNs are introduced for the first time, but that classification performance is evaluated as a forward channel-quality indicator for

modulation-order selection under realistic channel variability. At the same time, the chapter deliberately exposes a limitation: the CNN can classify constellation images effectively, but its internal representation is not structurally linked to channel mechanisms. This limitation is the bridge from Chapter 2 to IINN [4], [7], [11].

1.2.2 Learning-Based Wireless Propagation Modelling

Wireless propagation modelling has also seen extensive use of ML and DL predictors. Empirical models and deterministic ray-tracing tools remain important, but they can be limited by environment specificity, high simulation cost, or incomplete measurement coverage. ML-based received-power and path-loss predictors have therefore been studied using tree ensembles, boosting methods, DNNs, image-based models, satellite or map features, MDT data, and hybrid feature-engineering pipelines [5], [6], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46]. These works demonstrate that data-driven predictors can achieve strong accuracy, especially when training and testing conditions are well aligned.

The remaining gap is not simply replacing empirical propagation models with more accurate regressors. The more demanding requirement is to build predictors that remain traceable and robust when deployment conditions shift. A model used for radio planning, anomaly diagnosis, or digital-twin-assisted optimisation should expose whether an error or coverage change is associated with path loss, antenna alignment, site geometry, clutter, or residual effects. Conventional black-box models do not provide this inference pathway. This motivates the IINN framework developed in Chapter 3 [12], [14], [16].

1.2.3 Explainability, Physics-Guided Learning, and Structural Traceability

Post-hoc explainability methods such as SHAP, LIME, saliency maps, and perturbation analysis can estimate influential inputs after a model has been trained [19], [20], [47], [48], [49]. These methods are useful for communicating feature influence, but they do not necessarily reveal the actual sequence of internal operations that produced a prediction. In wireless propagation modelling, this distinction matters: knowing that distance is important is not the same as tracing how distance-dependent path loss, antenna attenuation, and site/environment terms combine during inference [7], [12], [14].

Other structured approaches introduce domain knowledge through physics-informed loss terms, residual constraints, deep unfolding, neuro-symbolic structure, Equation Learner, or Kolmogorov–Arnold networks [10], [25], [26], [27], [50], [51], [52], [53]. These approaches show that structure can improve learning, but many do not organise the inference pathway as an inspectable received-power decomposition. The gap addressed by IINN is therefore architecture-level structural traceability: analytical variables, coefficients, and operators are mapped into trainable NN components so that intermediate propagation terms remain visible during inference [8], [9], [12], [54].

1.2.4 Limited Data, OOD Robustness, and Transfer Learning

Wireless datasets are often limited, biased, or tied to a particular simulator, city, carrier frequency, or measurement campaign. Data augmentation, GAN-based propagation synthesis, calibration, and domain-adaptation methods can reduce sample requirements or improve transfer [23], [24], [55], [56], [57]. These methods are complementary to IINN, but they do not by themselves determine which internal components of a model should change when a target deployment changes [12], [21], [22].

Conventional transfer learning often freezes and fine-tunes layers using empirical search. This can be effective, but it is not naturally connected to wireless-domain variables such as base-station height, antenna tilt, sector orientation, or frequency. Chapter 5 addresses this gap by using the interpretability of IINN to guide layer selection: invariant components are frozen and domain-variable components are fine-tuned. Thus, the transfer-learning contribution is not a new propagation architecture; it is an experiment showing that the IINN architecture introduced in Chapter 3 and evaluated in Chapter 4 can make adaptation more principled [13], [28], [29], [30].

1.2.5 Central Knowledge Gap

Across these literatures, the central knowledge gap is the absence of a wireless learning framework that jointly provides: (i) practical predictive capability, (ii) architecture-level interpretability rather than only post-hoc explanation, (iii) robustness under limited-data and OOD deployment shifts, (iv) component-level structural traceability, and (v) a principled basis for transfer learning. This thesis addresses that gap by progressing from a conventional CNN spectrum-awareness model, to IINN for structurally traceable propagation modelling, to IINN-guided transfer learning [11], [12], [13], [16].

1.3 Research Problem

The thesis addresses the following overarching research problem [1], [5], [7], [12]:

How can wireless learning systems be designed so that they remain accurate, data efficient, interpretable, and adaptable under limited-data and shifted-deployment conditions, while preserving a meaningful connection between internal computation and domain physics?

This problem is decomposed into four specific challenges [5], [17], [21], [28].

- P1: **Data dependency.** Conventional DNNs require representative labelled data to learn physical relationships that are already partially known from wireless theory [5], [6], [18].
- P2: **Opacity.** The internal computations of black-box models do not correspond to identifiable wireless propagation or network-control mechanisms [7], [19], [20].
- P3: **Distribution shift.** Models trained in one environment may degrade when tested under different city, frequency, antenna, clutter, or propagation-label conditions [21], [22], [23].
- P4: **Unprincipled adaptation.** Conventional transfer learning often relies on trial-and-error layer selection rather than domain knowledge about what actually changes between source and target environments [13], [28], [29].

1.4 Research Aim and Questions

The aim of this thesis is to develop and evaluate a structure-aware neural modelling framework for wireless systems, in which domain knowledge is incorporated into the model architecture before training. The thesis is guided by the following research questions [8], [9], [12].

- RQ1: Can conventional deep learning provide practical wireless intelligence, and what limitations does it expose when interpretability and deployment robustness are required [4], [7], [11]?
- RQ2: Can analytical wireless-domain knowledge be converted into a NN architecture whose internal components remain physically meaningful and trainable [8], [9], [12]?
- RQ3: Does a structure-aware NN architecture improve data efficiency and robustness under limited-data and out-of-distribution conditions [12], [18], [21]?
- RQ4: Can architectural interpretability be used not only for explanation, but also to make transfer learning more principled and efficient [13], [28], [30]?

1.5 Thesis Overview

The thesis develops its argument through four technical contribution chapters [11], [12], [13], [16].

Chapter 2 begins with a conventional data-driven wireless AI problem: spectrum awareness for vehicular communication. A lightweight deep learning modulation-classification framework is used to support modulation-order selection under different channel conditions. This chapter demonstrates that deep learning can be practically useful in wireless decision-making. At the same time, it exposes the limitation that motivates the rest of the thesis: the CNN classifier can perform well, but its internal reasoning is not physically transparent [4], [7], [11].

Chapter 3 introduces the Innately Intelligent Neural Network (IINN) construction methodology. IINN is a domain-informed modelling approach in which analytical relationships are decomposed into neural components. Instead of using one undifferentiated feature vector, scalar physical variables are processed through structured computational branches corresponding to known mathematical operations. This chapter explains the motivation, related work, construction rules, and wireless propagation instantiation of IINN [8], [9], [12].

Chapter 4 evaluates the IINN propagation model under data scarcity, city shift, frequency shift, and propagation-label shift. It also presents the structural-traceability evidence showing how IINN exposes intermediate propagation terms, learned analytical coefficients, and section-wise calculation pathways [5], [12], [21].

Chapter 5 extends IINN to transfer learning. Because IINN layers correspond to physical terms, the model can identify which components are expected to change between source and target domains. This enables domain-aware fine-tuning: invariant components are frozen and variable components are adapted. The result is a transfer learning procedure that is not based on exhaustive layer-search heuristics, but on the physical meaning of the model components [13], [28], [29].

Table 1.1: Role of each chapter in the thesis argument.

Chapter	Technical focus	Main evidence	Role in thesis argument
Chapter 2	Deep learning spectrum awareness for vehicular communication	CNN-based modulation classification supports high detection probability under varying channel settings	Demonstrates the practical value of correlation-driven AI, while exposing the absence of structural interpretability.
Chapter 3	IINN construction methodology	Equation-to-architecture design rules and wireless propagation instantiation	Introduces architecture-level domain embedding and explains how physically meaningful components are built before training.
Chapter 4	IINN propagation experiments and ablation	Data-scarcity, city-shift, frequency-shift, propagation-label-shift, and structural-traceability evidence	Shows how IINN can remain structurally traceable without being tied to one propagation-label model.
Chapter 5	IINN-based transfer learning	Domain-aware fine-tuning improves adaptation under base-station height and sector changes	Shows that interpretability is operationally useful because it enables principled adaptation.

1.6 Contributions

The thesis makes four main contributions [11], [12], [13], [16].

- C1: A diagnostic demonstration of correlation-driven wireless AI.** The thesis first demonstrates that a CNN-based spectrum-awareness system can support modulation classification and CQI-oriented decision-making in vehicular communication. This establishes the practical usefulness of deep learning while motivating the need for models that are more interpretable and physically traceable [4], [7], [11].
- C2: A domain-informed IINN architecture design methodology.** The thesis proposes the IINN framework, in which physical variables, analytical relationships, and mathematical operators are mapped into neural architecture before training. This shifts the design philosophy from heuristic architecture selection to equation-informed model construction and enables structural traceability of the prediction pathway [8], [9], [12].

- C3: Robust and traceable wireless propagation modelling under shift.** The thesis evaluates IINN under city, frequency, data-scarcity, and propagation-label-shift settings. It also shows that structural traceability can expose intermediate propagation terms and learned coefficients, giving the model an inspectable prediction pathway rather than only a final RSRP estimate [5], [12], [21].
- C4: A domain-aware transfer learning strategy.** The thesis shows that IINN interpretability can guide transfer learning. Instead of choosing layers to fine-tune by trial and error, the proposed approach identifies variable and invariant model components from domain knowledge and adapts only the physically relevant layers [13], [28], [30].

1.7 Scope of the thesis

The claims of this thesis are intentionally bounded. IINN is not presented as a universal replacement for all machine learning methods. Flexible models such as LightGBM, XGBoost, and deep ensembles may remain preferable when large, representative datasets are available and the primary objective is raw predictive accuracy. IINN is most relevant when three conditions hold: the task is governed by at least approximate physical or empirical structure, labelled data are limited or expensive, and interpretability or controlled adaptation is operationally important [5], [12], [58], [59].

The thesis also distinguishes between two forms of interpretability. Post-hoc explainability helps communicate or approximate the behaviour of an already-trained model. Architecture-level interpretability changes how the model is constructed. IINN belongs to the second category. Its goal is not merely to explain a black box after training, but to reduce the black-box character of the model by design [7], [12], [15], [17].

Finally, the evaluation is largely simulation based. This enables controlled comparisons across data scarcity, deployment shift, and propagation-label shift, but it does not replace validation on live-network measurements. Real-measurement validation is therefore treated throughout the thesis as a critical future step rather than as an already established result [3], [5], [6].

1.8 Structure of the thesis

The chapters are organised as follows [11], [12], [13], [16].

Chapter 2 presents a deep learning-based spectrum-awareness approach for vehicular communication. The chapter introduces spectrum monitoring, modulation classification, a CNN-based classifier, and a CQI-oriented decision mechanism. It demonstrates the practical value of data-driven AI in wireless systems while motivating the need for more interpretable learning architectures [4], [11].

Chapter 3 introduces the IINN framework for wireless propagation modelling. The chapter develops the equation-to-architecture design process and constructs the propagation-oriented IINN methodology [8], [12].

Chapter 4 evaluates the IINN framework. The chapter discusses the simulation data, benchmark protocol, OOD robustness experiments, propagation-label-shift evaluation, training-stability check, and structural-traceability evidence [12], [21].

Chapter 5 extends IINN to transfer learning. The chapter proposes domain-aware fine-tuning and shows how architectural interpretability can be used to decide which layers should be adapted under changes in base-station height, sector configuration, and deployment conditions [13], [28].

Chapter 6 synthesizes the four technical contribution chapters and concludes the thesis. It summarizes the thesis claims, states the main technical findings, discusses limitations and scope, and presents future research directions including real-measurement validation, cross-domain IINN construction, automated equation decomposition, and IINN-LLM explanation interfaces [3], [9], [60].

1.9 Publications

The thesis is supported by the following publications and manuscripts [11], [12], [13], [16], [56].

Published Work

- **S. B. A. Zaidi**, M. A. Shawky, A. Taha, Q. H. Abbasi, M. A. Imran, and S. Ansari, “*An Efficient Deep Learning-Based Spectrum Awareness Approach for Vehicular Communication*,” in Proc. IEEE Wireless Communications and Networking Conference (WCNC), Glasgow, UK, 2023, pp. 1–6, doi: 10.1109/WCNC55385.2023.10118615 [11].
- **S. B. A. Zaidi**, W. Raza, H. N. Qureshi, M. A. Imran, A. Imran, and S. Ansari, “*A Domain-Aware Framework for Interpretable and Resilient Propagation Models: Enabling Digital Twins for Wireless Networks*,” in Proc. IEEE 99th Vehicular Technology Conference (VTC-Spring), Singapore, 2024, pp. 1–6, doi: 10.1109/VTC2024-Spring62846.2024.1068364 [16].
- **S. B. A. Zaidi** and S. Ansari, “*Designing Interpretable AI-Driven Propagation Models for Advanced Wireless Networks*,” in Proc. IEEE AP-S/INC-USNC-URSI, Florence, Italy, 2024, pp. 1847–1848, doi: 10.1109/AP-S/INC-USNC-URSI52054.2024.10687134 [54].
- W. Raza, **S. B. A. Zaidi**, M. U. B. Farooq, H. N. Qureshi, and A. Imran, “*Refining Wireless Propagation Models Using Domain-Informed GANs Amid Data Scarcity*,” in Proc. IEEE 100th Vehicular Technology Conference (VTC-Fall), Washington, DC, USA, 2024, pp. 1–6, doi: 10.1109/VTC2024-Fall63153.2024.10757539 [56].
- **S. B. A. Zaidi**, W. Raza, M. U. B. Farooq, S. Ansari, A. Imran, and M. A. Imran, “*A Novel Innately-Intelligent Transfer Learning Framework for Wireless Networks and Beyond*,” in Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Istanbul, Türkiye, 2025, pp. 1–6, doi: 10.1109/PIMRC62392.2025.11274900. [13].

Accepted / Under Review

- **Zaidi, S. B. A.**, et al., “*IINN: Innately Intelligent Neural Networks for Interpretable and Robust Wireless Propagation Modelling*,” under revision for IEEE Transactions on Mobile Computing [12].

Conference Posters / Presentations

- **Zaidi, S. B. A.**, “*An Efficient Deep Learning-based Spectrum Awareness Approach for Vehicular Communication*,” Poster presentation, SANS Conference, [James Watt School of Engineering, University of Glasgow], [2023] [11].

1.10 Chapter Summary

This chapter has introduced the central motivation of the thesis and reviewed the literature needed to position the work: spectrum awareness, learning-based propagation modelling, post-hoc explainability, domain-informed learning, limited-data robustness, and transfer learning. The resulting gap is the need to move beyond purely correlation-driven wireless AI toward models that are accurate, interpretable, data-efficient, structurally traceable, and robust under deployment shifts. The next chapter begins this progression by examining a representative deep learning solution for wireless spectrum awareness, establishing both the practical value and the limitations of conventional data-driven NN models [4], [5], [7], [11].

Chapter 2

An Efficient Deep Learning-based Spectrum Awareness Approach for Vehicular Communication

Intelligent transportation systems require reliable information exchange between network terminals in different vehicular communication environments. Efficient use of the dedicated spectrum is crucial for maximising communication performance, and this requires adapting the modulation order to changing channel conditions. This chapter proposes a lightweight spectrum-awareness methodology that combines wideband spectrum monitoring with deep learning-based modulation classification. The channel-quality indicator block uses the probability of detection, denoted consistently in this chapter as P_d , as a forward indicator for modulation-order selection. Using a 3D stochastic vehicular channel, the classifier is evaluated under different speed, variance, signal-to-noise ratio (SNR), and rural/urban channel conditions. The results show that conventional CNN-based wireless intelligence can be practically useful, while also motivating the shift toward the structurally traceable IINN approach developed in Chapter 3.

2.1 Preliminaries and Overview

This section reviews spectrum monitoring and modulation-classification methods from the perspective of the contribution made in this chapter. The aim is not only to list prior work, but to identify the technical gap that motivates a lightweight image-based CNN classifier for vehicular spectrum awareness.

2.1.1 Spectrum monitoring techniques

According to the literature presented by Gupta et al. [61], two major schemes are commonly used for spectrum monitoring: receiver-statistics-based monitoring and energy-ratio-based monitoring. Receiver-statistics approaches detect the presence of primary users by comparing quantities such as bit-error counts against a threshold. Their main limitation is that the statistic can be distorted by receiver hardware impairments, noise uncertainty, synchronization error, and channel fading; the detector can therefore confuse a poor receiver condition with true spectrum occupancy. Energy-ratio approaches avoid some decoding assumptions by observing subcarrier or band energy, often through sliding windows. Their limitation is that energy alone does not identify the modulation order or the suitability of a channel for a particular link adaptation decision.

More recent systems extend monitoring to wider frequency ranges. Robert et al. [62] propose a real-time monitoring system that gathers data from a three-axis antenna on three synchronised receiving channels. Shiba et al. [63] introduce a multi-frequency sampling network for wideband spectrum monitoring in internet-of-things applications. These works demonstrate that wideband observation is useful, but they do not by themselves solve the modulation-classification problem needed by the CQI engine in this chapter. Therefore, this work combines wideband monitoring with an image-based modulation classifier. The NI-LabView wideband monitoring block is used as the spectrum-observation stage [64], while the CNN classifier provides the modulation-awareness stage.

2.1.2 Modulation classification techniques

The two major categories of modulation classification techniques are likelihood-based (LB) [31], [32] and feature-based (FB) [33], [34], as summarised in Fig. 2.1. LB methods formulate modulation recognition as a multiple-hypothesis test. They can be statistically strong when the signal and channel assumptions are accurate, but they are computationally expensive and often require prior knowledge of channel parameters. This is a serious limitation in V2V environments where Doppler, scatterer motion, and multipath structure vary quickly.

FB methods reduce complexity by first constructing signal descriptors and then applying a classifier. Shi et al. [31] reduce LB complexity using Gauss-Legendre and Gauss-Hermite quadrature rules, while Zheng et al. [32] use a maximum average likelihood algorithm for OFDM modulation-order detection. These methods remain sensitive to modelling assumptions

and active-subcarrier patterns. Lee et al. [33] convert wireless-signal characteristics into 2D images and classify them using a CNN; this supports the image-based direction used here, but their evaluation does not cover all signal forms and does not fully stress vehicular channel variation. Zhang et al. [34] use feature selection with a tree-like structure and a multi-layer smooth SVM, but this depends strongly on the selected features.

The limitation across much of this literature is that modulation classification is often evaluated mainly as an SNR problem. In vehicular links, however, SNR is only one source of uncertainty. Mobility, scatterer speed, channel variation, and urban/rural multipath structure also change the constellation geometry. The unique contribution of this chapter is therefore not the claim that CNNs are new for modulation classification; rather, it is the integration of constellation-image classification into a spectrum-awareness pipeline and the evaluation of detection probability under channel variation, vehicle speed, SNR, and urban/rural propagation conditions.




Figure 2.1: Techniques for modulation classification.

2.2 Role of this Chapter in the Thesis


The broader literature on learning-based coverage and propagation modelling is reviewed in Section 1.2.2, because it motivates the IINN propagation-modelling contribution developed in Chapter 3. This chapter is intentionally narrower. It focuses on spectrum awareness for vehicular communication and uses a CNN-based modulation classifier to show both the practical value and the limitation of conventional data-driven wireless AI [4], [7], [11].

The model and dataset used here are different from the IINN studies in Chapter 3 and Chapter 4. This chapter performs modulation classification from constellation-image inputs generated under a 3D stochastic V2V channel. In contrast, Chapter 3 and Chapter 4 address RSRP prediction from radio-planning and propagation features. The comparison across chapters is therefore not a like-for-like model benchmark. Instead, Chapter 2 establishes the thesis motivation: a CNN can be useful for wireless decision-making, but its internal reasoning is not physically traceable. The later IINN chapters respond to that limitation by changing the model-construction principle [5], [11], [12].


2.3 Introduction

The adoption of intelligent transportation systems contributes to decreasing the number of road fatalities and improving transportation safety [65]. By allowing wireless communication between different vehicular ad-hoc network (VANET) terminals, traffic-related messages are shared among vehicles[66]. These messages include critical traffic information such as location, speed, heading, etc. VANET is a mobile communication technology used in the vehicle domain that facilitates vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications[65], [66]. Due to the recent influx of wireless technologies, the dedicated short-range communication (DSRC) spectrum (5.85 to 5.925 GHz) becomes fully occupied, limiting the communication throughput, especially in traffic congestion scenarios. This matter motivated researchers to turn to the unoccupied frequency spectrum to increase the network capacity [67]. Optimizing communication performance in VANET requires terminals to have a good observation of the channel spectrum, referred to as “spectrum awareness.” In general, spectrum awareness includes interference environment identification and modulation classification [61]. Accordingly, it is imperative to implement reliable spectrum monitoring techniques in order to ensure reliable communication.

Successful classification of the modulation order helps identify the most suitable modulation choice for the current channel condition. To keep notation consistent, this chapter reports classification performance as the probability of detection, P_d . The value of P_d is used as a forward channel-quality indicator (CQI) in the proposed engine, as shown in Fig. 2.2(a).



(a) Block diagram of the modulation classification.



(b) CQI optimisation engine.

Figure 2.2: The proposed spectrum awareness approach.

One important factor that can affect data transmission is the modulation order, which refers to the number of possible states that the carrier signal can take. A higher modulation order allows more data to be transmitted, but it can also make the signal more susceptible to noise, fading, and interference. In this chapter, the classified modulation classes are 2-PSK, 4-PSK, 8-PSK, 16-PSK, 8-QAM, and 16-QAM. By using a CQI engine to adapt the modulation order, it is possible to improve the quality and reliability of data transmission over the forward channel. The CQI decision is based on P_d rather than the traditional integer CQI scale ranging from 0 (poorest channel quality) to 15 (best channel quality). A high value of P_d indicates a better channel condition and supports maintaining or increasing the modulation order, while a lower value suggests that a more robust modulation should be selected, as shown in Fig. 2.2(b). Due to vehicular channel fluctuation between rural and urban areas, the modulation order must be optimised between communicating terminals while keeping the probability of error, P_e , acceptable. The channel's unpredictable behaviour, including line-of-sight and non-line-of-sight variation, carrier-frequency offset, and additive noise, increases P_e and motivates a reliable modulation-classification technique for fast- and slow-fading vehicular channels.

2.3.1 Evaluation Metrics and Modulation-Switching Rule

For a given modulation class k , the probability of detection is defined as

$$P_d(k) = \frac{N_{\text{correct}}(k)}{N_{\text{test}}(k)}, \quad (2.1)$$

where $P_d(k) \in [0, 1]$, $N_{\text{correct}}(k)$ is the number of correctly classified test samples for class k , and $N_{\text{test}}(k)$ is the total number of test samples for that class. Tables in this chapter report $100P_d(k)$ as a percentage for readability. When the same test population is considered, the corresponding class error probability is

$$P_e(k) = 1 - P_d(k). \quad (2.2)$$

The remainder of the chapter therefore uses P_d consistently for reported classifier performance. In the CQI engine, P_d is used as a reliability indicator rather than as a direct bit-error-rate estimate. A practical modulation-switching rule can be defined as follows: if $P_d \geq 0.90$, the current or a higher-throughput modulation can be maintained; if $0.80 \leq P_d < 0.90$, the system

should avoid increasing the modulation order and may hold the current order; and if $P_d < 0.80$, the link should switch to a more robust lower-order modulation. These thresholds are design thresholds for the proposed engine and can be tuned for a target packet-error-rate or latency requirement.

The current state-of-the-art of modulation optimization for mobile and vehicular communication depends on channel probing [68], which consumes high communication overhead. Pilot-based channel estimation schemes are characterized by a high amount of communication overhead and low spectrum utilization. In this challenging scenario, the author contributes the following:

1. For effective observation of unoccupied channels, we apply the spectrum aggregation-based ultra-wideband spectrum monitoring method.
2. By using deep learning, this work develops a lightweight feature-based modulation classification approach that can detect modulation order with a high P_d and an acceptable P_e .
3. Based on a 3-dimension (3D) stochastic vehicular channel [69], the proposed approach was evaluated at various vehicle speeds, signal-to-noise ratios (SNRs), and channel variations in urban and rural areas.

This section reviews the V2V channel modeled in [69] and discusses the designed deep learning model.

2.3.2 Review of the 3D stochastic V2V channel

This study adopts the stochastic vehicular channel modeled in [69]. Consider a scenario in which two vehicles, Alice and Bob, wirelessly communicate at a central frequency f_c . In this case, Bob's received signal is the combination of a number of L components coming from different moving and fixed scatterers, see Fig. 2.3. The l^{th} multipath component has a different fading coefficient and phase delay denoted by a_l and ϕ_l , respectively. Hence, the channel observation at the side of Bob at time t is represented by

$$H_B(t) = \sum_{l=1}^L |a_l| \exp(j\phi_l) \exp(j2\pi v_l t) \quad (2.3)$$

Table 2.1: Channel parameter settings

Channel Parameters	Value
Number of multipath components (L)	Urban: 16, Rural: 5
Speed of Tx/Rx	10, 20, 30 m/s
Speed of the scatterer	30 m/s
Azimuth departure/arrival angles ($\alpha_{A(B),l}$)	$U[-\pi, \pi]$
Departure/arrivals' elevation angles ($\beta_{A(B),l}$)	$U[-\pi, \pi]$
Incident reflected scatterers' angles ($\alpha_{1(2),l}$)	$U[0, \pi/3]$
Weibull distribution - scale coefficient (ρ)	2.985
Weibull distribution - shape coefficient (a_w)	0.428

where the doppler parameter v_l combines Alice's, Bob's, and the l scatterer's doppler shifts, denoted by $v_{A,l}$, $v_{B,l}$, and $v_{S,l}$, respectively, as follows.

$$v_l = v_{A,l} + v_{B,l} + v_{S,l} \quad (2.4)$$

where

$$\begin{aligned} v_{A(B),l} &= u_{A(B)\max} \frac{f_c}{c} \cos \alpha_{A(B),l} \cos \beta_{A(B),l} \\ v_{S,l} &= u_S \frac{f_c}{c} (\cos \alpha_{1,l} + \cos \alpha_{2,l}) \end{aligned} \quad (2.5)$$

Here, $H_B(t)$ is a scalar complex channel coefficient observed at time t , while the sum over l accounts for the vector of L multipath contributions. Similarly, v_l in Equation (2.4) is a scalar Doppler contribution for the l^{th} path, not a vector quantity. The parameters $u_{A(B)\max}$, $\alpha_{A(B),l}$, and $\beta_{A(B),l}$ denote Alice's and Bob's maximum vehicle speeds, azimuth angles, and elevation angles of departure/arrival, respectively, while $\alpha_{S,l}$ and $\beta_{S,l}$ denote the scatterers' incident and reflected angles. According to [69], the randomness of scatterer speed u_S follows the Weibull distribution denoted by

$$p_{u_S}(u_S) = \rho u_S^{a_w-1} \exp(-\rho u_S^{a_w}/a_w) \quad (2.6)$$

where ρ and a_w are the Weibull scale and shape parameters, respectively. The symbol a_w is used here to avoid confusion with the multipath fading coefficient a_l in Equation (2.3). In this study, we modeled the V2V channel with parameter settings listed in Table I.




Figure 2.3: Vehicular channel model representation for spectrum awareness.

2.3.3 The proposed image-based deep learning model

This subsection details the CNN classification block highlighted in yellow in Fig. 2.2(a). The received complex baseband symbols are first represented in the in-phase/quadrature plane and saved as constellation images. This image representation is used because channel impairments such as fading, Doppler spread, phase distortion, and additive noise appear as visible geometric changes in the constellation pattern. A CNN is therefore appropriate because local convolutional filters can learn spatial distortions of constellation clusters without requiring hand-crafted signal features. In contrast, using the raw sample sequence alone would require the network to learn invariances to rotation, spread, and cluster displacement directly from a small dataset [4], [11].

The constellation images used in this study are black-and-white intensity images rather than colour images. Colour does not encode an additional physical variable in a constellation plot; the useful information is the location, spread, and density of points in the I/Q plane. Using grayscale therefore reduces unnecessary input variation and prevents the classifier from learning artefacts of plotting colour maps. If the implementation toolbox stores the image in an RGB-compatible format, the colour channels are treated only as replicated intensity information and not as separate physical features [4].

Fig. 2.4 presents the flowchart of the proposed image-based model. The input layer accepts constellation images with dimensions 273×328 pixels, followed by six consecutive 2D convolution layers with filter numbers $\{16, 24, 32, 48, 64, 96\}$. Each convolution layer is followed by batch normalization, ReLU activation, and 2D max pooling. The increasing number of filters allows the network to progress from local constellation edges and cluster shapes to higher-level class-specific patterns, while max pooling reduces sensitivity to small displacements caused by noise and fading. The final max-pooling output is connected to a fully connected layer with six classes, corresponding to 2-PSK, 4-PSK, 8-PSK, 16-PSK, 8-QAM, and 16-QAM. The softmax layer is used because the task is single-label multiclass classification: each input image belongs to exactly one modulation class, and the output must be a normalized probability distribution over the six classes. The classification layer is trained using categorical cross-entropy loss:

$$\mathcal{L}_{\text{CE}} = - \sum_{k=1}^6 y_k \log(\hat{y}_k), \quad (2.7)$$

where y_k is the one-hot class label and $\hat{y}_k$ is the predicted probability of class k . Cross-entropy is appropriate because it penalizes confident wrong class assignments more strongly than small probability errors and directly matches the softmax output. The network has 28 layers in total. The Adam optimiser is used with an initial learning rate of 0.001, validation frequency of 50 iterations, maximum 30 epochs, and mini-batch size of 64.

This architecture was chosen as a lightweight CNN suited to the available image dataset. It is deeper than a shallow two- or three-layer classifier, which may not capture the progressive distortion patterns across the six modulation classes, but substantially smaller than generic image-recognition backbones. Compared with LSTM-style sequence models, the CNN is also better aligned with the chosen input representation: once I/Q samples are converted to constellation images, the relevant structure is spatial rather than temporal. LSTM models can be useful when long time sequences are retained, but they add recurrent-state computation and are less natural for static constellation snapshots. The selected CNN therefore represents a trade-off between accuracy, computational cost, and interpretability of the input representation [1], [4].

For repeatability, the dataset is stratified by modulation class before training and validation. Model selection should be checked using stratified k -fold cross-validation; in this chapter, a five-fold protocol is used to verify that each class contributes samples to every fold, while the final reported experiments use the fixed channel-condition test settings in Table 2.2–Table 2.6. The number of epochs, mini-batch size, learning rate, validation frequency, and channel-condition settings are stated explicitly so that the experiment can be repeated with the same train/validation/test construction.




Figure 2.4: The proposed lightweight CNN architecture for constellation-image modulation classification. The convolutional blocks learn local I/Q cluster geometry, while the softmax output produces six class probabilities.

2.4 Spectrum awareness approach

This section describes the overall spectrum-awareness pipeline used in the chapter. The approach has two linked functions. First, spectrum monitoring observes a frequency span and identifies whether energy is present around the intended DSRC band. Second, modulation classification takes the received baseband symbols from the candidate channel and predicts the modulation class so that the CQI engine can decide whether the current modulation order is reliable [4], [11], [61].

In this context, wideband monitoring refers to observing a frequency span that is wider than the individual communication channel of interest, so that occupied and unoccupied portions of the band can be compared. Narrowband monitoring refers to observing only the channel immediately around a specific carrier. Wideband monitoring is useful for spectrum awareness because it gives a broader occupancy view, while narrowband processing is useful for detailed demodulation and classification once a candidate carrier has been selected. For the implementation, we use the NI-LabView wideband monitoring example [70] with the USRP X310 first RF channel and transmit at $f_c = 5.9$ GHz using the second RF channel. The spectrum-monitoring plot previously shown for 5.88–5.91 GHz has been removed because the main contribution of this chapter is the modulation-classification and CQI decision pipeline, not visual inspection of a single spectrum snapshot.

2.4.1 The training phase

In this phase, the training dataset setting, depicted in Fig. 2.5, is adjusted for acquiring the dataset used for training the CNN. It is noteworthy that the training dataset includes the internal additive complex Gaussian noise of the USRP X310 receiver channel. We use an OFDM communication system at $f_c = 5.9$ GHz for DSRC, with 256 subcarriers, 64 cyclic-prefix samples, and 125 subcarriers holding the transmitted data. The six modulation classes are 2-PSK, 4-PSK, 8-PSK, 16-PSK, 8-QAM, and 16-QAM. For each modulation class, 150 constellation images with dimensions 273×328 pixels are acquired. Training samples for different constellations in the polar coordinates (in-phase and quadrature axes) are presented in Fig. 2.6.

The CNN is trained with the categorical cross-entropy loss in Equation (2.7). Validation loss is checked every 50 iterations during training to monitor convergence and guard against overfitting. The maximum number of epochs is 30, the mini-batch size is 64, and Adam is used with an initial learning rate of 0.001. For repeatability, the class split is stratified so that each modulation type is represented in training and validation subsets. A five-fold cross-validation protocol is used for model-selection sanity checking, while the final reported results are produced using the fixed channel-condition tests described below. Based on the obtained data, the total training time was [74:13] minutes using a Core-i7 CPU at 2.7 GHz with 16 GB RAM.




Figure 2.5: Implementation block diagram.

2.4.2 The classification phase

In this phase, using the testing dataset setting depicted in Fig. 2.5, we simulated the 3D stochastic V2V channel reviewed above. Using the channel parameter settings listed in Table 2.1 and the complex additive Gaussian noise block, we evaluated the classifier's performance at different SNRs, vehicle speeds, and channel variations in urban and rural areas. For urban and rural areas, we set L to 16 and 5 multipath components, respectively. These components are convoluted with the received OFDM symbol at the receiver side.

The noise variation is controlled through the additive complex Gaussian noise power and the resulting SNR. Lower SNR increases the spread of constellation points around their ideal locations, while higher mobility and channel variation introduce additional phase rotation, Doppler spread, and cluster displacement. The classifier is therefore tested under separate changes in channel variation, speed, SNR, and urban/rural multipath structure to isolate which impairment most strongly reduces P_d .




Figure 2.6: Black-and-white constellation-image training samples for the six modulation classes. The grayscale representation preserves I/Q geometry while avoiding colour features that have no physical meaning in this task.

2.5 Performance evaluation

In this section, we examine how classification performance is affected under different test parameters: channel variation $var.$, vehicle speed, SNR, and urban/rural multipath structure. Four controlled experiments were conducted. In each experiment, three parameters are kept fixed while one parameter is varied. This design does not replace a full factorial statistical study, but it allows the main physical effect of each impairment to be interpreted separately. The desired behaviour is that P_d remains high on the diagonal class decision and that off-diagonal confusion remains low enough for the CQI engine to avoid selecting an unreliable modulation order.

2.5.1 Experiment 1: P_d at different channel variations $var.$

In this experiment, we set the maximum vehicle speeds $u_{A(B)\max}$ in Equation (2.5) to 30 m/s and the SNR to 25 dB in an urban scenario, where $L = 16$ multipath components. The channel variation parameter $var.$ is then set to 0.1, 0.3, and 0.6. Table 2.2 shows that P_d decreases as $var.$ increases. For example, for 8-PSK, the reported value of $100P_d$ decreases from 90.1% at $var. = 0.1$ to 87.1% and 86.5% at $var. = 0.3$ and $var. = 0.6$, respectively.

The physical rationale is that higher channel variation spreads constellation clusters and makes the observed image less similar to the clean class geometry. The effect is stronger for modulations whose decision regions become visually close under distortion. In contrast, 16-QAM remains at 100% in this experiment because its grid-like amplitude/phase structure remains distinctive under the tested variation range. The practical implication is that when channel variation pushes P_d below the design threshold, the CQI engine should avoid increasing the modulation order even if the nominal SNR remains high.

2.5.2 Experiment 2: P_d at different vehicle speeds $u_{A(B)\max}$

In this experiment, the channel variation is fixed at $var. = 0.3$ and the SNR is fixed at 25 dB in an urban scenario. The maximum vehicle speed $u_{A(B)\max}$ is varied between 10, 20, and 30 m/s. Table 2.3 shows that increasing speed generally reduces P_d . For example, for 8-PSK, the reported value of $100P_d$ decreases from 89.1% at 10 m/s to 88.6% and 87.1% at 20 and 30 m/s, respectively.

The reason is that higher speed increases Doppler-related distortion and changes the apparent constellation during the observation interval. The decrease is moderate rather than catastrophic, suggesting that the CNN has learned some robustness to mobility-induced spread. However, this robustness is not uniform across classes: lower-order PSK and mixed amplitude/phase constellations remain more vulnerable because their distorted image patterns can overlap with other classes. This reinforces the need for a CQI rule based on P_d rather than assuming that a high-SNR link always supports a high modulation order.

2.5.3 Experiment 3: P_d at different SNRs

In this experiment, the channel variation is fixed at $var. = 0.3$ and $u_{A(B)\max}$ in Equation (2.5) is fixed at 30 m/s in an urban scenario. The SNR is varied between 15, 20, and 25 dB. Table 2.4 shows that P_d increases with SNR. For example, for 8-PSK, the reported value of $100P_d$ increases from 82.3% at 15 dB to 84.5% and 87.1% at 20 and 25 dB, respectively.

This trend is expected because higher SNR reduces the additive noise spread around constellation points. The result also shows that the classifier is not equally sensitive for all classes. 2-PSK remains relatively weak even at 25 dB because, in the simulated image representation, its distorted clusters can be confused with amplitude-separated classes such as 8-QAM. 16-QAM

remains highly separable in the tested cases, which indicates that the CNN is exploiting the structured grid pattern of the constellation image. The trade-off is that strong image separability in this controlled setting does not guarantee robustness under untested impairments such as hardware nonlinearity or synchronization errors.

2.5.4 Experiment 4: P_d in urban and rural areas

In this experiment, the channel variation is fixed at $var. = 0.3$, $u_{A(B)\max}$ in Equation (2.5) is fixed at 30 m/s, and SNR is fixed at 25 dB. The number of multipath components L in Equation (2.3) is set to 16 for the urban scenario and 5 for the rural scenario, following the stochastic V2V channel model in [69]. Table 2.5 shows that P_d is higher in the rural scenario for all nontrivial classes. For example, for 8-PSK, the reported value of $100P_d$ is 98.5% in the rural scenario and 87.1% in the urban scenario.

The rationale is that the rural setting has fewer multipath components and therefore produces less severe constellation distortion. The urban setting contains richer scattering, which increases overlap between constellation clusters and makes classification harder. This result is important because it shows that the classifier is not only responding to SNR; it is also sensitive to propagation morphology. The limitation is that only two stylized environment types are tested. A broader dataset with measured urban canyons, suburban roads, tunnels, and intersections would be required before using the threshold rule as an operational standard.

Finally, we summarise a case experiment in the form of a confusion matrix, as presented in Table 2.6. These results are obtained at $var. = 0.3$, $u_{A(B)\max} = 30$ m/s, and SNR = 25 dB in an urban scenario. The desirable outcome is a dominant diagonal, meaning that each true modulation class is assigned to its own predicted class with high probability. Off-diagonal entries indicate modulation-switching risk: if the classifier confuses a robust low-order modulation with a higher-order class, the CQI engine may choose a modulation that is too aggressive for the channel.

Table 2.2: P_d at different channel variations

Modulation order	SNR = 25 dB, Speed = 30 m/s, Urban		
	$var. = 0.1$	$var. = 0.3$	$var. = 0.6$
2-PSK	76.9%	76.2%	72.7%
4-PSK	98.7%	98.6%	98.1%
8-PSK	90.1%	87.1%	86.5%
16-PSK	87.5%	86.5%	85.9%
8-QAM	79.3%	78.7%	77.3%
16-QAM	100%	100%	100%

Table 2.3: P_d at different vehicle speeds

Modulation order	SNR = 25 dB, $var. = 0.3$, Urban		
	Speed = 10 m/s	Speed = 20 m/s	Speed = 30 m/s
2-PSK	77.8%	76.9%	76.2%
4-PSK	98.9%	98.9%	98.6%
8-PSK	89.1%	88.6%	87.1%
16-PSK	88.7%	88.1%	86.5%
8-QAM	81.5%	79.5%	78.7%
16-QAM	100%	100%	100%

As shown in the matrix, there are two major confusing cases. In case 1, the network confuses 2-PSK with 8-QAM. In case 2, the network confuses 16-PSK with 16-QAM. Both errors are plausible because noise, mobility, and multipath spread the constellation points, reducing the separation between nearby phase/amplitude structures. In a practical CQI engine, these cases should trigger conservative adaptation: when P_d falls below the target threshold or when confusion is concentrated toward a higher-order class, the system should switch down rather than up.

The lower-SNR effect is visible in Table IV even without a separate off-diagonal confusion matrix. At SNR = 15 dB, P_d decreases for all nontrivial classes: 2-PSK drops to 68.3%, 8-QAM to 73.4%, and 8-PSK to 82.3%. The exception is 16-QAM, which remains at 100% in this experiment because its constellation geometry remains more distinctive under the simulated impairment settings. These results indicate that the unlikely off-diagonal cases become more likely when SNR decreases, particularly for constellations whose noisy clusters overlap in the in-phase/quadrature image plane.

Table 2.4: P_d at different signal-to-noise ratios

Modulation order	Speed = 30 m/s, $var. = 0.3$, Urban		
	SNR = 15 dB	SNR = 20 dB	SNR = 25 dB
2-PSK	68.3%	72.4%	76.2%
4-PSK	91.7%	97.7%	98.6%
8-PSK	82.3%	84.5%	87.1%
16-PSK	84.9%	85.7%	86.5%
8-QAM	73.4%	74.7%	78.7%
16-QAM	100%	100%	100%

Table 2.5: P_d in urban and rural areas

Modulation order	Speed = 30 m/s, $var. = 0.3$, SNR = 25 dB	
	Urban	Rural
2-PSK	76.2%	99.1%
4-PSK	98.6%	100%
8-PSK	87.1%	98.5%
16-PSK	86.5%	100%
8-QAM	78.7%	97.9%
16-QAM	100%	100%

2.5.5 Data-set publication

The data collected during experiment as shown in Fig. 2.7 has been published with the following DOI and is made public for reproducibility and reuse with the DOI: 10.5525/gla.researchdata.1420.

Study
Research
About us
Student life
Alumni
Support us
Contact


University of Glasgow
Subjects A-Z
Staff A-Z
Academic units A-Z

Home > Enlighten > Research Data >

Enlighten Research Data

About
Latest Additions
Search
Browse
Research Data Management
My Research Data

An Efficient Deep Learning-based Spectrum Awareness Approach for Vehicular Communication

[Zaidi, S. B. A.](#), [Shawky, M.](#) , [Taha, A.](#) , [Abbasi, Q.](#) , [Imran, M.](#)  and [Ansari, S.](#)  (2023) *An Efficient Deep Learning-based Spectrum Awareness Approach for Vehicular Communication*. [Data Collection]

Datacite DOI: [10.5525/gla.researchdata.1420](https://doi.org/10.5525/gla.researchdata.1420)

Collection description

Constellation images of different wireless modulation orders (PSK and QAM)

Date Deposited:	14 Apr 2023 13:52
URI:	http://researchdata.gla.ac.uk/id/eprint/1420

 [Additional details](#)

Available Files

Data

 [Readme_1420.docx](#)

Read me

 [Data_1420.zip](#)

Repository Staff Only: [Update this record](#)

Figure 2.7: Dataset deposit; DOI: 10.5525/gla.researchdata.1420

Table 2.6: The confusion matrix at Speed = 30 m/s, $var. = 0.3$, and SNR = 25 dB in urban area

True	2-PSK	76.2	0	0	0	23.8	0
	4-PSK	0	98.6	0	0	0	1.4
	8-PSK	0	0	87.1	12.8	0	0.01
	16-PSK	0	0	0.3	86.5	0	13.2
	8-QAM	0	0	1.3	9.2	78.7	1.8
	16-QAM	0	0	0	0	0	100
		2-PSK	4-PSK	8-PSK	16-PSK	8-QAM	16-QAM
		Probability of detection (%)					

2.6 Chapter Summary

This chapter has demonstrated that deep learning can provide effective wireless spectrum-awareness capabilities under realistic vehicular communication conditions. The proposed CNN uses black-and-white constellation images to classify six modulation classes and uses the resulting probability of detection P_d as a CQI-oriented reliability indicator. The experiments show that P_d is affected not only by SNR, but also by channel variation, vehicle speed, and urban/rural multipath structure. This supports the central technical message of the chapter: modulation selection in vehicular communication should account for channel dynamics rather than rely only on a static SNR view [4], [11].

At the same time, this chapter also exposes important limitations. First, the CNN is a black-box classifier. Although the input image has a clear physical interpretation, the learned convolutional filters are not explicitly tied to channel mechanisms such as Doppler, scatterer motion, phase rotation, or multipath delay. Second, the threshold rule based on P_d is a design rule for this CQI engine and would need calibration against packet-error-rate or bit-error-rate targets before deployment. Third, the dataset is generated from a stochastic V2V model and a controlled USRP-based setup; broader validation with measured vehicular data, additional hardware impairments, and more environment types would be required for operational use. Fourth, the model solves a multiclass modulation-classification problem, whereas the later IINN chapters solve a different RSRP regression and propagation-modelling problem [4], [5], [7].

These limitations motivate the shift made in the next two chapters. The thesis does not move to IINN because the spectrum-awareness CNN is ineffective; it moves to IINN because useful wireless AI also needs structural traceability, robustness under deployment shift, and a way to connect model decisions to physical mechanisms. Chapter 3 therefore changes the design question from “Which generic neural architecture should classify or predict from data?” to “How can known domain structure be built into the architecture before training begins?” [7], [8], [12].

Chapter 3

IINN Construction Methodology and Modelling Framework

This chapter develops the construction methodology of Innately Intelligent Neural Networks (IINNs). The central idea is that a neural network does not need to begin as a generic black-box function approximator when reliable domain structure is available. Instead, physical variables, analytical equations, coefficients, and mathematical operators can be translated into a neural architecture before training begins. The resulting model remains trainable, but its internal pathway is aligned with physically meaningful components [8], [9], [12].

The purpose of this chapter is methodological. It explains why IINN is needed, how it is positioned relative to existing interpretable and physics-guided learning approaches, and how an analytical wireless propagation model can be decomposed into neural components. The data-generation procedure, benchmark evaluation, OOD experiments, ablation study, and structural-traceability audit are presented separately in Chapter 4. This split keeps the construction logic distinct from the experimental evidence and addresses the concern that the original IINN chapter was too long [5], [7], [12].

3.1 Motivation and Problem Definition

Conventional DNNs can learn useful mappings for wireless prediction tasks, but their internal representations are usually not tied to physical mechanisms such as path loss, antenna alignment, site height, clutter, or residual propagation effects. This creates three practical problems. First, the model may require large representative datasets because it must infer all relationships

from examples. Second, the model may be brittle under deployment shift because learned correlations can be specific to the training city, frequency, or simulator. Third, the model is difficult to audit because an erroneous prediction cannot be traced through physically named computational components [5], [7], [21].

IINN addresses these problems by moving domain knowledge from post-hoc explanation into the model architecture itself. The model is not merely regularised by physics after training; it is constructed so that the computation path follows the structure of the domain. For wireless propagation, this means that distance-frequency path loss, antenna-gain terms, height-related effects, clutter terms, and residual corrections can be represented as explicit neural components. The resulting network is interpretable by construction because each intermediate branch has a physical meaning before training starts [8], [9], [12], [54].

This chapter answers the following methodological questions:

- How can analytical domain knowledge be converted into neural-network structure?
- Which parts of a propagation equation should become fixed operators, trainable parameters, or bounded residual components?
- How does IINN differ from conventional feature engineering, post-hoc explainability, PINNs, and generic DNNs?

The chapter first reviews related work, then presents the general IINN design rules, and finally instantiates the method for wireless received-power modelling. The experimental validation of these design choices follows in Chapter 4 [5], [12].

3.2 Related Work

This section reviews prior work on explainability, interpretable models, and domain-informed learning, and highlights the limitations that motivate the proposed approach [7], [8], [9], [17].

3.2.1 Post-Hoc Explainability in Wireless and Data-Driven Models

Post-hoc explainability has been widely adopted to interpret data-driven wireless models after training, particularly through feature-attribution and surrogate explanation methods. In a recent study [71], multiple algorithms including linear regression, k-nearest neighbors, decision tree, random forest, extremely randomized trees, adaptive enhancement, extreme gradient enhancement, LightGBM, categorical boosting, and DNN were evaluated for large-scale path loss prediction using a radio propagation dataset. However, in their study, SHAP analysis was used

as a post-hoc feature importance approximation method, meaning that the explanations were addressed after algorithm development to identify influential features. Although SHAP analysis lists the influential features in the predictions, it offers no structural transparency into how those predictions were generated, leaving the model opaque during both training and inference.

Similarly, while the work in [72] makes an important contribution by integrating explainability into DRL-based V2X resource allocation and optimizing model complexity, it still relies on SHAP-based post-hoc methods for interpretability. This reflects the ongoing reliance on external explanation tools rather than inherently interpretable model designs. For instance, the XAI-CHEST framework [73] enhances interpretability in channel estimation by identifying relevant and irrelevant subcarriers through a perturbation-based noise model. While this approach provides valuable input-level insights, it still relies on externally analyzing a pre-trained closed-box model, limiting transparency into the model’s internal computational processes. This reliance on post-hoc explainability is not unique to wireless communication; similar techniques are widely employed in other domains, such as healthcare, where sensitivity analysis and feature attribution methods are applied after training to interpret model predictions [74].

Overall, these studies demonstrate the usefulness of post-hoc explanation tools, but they do not make the underlying model inherently interpretable. The reasoning process remains external to the trained model and therefore cannot fully reveal the internal computational pathway from input variables to prediction.

3.2.2 Intrinsically Interpretable and Knowledge-Distilled Models

Alternatively, intrinsic or pre-training interpretability methods aim to build transparency into the model architecture from the outset. These models are designed so that explainability and interpretability emerge naturally from the algorithm itself. For example, [75] proposed an adaptive multiple linear regression framework that offers intrinsic interpretability via explicit local and global effect calculations. However, its reliance on adjusting linear relationships limits its ability to capture complex nonlinear dependencies, and its coefficient-based explanations may oversimplify variable interactions. In another example, [76] used feature attribution-based knowledge distillation to enhance the performance of intrinsically interpretable models by distilling feature attribution knowledge from deep networks. However, its interpretability remains partial, as the student’s reasoning reflects approximations of a closed-box teacher’s attributions, potentially inheriting biases or inaccuracies.

Thus, although intrinsically interpretable and distilled models improve transparency compared with purely post-hoc methods, they may still rely on simplified functional assumptions or inherit approximations from closed-box models.

3.2.3 Data Scarcity, Representation Bias, and Synthetic Data Generation

As can be concluded from the above discussion, present-day methods of introducing interpretability remain fundamentally limited, relying on approximations or indirect constraints that fail to fully explain the true decision logic of a model. As such, they fail to deliver models with transparent reasoning aligned with domain principles, leading to a prevailing reliance on purely data-driven learning, where all knowledge is extracted from training data [71], [77]. The current DNN paradigm assumes abundant and representative training data, an assumption often challenged in complex real-world domains. Traditional DNNs are highly dependent on large, high-quality datasets and their performance often deteriorates with limited samples, as they lack strong inductive inferences and must learn domain relationships from available data [78], [79].

Although DL benefits from hierarchical feature extraction in large datasets, its dependence on data makes it fragile in data-scarce settings, particularly in wireless networks, where high-quality measurements require costly and time-consuming drive-test campaigns, and experimentation on live networks is often impractical or impossible [80], [81]. In this context, recent work has explored strategies to mitigate data scarcity in wireless networks, such as increasing tabular datasets using multistage generative adversarial networks (GANs) with an attention-through-segmentation training scheme [55], or initializing GAN generators with data derived from analytical propagation equations to improve synthetic data quality [56]. Although these approaches alleviate the shortage of training samples, they remain dependent on GAN-based data generation and do not offer inherent interpretability or explainability of the resulting models.

Beyond training data scarcity, another challenge lies in acquiring representative and comprehensive datasets, particularly in wireless networks, where factors such as the use of gold-standard parameter ranges in most deployment settings, seasonal variations, hardware diversity, changing mobility patterns, and network topology changes make such collection infeasible [77]. In [82], this is defined as representation bias within the broader context of data imbalance, referring to situations where certain groups are underrepresented due to biased data collection or an uneven underlying distribution.

In [83], an LLM-based framework is presented for interpretable log-based fault diagnosis. By combining model fine-tuning with a fault-oriented log summarization module, the method still depends on domain-specific training data and the reliability of generated explanations may vary across unseen systems and evolving log patterns. In [52], an Equation Learner (EQL) network is integrated with deep learning architectures to perform symbolic regression in an end-to-end manner. The proposed approach depends on predefined functional operators and may face difficulty in learning highly complex real-world relationships that do not match the assumed symbolic structure. In [53], a Kolmogorov-Arnold Networks (KAN)-based method is presented for eavesdropping node localization in Power Internet of Things (PIoT) networks by learning nonlinear mappings from node features to attack locations. However, its performance depends on high-quality feature engineering and controlled data conditions, which can limit robustness in highly dynamic real-world grid environments.

These studies address important aspects of data scarcity, representation bias, and nonlinear modelling. However, most of them either remain dependent on large or carefully curated datasets, rely on synthetic data generation, or provide interpretability only indirectly. They do not directly construct a neural architecture whose internal computation mirrors domain-level physical reasoning.

3.2.4 Physics-Informed and Domain-Guided Neural Learning

However, the argument is to create architectures that endow domain intelligence, where the reasoning process is guided by embedding analytical knowledge rather than learning solely from data. Such models can leverage prior domain constraints to maintain stable performance even when data imperfections occur, thereby mitigating the risks posed by corrupted or inconsistent training inputs. A recent example in this context is the adoption of physics-informed neural networks (PINNs), which integrate known physical laws into the learning process as soft constraints during training [50].

For example, in [84], while the use of PINNs introduces physics-based constraints that could enhance interpretability, the paper does not verify whether learned parameters align with physical expectations or provide sensitivity analyses to illustrate their behavior under controlled variations. Similarly, [85] proposed a PINN-based approach that reduces reliance on labeled data by embedding PDE constraints. However, its interpretability remains implicit, focusing on physics compliance without exposing the model’s internal reasoning. Moreover, its perform-

ance still depends on dense spatial-temporal sampling, which limits applicability in wireless scenarios with sparse or irregular data. While PINNs promote physical consistency and have demonstrated improved performance in low-data regimes, they still rely on closed-box neural architectures. This limits their ability to offer fully interpretable insights into the internal flow of computations.

Therefore, while PINNs and related methods improve physical consistency, they usually impose physics through the optimization objective rather than through an explicitly interpretable computational topology. This limits their ability to provide step-by-step structural explanations of the model’s internal reasoning.

3.2.5 Toward Inherently Interpretable Domain-Embedded Architectures

However, when a model is designed from the outset with built-in self-explanatory mechanisms, both interpretability and explainability become intrinsic to its structure. Such models allow direct inspection, auditing, and tracing of the computational pathway from input to output, fostering greater transparency and trust. As interpretability increases, the model functions as a primary source of knowledge rather than relying solely on data, making its reasoning more transparent, sometimes with a trade-off in predictive accuracy [77].

An analogous principle can be observed in biological neural systems, where useful structural organization exists prior to postnatal learning. Motivated by this perspective, we note two desirable properties that remain challenging for many conventional DNNs: 1) the presence of useful task-relevant priors before exposure to training data, arising from purposeful architectural structure and initialization; and 2) efficient few-shot learning with limited representative data, enabling faster adaptation to new environments. These principles inspire the proposed IINN framework, where domain knowledge is embedded into the network architecture prior to data-driven optimization.

In [57], a dual strategy is presented that combines expert knowledge-based feature compression and disentangled representation learning for low-altitude network coverage prediction under limited-data conditions. However, the method relies on handcrafted domain knowledge and sparse measurement data, which may restrict scalability and adaptability across different deployment scenarios.

These observations motivate a design philosophy in which interpretability is not added after model training, but is embedded into the architecture itself through physically meaningful variables, operators, and computational pathways.

3.2.6 Positioning of IINN Relative to Existing Paradigms

Inspired by these requirements, this work introduces a framework for constructing IINNs. The core idea is to embed domain knowledge directly into the architecture by incorporating pre-known functional dependencies, mathematical operators, and domain-specific relationships. This results in a network whose structure reflects the underlying physics, providing it with an innate inductive bias and aligning its internal reasoning with established domain principles. By starting from a model-aware architecture rather than a randomly initialized one, the learning process becomes guided from inception, reducing the dependence on large annotated datasets for effective training.

This design philosophy yields five core advantages, collectively referred to as the PRIME properties:

- Purposeful – rather than relying on arbitrary design choices, the IINN architecture is intentionally constructed using established domain knowledge rather than heuristic choices;
- Resilient – the embedded physics enable the model to remain robust under training data scarcity and distributional shifts;
- Interpretable – each internal operation corresponds to a meaningful physical computation, ensuring transparent learning pathways;
- Mathematical – the network architecture explicitly mirrors the domain’s functional and analytical relationships; and
- Explainable – intermediate computations can be articulated step-by-step, enabling clear attribution of model behaviour to underlying physical terms.

Together, these properties position IINN as an inherently interpretable alternative to conventional closed-box NNs.

While prior structured-learning approaches often separate heterogeneous propagation factors through specialized subnetworks, IINN instead assigns architectural components according to known analytical relationships, resulting in a more intuitive equation-based design.

Table 3.1: Conceptual and architectural comparison of IINN with related learning paradigms.

Dimension	IINN	Deep Unfolding	PINNs	Classical ML/DL
Foundational idea	Equation-native architecture motivated by pre-structured neural systems and useful priors before learning.	Unrolls iterative optimization algorithms, such as ISTA or AMP, into finite-depth NNs [26], [27].	Approximates governing PDE solutions using physics residuals during training [10], [25].	Learns statistical mappings using trees or NNs without explicit physical constraints [58], [59].
Prior-knowledge mechanism	Analytical relationships are encoded directly into topology and operators.	Algorithmic structure is preserved from an optimization procedure.	Physical laws enter through the loss function using PDE residuals and boundary conditions.	Prior knowledge enters indirectly through feature engineering or input selection.
Nature of domain knowledge	Algebraic and empirical models used in wireless engineering, including propagation and antenna relationships.	Procedural knowledge describing iterative numerical solvers.	Differential operators describing continuous spatiotemporal systems.	Statistical correlations inferred from observed data distributions.
Network structure	Directed acyclic graph reflecting symbolic equation decomposition.	Sequential layers corresponding to solver iterations.	Generic MLP/CNN architectures augmented with automatic differentiation.	Decision trees, ensembles, or fully connected NNs with unconstrained topology.
Trainable parameters	Physically interpretable coefficients, such as path-loss exponents and correction terms.	Learnable algorithmic parameters, such as step sizes and thresholds.	Neural weights approximating solution fields, typically without direct physical meaning.	Statistical weights or split rules without explicit physical interpretation.

Drawing from the key insights identified in the literature survey, Tables 3.1 and 3.2 position IINN relative to deep unfolding, PINNs, and representative data-driven models, highlighting that IINN embeds analytical domain knowledge directly into network topology rather than introducing physics through optimization structure, loss-level constraints, or post-hoc explanations.

To clarify the conceptual positioning of the proposed framework, Table 3.1 and Table 3.2 compare IINN with deep unfolding, PINNs, and classical ML/DL paradigms. The comparison is split into two parts for readability. Table 3.1 focuses on how prior knowledge is incorporated into the model architecture, while Table 3.2 summarizes interpretability, data efficiency, generalization, scalability, and relevance to radio propagation.

Table 3.2: Interpretability, generalization, and deployment comparison of IINN with related learning paradigms.

Dimension	IINN	Deep Unfolding	PINNs	Classical ML/DL
Interpretability	Structural interpretability: internal components correspond to identifiable physical terms.	Limited to optimization dynamics rather than physical subsystems.	Limited to physics-constraint satisfaction; internal representations remain opaque [25].	Usually requires post-hoc tools such as SHAP [71], [86].
Data efficiency	High data efficiency due to a restricted physics-guided hypothesis space.	Moderate; depends on learning stable solver parameters.	Often data-intensive due to collocation-point requirements and loss-balancing challenges.	Typically data-hungry and degrades under severe data scarcity [71].
Generalization under NID	Improved extrapolation due to enforced physical dependencies, such as distance, angle, and height.	Limited; algorithmic invariance does not guarantee physical consistency.	Sensitive to model mismatch and PDE specification errors [25].	Generally weak unless domain shifts are explicitly represented in features.
Computational scalability	Lightweight architecture with small parameter count, suitable for edge deployment.	Scales with the number of unrolled iterations; depth can affect stability.	Computationally expensive for large-scale or high-dimensional domains.	Highly scalable with mature software stacks, such as LightGBM and XGBoost.
Radio-propagation relevance	Aligned with standardized propagation and antenna models used in RAN planning.	Limited because closed-form path-loss equations are not iterative solvers.	Ill-suited for city-scale propagation unless full-wave PDE solvers are used.	Can achieve strong empirical accuracy but lacks physical decomposition [71].
Primary limitation	Requires at least an approximate analytical or empirical model.	Restricted to problems with known iterative algorithms.	Training instability and scalability challenges for large domains.	Poor robustness under data scarcity and distributional shift [71].

3.3 Innately Intelligent Neural Networks: Modeling Framework

3.3.1 General Framework for Designing IINN

This section first presents the general design methodology of the Innately Intelligent Neural Network (IINN) framework. The methodology is then instantiated through a wireless propagation use case, where analytical propagation relationships are decomposed into interpretable neural components. In this formulation, the objective is not to construct a generic closed-box neural network, but to design a model whose internal computational structure reflects known domain relationships before training begins.

3.3.1.1 Domain-Informed Equation Formulation

Let $\mathbf{x} \in \mathbb{R}^n$ denote the set of physical variables, including controllable parameters and exogenous environmental factors. A domain-informed model expresses the target variable as

$$y = f_{\text{dom}}(\mathbf{x}; \boldsymbol{\psi}, \boldsymbol{\kappa}) + \varepsilon, \quad (3.1)$$

where $\boldsymbol{\psi}$ denotes adjustable coefficients that can be learned from data, $\boldsymbol{\kappa}$ denotes fixed constants obtained from domain knowledge or standards, and ε captures unmodelled effects, measurement noise, or mismatch between the analytical approximation and the observed system.

3.3.1.2 Analytical Decomposition into Physical Components

After formulating the domain equation, the next step is to decompose it into its constituent variables, coefficients, constants, and mathematical operators. Each physical parameter is represented as an independent scalar input stream, in contrast to conventional neural networks that typically concatenate all features into a single vector and immediately mix their roles through dense weight matrices. This scalar-based representation preserves the semantic and mathematical identity of each variable within the analytical formulation.

Let the input variables be denoted as

$$x_1, x_2, \dots, x_n \in \mathbb{R}, \quad (3.2)$$

where each x_i represents a distinct physical parameter. These scalar inputs are individually processed through dedicated affine mappings. In this work, such mappings are referred to as parametric linear operations (PLOs), since they implement trainable linear transformations while preserving the physical role of the corresponding input variable.

The set of analytical operations can be represented as

$$\mathcal{O} = \{\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_n\}, \quad (3.3)$$

where each function $\mathcal{O}_i : \mathbb{R}^n \rightarrow \mathbb{R}$ corresponds to an operation such as addition, subtraction, multiplication, division, logarithm, minimum, or maximum. In compositional form, the target can be expressed as

$$y = (\circ_{i=1}^n \mathcal{O}_i)(\mathbf{x}) = \mathcal{O}_n(\mathcal{O}_{n-1}(\dots \mathcal{O}_1(\mathbf{x}) \dots)). \quad (3.4)$$

Thus, the output is obtained by applying a sequence of physically meaningful transformations to the scalar input variables. This decomposition provides the blueprint for constructing an interpretable neural architecture.

3.3.1.3 Mapping Analytical Operations to Neural Components

In conventional deep learning, network architecture is often designed heuristically, with the number of layers, neurons, activation functions, and other hyperparameters selected through trial and error or hyperparameter optimization. IINN adopts a different design principle. Since the domain equation specifies how variables interact through mathematical operations, the architecture is constructed purposefully so that intermediate computations correspond to known physical relationships.

Each analytical operation $\mathcal{O}_i$ is mapped either to a trainable affine PLO, when the operation represents an adjustable coefficient, or to a fixed differentiable operator layer, when the operation corresponds to a known mathematical transformation such as addition, subtraction, multiplication, division, logarithm, squaring, minimum, or maximum. For a trainable affine layer, the operation can be written as

$$L_i(x) = w_i x + b_i, \quad (3.5)$$

where w_i and b_i are the trainable weight and bias of the corresponding PLO. The adjustable coefficients of the analytical model are denoted as

$$\Psi = \{\psi_1, \psi_2, \dots, \psi_m\}. \quad (3.6)$$

These coefficients are incorporated into the neural model as trainable parameters:

$$\begin{cases} w_i, b_i & \text{for } i = 1, 2, \dots, L, \\ \{w_i, b_i\} \approx \Xi_i, \end{cases} \quad (3.7)$$

where Ξ_i denotes the analytical coefficient or prior value associated with the corresponding operation. During training, these parameters are updated through backpropagation, allowing the model to adapt to data while preserving the structural information provided by the analytical formulation.




Figure 3.1: Workflow for constructing an IINN model. The example illustrates how a portion of a radio propagation equation can guide interpretable and domain-aligned architecture design.

3.3.1.4 General Construction Rules for IINN Design

The transformation from an analytical model to an IINN architecture follows a systematic set of construction rules:

1. Known constants or predefined analytical coefficients are implemented as fixed non-trainable nodes or layers.
2. Parameters that vary across environments or are only approximately known are represented as trainable weights.
3. Analytical operations such as addition, subtraction, multiplication, division, logarithms, exponentials, minimum, maximum, or attenuation laws are mapped to differentiable computational blocks.
4. Distinct physical effects, such as distance-dependent loss, antenna alignment, clutter loss, and height-dependent correction, are represented through separate branches or modules.
5. Outputs of individual branches are combined according to the structure of the original analytical model.
6. Optional trainable refinement layers may be added to capture mismatch between the approximate analytical model and measured or simulated data.

The workflow for constructing an IINN model is illustrated in Fig. 3.1.

3.3.2 Wireless Propagation Instantiation

Wireless propagation provides a suitable domain for evaluating IINN because it is governed by well-established physical principles, empirical propagation models, and standardized analytical formulations. In contrast to conventional closed-box models, which must infer propagation relationships entirely from data, an IINN can embed known dependencies such as distance attenuation, frequency dependence, antenna alignment, height effects, and environmental clutter directly into its computational structure.

In this work, the IINN instantiation uses the 3GPP TR 38.900 UMa formulation as an analytical scaffold for constructing the architecture. This use of the 3GPP equation should not be interpreted as restricting the IINN framework to a single propagation model. Rather, it provides a concrete domain equation from which interpretable neural components can be derived and later audited through the structural-traceability analysis.

3.3.2.1 Received Power Formulation

For a mobile network, the received signal power at a UE can be expressed as

$$P_r[\text{dBm}] = P_t - PL + G - L_{\text{feeder}} + L_{\text{other}}, \quad (3.8)$$

where P_r is the received power, P_t is the gNodeB transmit power, PL denotes propagation path loss, G represents the antenna gain, L_{feeder} denotes feeder or cable loss, and L_{other} accounts for additional system-level losses.

3.3.2.2 Path-Loss and Antenna-Gain Components

Following 3GPP TR 38.900 Release 14 [87], the path loss for a 3D Urban Macro (UMa) scenario is expressed as

$$PL[\text{dB}] = \max(PL_{\text{LoS}}, PL_{\text{NLoS}}). \quad (3.9)$$

The line-of-sight path loss is given by

$$PL_{\text{LoS}}(d, f_c) = 22 \log_{10}(d) + 20 \log_{10}(f_c) + 28, \quad (3.10)$$

where d is the transmitter–receiver distance and f_c is the carrier frequency.

The non-line-of-sight component is expressed as

$$PL_{\text{NLoS}}(W, h, h_{bs}, f_c, d, h_{ue}) = 161.04 - 7.1 \log_{10}(W) + 7.5 \log_{10}(h) - \left(24.37 - 3.7 \left(\frac{h}{h_{bs}} \right)^2 \right) \log_{10}(h_{bs}) + 20 \log_{10}(f_c) + (43.42 - 3.1 \log_{10}(h_{bs})) (\log_{10}(d) - 3) - 0.6(h_{ue} - 1.5) - (3.2(\log_{10}(17.625))^2 - 4.97), \quad (3.11)$$

where W denotes street width, h denotes average building height, h_{bs} denotes base-station height, and h_{ue} denotes UE height.

A simplified form of the NLoS expression is

$$PL_{\text{NLoS}}(f_c, W, h, h_{bs}, h_{ue}, d) = 20 \log_{10}(f_c) - 7.1 \log_{10}(W) + 7.5 \log_{10}(h) - 15.06 \log_{10}(h_{bs}) + 3.7 \log_{10}(h_{bs}) \left(\frac{h}{h_{bs}} \right)^2 - 0.6 h_{ue} - 3.1 \log_{10}(d) \log_{10}(h_{bs}) + 43.42 \log_{10}(d) + 31.69. \quad (3.12)$$

Expanding the squared height-ratio term yields

$$PL_{\text{NLoS}}(f_c, W, h, h_{bs}, h_{ue}, d) = 20 \log_{10}(f_c) - 7.1 \log_{10}(W) + 14.9 \log_{10}(h) - 18.76 \log_{10}(h_{bs}) - 0.6 h_{ue} - 3.1 \log_{10}(d) \log_{10}(h_{bs}) + 43.42 \log_{10}(d) + 31.69. \quad (3.13)$$

The antenna gain component is defined using the antenna efficiency ζ and the horizontal and vertical beamwidths B_h and B_v :

$$G_{\text{max}}(\zeta, B_h, B_v) = \zeta D = \zeta \frac{4\pi}{B_h B_v} = \log_{10} \zeta + [\log_{10} 4 + \log_{10} \pi] - [\log_{10} B_h + \log_{10} B_v]. \quad (3.14)$$

The directional antenna attenuation is approximated as

$$A_{\text{att}}(B_h, B_v, \phi_{ue}, \phi_{bs}, \theta_{ue}, \theta_{bs}, \lambda_h, \lambda_v, A_h, A_v) = \lambda_v \min \left[12 \left(\frac{\phi_{ue} - \phi_{bs}}{B_v} \right)^2, A_v \right] + \lambda_h \min \left[12 \left(\frac{\theta_{ue} - \theta_{bs}}{B_h} \right)^2, A_h \right], \quad (3.15)$$

$$\begin{aligned}
P_r(d, f_c, W, h, h_{ue}, h_{bs}, B_h, B_v, \phi_{ue}, \phi_{bs}, \theta_{ue}, \theta_{bs}, \lambda_h, \lambda_v, A_h, A_v) \\
= P_t - \max \left[22 \log_{10}(d) + 20 \log_{10}(f_c) + 28, 20 \log_{10}(f_c) - 7.1 \log_{10}(W) + 14.9 \log_{10}(h) - 18.76 \log_{10}(h_{bs}) \right. \\
\left. - 0.6 h_{ue} - 3.1 \log_{10}(d) \log_{10}(h_{bs}) + 43.42 \log_{10}(d) + 31.69 \right] + 10 \log_{10}(\zeta) + 10 \log_{10} \left(\frac{4\pi}{B_h B_v} \right) \\
+ \lambda_v \min \left[12 \left(\frac{\phi_{ue} - \phi_{bs}}{B_v} \right)^2, A_v \right] + \lambda_h \min \left[12 \left(\frac{\theta_{ue} - \theta_{bs}}{B_h} \right)^2, A_h \right].
\end{aligned} \tag{3.16}$$

Figure 3.2: Received signal power model P_r for a UE in a 3D urban macro environment, based on 3GPP TR 38.900. The expression captures distance-dependent path loss, LoS/NLoS selection, building and street characteristics, antenna heights, and 3D antenna-gain patterns. This analytical form provides the scaffold for constructing the IINN architecture.

where ϕ_{bs} and ϕ_{ue} denote the gNodeB and UE tilt angles, respectively, while θ_{bs} and θ_{ue} denote the gNodeB and UE azimuth angles. The terms λ_h and λ_v denote horizontal and vertical weighting factors, while A_h and A_v denote the maximum horizontal and vertical attenuations.

The resulting received-power expression is summarized in Equation (3.16).

3.3.2.3 Feature Identification and Physical Inputs

The next step is to identify the features that represent the physical and geometric characteristics of the environment. For RSRP prediction, these include propagation parameters, antenna configuration, geometric alignment, height-related features, and environmental descriptors. The identified parameters are listed in Table 4.3. These variables are then mapped into scalar input streams, preserving their physical meaning throughout the architecture.

3.3.3 Illustrative Sub-Network: Vertical Antenna Alignment

To illustrate how the IINN framework constructs intermediate layers and determines appropriate mathematical operations, we focus on the vertical directional antenna alignment term from Equation (3.16). This component computes attenuation due to angular misalignment between the UE and gNodeB antennas. As shown in Fig. 3.3, the analytical term is systematically mapped to an IINN sub-network.

This component involves four variables: ϕ_{ue} , ϕ_{bs} , B_v , and A_v . These variables are defined as distinct input layers and processed through dedicated PLOs. The configuration of each PLO depends on the operation being represented. Since subtraction of two scalar variables yields a scalar output, the corresponding PLOs use a single neuron with linear activation. This preserves

the original scalar values and avoids nonlinear distortion. The linear mapping is expressed as

$$f(x) = wx + b, \quad (3.17)$$

where w and b are trainable parameters.

For example, in the vertical antenna attenuation component, if $\phi_{ue} = -7.4717^\circ$ and $\phi_{bs} = 0^\circ$, the expected output after subtraction is -7.4717° . A linear operation preserves this physical value. In contrast, a nonlinear activation such as the sigmoid function,

$$f(x) = \frac{1}{1 + e^{-x}}, \quad (3.18)$$

would distort the value and produce an output inconsistent with the intended analytical operation.

The effective tilt angle is computed by subtracting ϕ_{bs} from ϕ_{ue} . This difference is then normalized by B_v . The B_v branch is configured without a bias term to prevent unintended offsets in the denominator. The normalized output is squared, optionally scaled according to the physical model, and then passed through a minimum operation with A_v to enforce the maximum attenuation bound specified by the antenna model.

3.3.4 Overall IINN Architecture and Interpretability

The overall internal architecture of the proposed IINN model is depicted in Fig. 3.4. The design originates from the analytical received-power expression and maps each physical dependency to a corresponding neural component. Input variables are first processed through dedicated PLO layers that capture their independent trainable contributions. These outputs are then combined through arithmetic and nonlinear operator layers, including addition, subtraction, division, logarithmic transformation, minimum, and maximum selection.

A layer-wise breakdown of this construction is presented in Table 3.3. PLO layers implement scalar affine mappings, while operator layers reproduce mathematical transformations from the analytical formulation. Custom Lambda layers implement operations such as logarithms or squared normalized ratios, while merge layers implement addition, subtraction, minimum, and maximum operations. Together, Fig. 3.4 and Table 3.3 establish a one-to-one correspondence between analytical propagation components and NN implementation.




Figure 3.3: Illustration of the IINN sub-network constructed from the vertical directional antenna alignment component of the radio propagation equation. The network takes four input variables, ϕ_{ue} , ϕ_{bs} , B_v , and A_v , and processes them through mathematically aligned layers. Each operation, including subtraction, division, squaring, and minimum selection, mirrors the corresponding analytical computation.

Table 3.3: Layer types used in the Keras-based IINN architecture.

Layer type	Functional description
PLO	Dense layer implementing scalar affine mappings with configurable weight, bias, and linear activation.
Addition layers	Merge layers that perform element-wise summation of multiple inputs.
Subtraction layers	Merge layers that compute the difference between inputs.
Min./Max. layers	Merge layers that compute element-wise lower or upper bounds.
Lambda layers	Custom differentiable layers implementing operations such as logarithm, division, squaring, or squared normalized ratios.

By assigning physically meaningful variables to dedicated input streams and structuring the network layers to replicate analytical operations, the IINN model provides transparency into its internal computations. Intermediate activations can be inspected directly, allowing users to audit how individual components contribute to the final prediction. This makes the model inherently interpretable by design, rather than dependent on external post-hoc explanation tools.




Figure 3.4: Schematic diagram of the IINN model constructed from the 3GPP TR 38.900 received-power equation. Variables from the analytical equation are translated into scalar inputs and processed through dedicated PLO layers, followed by mathematical operators such as subtraction, logarithm, division, minimum, and maximum selection.

3.3.5 Structural Traceability and Component Audit

The transparency described above is referred to in this thesis as *structural traceability*. Structural traceability means that the inference pathway can be inspected as a sequence of named analytical and learned components, rather than only as a final prediction and a set of abstract hidden activations. During inference, each IINN branch can export its physical inputs, fixed operator outputs, learned coefficient corrections, and contribution to the final RSRP estimate.

This is different from post-hoc feature attribution. A method such as SHAP can indicate that distance or antenna tilt is influential, but it does not show the actual computation through which a trained model combined path loss, antenna alignment, clutter, and residual terms. In IINN, the explanation is provided by the model pathway itself. A forward trace can show how each propagation component contributes to $\widehat{\text{RSRP}}$, while a backward residual audit can identify which component contributes most strongly to a prediction error. This property is important for radio-planning and digital-twin workflows, where the practical question is often not only whether an error is large, but which physical mechanism is responsible for it.

The structural traceability audit used in this chapter therefore follows three steps:

1. export the intermediate component values for a given prediction;
2. compare the predicted RSRP with the simulator or measurement label to obtain the residual;
3. inspect which exposed component terms dominate the residual pathway and whether the error is associated with path loss, antenna alignment, site/environment contribution, clutter, or bounded residual correction.

This audit does not require training a separate explanation model. It follows directly from the IINN architecture.

3.3.6 Relation to Feature-Engineering and Conventional Neural Models

An alternative approach would be to compute handcrafted relative features, such as angular differences, normalized variables, or interaction terms, and then use a compact DNN to learn the final mapping. While such approaches may improve predictive performance compared with raw-input models, they still rely on hidden representations whose internal reasoning is generally difficult to interpret.

Table 3.4: Qualitative comparison between feature-engineering-plus-MLP approaches and the proposed IINN framework.

Aspect	Feature + MLP	IINN
Interpretability	Low	High
Hidden representations	Opaque	Physics-aligned
Number of parameters	Moderate	Lower
Low-data robustness	Moderate	Higher
Extrapolation ability	Limited	Better

In contrast, IINN embeds known analytical operators directly into the architecture. As a result, intermediate activations retain physical meaning, the number of free parameters can be reduced, and the model does not need to relearn relationships already available from domain knowledge. This is particularly beneficial under limited-data and shifted-environment settings. A qualitative comparison between feature-engineering-plus-MLP approaches and IINN is provided in Table 3.4.

The same construction principles are later revisited in the experimental chapter through a structural-traceability audit. There, learned analytical coefficients and section-wise prediction calculations are exported from the trained IINN, showing how the architecture preserves the traceable nature of the propagation computation.

3.4 Chapter Summary

This chapter has presented IINN as an architecture-level methodology for embedding domain knowledge into neural-network design. The construction process begins with an analytical or empirical equation, decomposes it into physically meaningful terms, maps variables and operations into neural components, and preserves component-level interpretability through named computational branches. The wireless propagation instantiation shows how received-power prediction can be structured around path-loss, antenna-alignment, site/environment, clutter, and residual components rather than treated as a single opaque regression problem.

The next chapter evaluates this construction experimentally. It describes the simulation datasets, benchmark models, OOD robustness experiments, propagation-label-shift analysis, training-stability evidence, and structural-traceability results drawn from the latest IINN paper.

Chapter 4

IINN Propagation Experiments and Structural Traceability

Chapter 3 developed the IINN construction methodology. This chapter provides the experimental evidence for that methodology using wireless propagation and received signal reference power (RSRP) prediction as the case study. It contains the data-generation discussion, benchmark protocol, evaluation metrics, OOD robustness experiments, propagation-label-shift analysis, and structural traceability results [5], [12], [21].

The chapter follows the latest IINN paper storyline. The main claim is not that IINN is universally the lowest-error predictor in every setting. Rather, the claim is that a domain-informed architecture can provide intrinsic structural traceability and stronger robustness under deployment shift and data scarcity, while remaining competitive with black-box baselines. The structural-traceability analysis shows that the model exposes intermediate propagation terms, learned analytical coefficients, and section-wise calculation pathways during inference [5], [12], [22].

4.1 Chapter Organisation

The chapter is organised as follows. The first section describes the simulation environment, dataset composition, urban morphology, terrain, and physical-consistency checks. The second section defines the benchmark models, training protocol, metrics, and OOD evaluation settings. The third section reports cross-city, joint city-frequency, data-scarcity, and propagation-label-shift results. The final sections present the IINN training-stability check, structural-traceability audit, and practical scope of the findings [3], [12], [21].

4.2 Network Simulation Setting for Data Generation and Collection

4.2.1 Motivation for the Wireless Propagation Case Study

The pathloss prediction problem is an ideal candidate to stress-test the expected benefits of IINN, since training data in wireless propagation tasks are often scarce, expensive to collect, and not fully representative of future deployment conditions. Moreover, solutions that aim to optimize network key performance indicators, including capacity, mobility, and throughput, require interpretable propagation models so that the network parameters responsible for performance changes can be identified and adjusted without compromising other objectives. For example, if a pathloss prediction model operates as a black box, as is the case for many representative benchmark models, its practical utility is limited to pathloss estimation only, primarily supporting initial cell planning and network deployment. In contrast, an interpretable or explainable pathloss model can support continuous and intelligent orchestration of network parameters in response to dynamic environmental conditions and traffic demands, thereby enabling substantially greater operational efficiency [5], [6], [7].

4.2.2 Simulation Environment and Deployment Scenario

To evaluate IINN under realistic wireless propagation conditions, we use the network simulation setting shown in Fig. 4.1. The simulation employs two high-fidelity 1-meter-resolution 3D map samples from Brussels, Belgium and Chicago, USA. These network topology configurations are constructed using Atoll by Forsk, a commercial system-level radio network planning and optimization tool [88]. The use of detailed 3D urban maps provides realistic morphology and spatial detail, ensuring that the propagation environment closely reflects practical deployment conditions. The system follows a 3GPP-compliant 5G NR cellular layout, consisting of five macrocell sites per city, each equipped with three-sector gNodeBs, covering a 1 km² area [3], [5].

The 1 km² region represents a standard analysis tile commonly used in professional radio-planning tools such as Atoll for scalable coverage modelling and performance evaluation. Such tiles provide consistent spatial units for signal prediction, interference estimation, and traffic analysis, while enabling larger metropolitan areas to be represented through multiple adjacent tiles. Therefore, the selected tile should be interpreted as a representative urban deployment segment rather than the full municipal extent of Brussels [6], [88].

Table 4.1 lists the simulation settings used to construct these network topologies.

Table 4.1: Simulation settings


Configuration	Parameter	Value
System	Radio access technology	5G NR
	Cell layout	5 Macrocell sites/ City
	Sectors per gNodeB	3 sectors
	Carrier frequency	3300, 3500, 3700 MHz
	Channel bandwidth	10 MHz
Geographic	Propagation model	3GPP 38.900 Urban Macro
	Simulation location	Brussels, Belgium & Chicago, USA
	Simulation area	1.0 km ²
	UE distribution	Poisson distribution
	Propagation matrix resolution	5 meters
	Average building height	20 meters
	Average street width	20 meters
gNodeB	Antenna model	3GPP antenna radiation pattern
	Height	20–37 meters
	Antenna tilt	0°–6°
	Antenna azimuth	360°
	Antenna gain	8 dBi
	Horizontal half-power beamwidth	65°
	Vertical half-power beamwidth	6°
	Transmit power	50 dBm

4.2.3 Dataset Composition and Feature Description

The data generated from the simulation are organized into three components. The first comprises gNodeB information, including geographical coordinates, antenna height, azimuth and tilt angles, transmit power, carrier frequency, and antenna type for all gNodeBs in the simulation area. The second captures environmental information through clutter type, reflecting radio propagation characteristics and terrain features that influence coverage prediction. The third component consists of UE measurements, including UE coordinates, height, azimuth and tilt angles, and the corresponding reference signal received power (RSRP) within the observation



(a) Brussels, 3.7 GHz.



(b) Chicago, 3.3 GHz.

Figure 4.1: Representative 1 m resolution 3D map samples over a 1 km² urban tile used for cross-city generalization experiments. The maps show simulated gNodeB deployment and UE distribution for Brussels at 3.7 GHz and Chicago at 3.3 GHz.

area. The spatial distribution of UEs is modeled using a Poisson distribution to emulate realistic UE placement patterns. The simulation area is divided into spatial bins, with the RSRP of each bin obtained by averaging the RSRP values of UEs located within that region. The dataset used in this paper is made publicly available on GitHub¹.

1. Syed Ali Basit, “IINN,” GitHub, accessed October 12, 2025, <https://github.com/syedalibasit/IINN>

4.2.4 Urban Morphology and Clutter Distribution




Figure 4.2: User distribution across clutter classes for the Brussels dataset.




Figure 4.3: User distribution across clutter classes for the Chicago dataset.




Figure 4.4: Clutter class distribution comparison between Brussels and Chicago datasets. Significant differences are observed in the prevalence of dense urban and high-rise environments, with Brussels exhibiting a higher concentration of complex propagation conditions.

Fig. 4.2 and Fig. 4.3 illustrate the distribution of UEs across clutter classes in the Brussels and Chicago datasets, respectively. The majority of Brussels user locations are concentrated in dense urban and high-rise building environments, indicating a propagation setting dominated by severe shadowing and multipath effects. This distribution highlights the structural complexity of the Brussels scenario used for evaluating model robustness. Compared to Brussels, the Chicago scenario exhibits a more balanced distribution across open, suburban, and urban environments, reflecting a less dense propagation setting. As shown in Fig. 4.4, this mismatch in environmental characteristics constitutes a primary source of domain shift in cross-city and out-of-distribution (OOD) evaluation scenarios.

4.2.5 Terrain, Obstruction, and Propagation Profiles

Fig. 4.5 shows the terrain and clutter profile between Site 0 and Site 3, highlighting both ground elevation and the type of obstacles present along the propagation path. This profile is valuable for understanding non-line-of-sight conditions and shadowing effects, since buildings and vegetation introduce significant variability in received power that cannot be explained by distance alone. By including these terrain projections, the dataset captures realistic propagation scenarios that incorporate both large-scale terrain features and local clutter effects.




Figure 4.5: Terrain and clutter profile between Site 0 and Site 3 for Brussels. The plot illustrates the elevation profile and intervening clutter, including buildings and vegetation, along the direct path, providing insight into propagation obstructions between two locations.

4.2.6 RSRP-Distance Characteristics and Physical Consistency




Figure 4.6: RSRP variation with distance for multiple sites and sectors in the Brussels scenario. The solid lines represent sector-wise median RSRP values, while shaded regions indicate the interquartile range (IQR), capturing variability due to local propagation conditions and clutter effects.




Figure 4.7: Terrain profile between a UE and its associated transmitter in Brussels for Site 1 Sector 2 over a distance of 245 meters.




Figure 4.8: Terrain profile between a UE and its associated transmitter in Brussels for Site 3 Sector 2 over a distance of 557 meters.

The plots in Fig. 4.6 illustrate the expected inverse relationship between RSRP and distance from the serving site. Across all sites, a clear monotonic decay trend is observed, consistent with the physics of path loss in cellular propagation. The sector-wise medians show distinct coverage lobes shaped by antenna azimuth and downtilt configurations, while the shaded IQR regions highlight the variability introduced by multipath propagation and environmental clutter. Notably, certain sites, such as Site 0 Sector 2 and Site 4 Sector 1, exhibit localized fluctuations in RSRP beyond smooth distance-dependent decay, particularly at mid- to long-range distances. These deviations correspond to terrain obstructions or clutter transitions, as confirmed from the terrain profiles shown in Fig. 4.7 and Fig. 4.8, which cause partial blockage and local propagation variability. Conversely, sectors with more favorable line-of-sight conditions, such as Site 2 across all sectors, demonstrate smoother attenuation curves with relatively narrower IQRs, indicating more stable link quality. Overall, these observations demonstrate both the expected large-scale path loss behavior and the influence of site-specific deployment factors, including antenna orientation, downtilt, surrounding clutter, and terrain obstruction. They also validate that the dataset captures realistic propagation dynamics across diverse environments, which is essential for training and evaluating the proposed IINN model.

4.3 Performance Evaluation

4.3.1 Experimental Protocol and Benchmark Models

The performance evaluation compares the proposed IINN with representative black-box and ensemble learning baselines. The evaluated models include Random Forest (RF), Extra Trees (ET), Histogram Gradient Boosting (HistGB), XGBoost, LightGBM, a stabilised DNN, and the proposed IINN model. These baselines are chosen because the objective is to evaluate whether embedding analytical propagation structure provides structural traceability and lower relative OOD degradation compared with modern learnable propagation models, rather than to compare only against fixed empirical propagation equations.

All benchmark models are trained and evaluated using the same experimental protocol to ensure a fair comparison. Specifically, all methods use the same Atoll-generated datasets, identical train/test splits, identical input features, and the same evaluation metrics. For the data-scarcity analysis, identical reduced training subsets are used for each model so that performance differences reflect model behavior rather than differences in data availability.

Table 4.2: Model configurations used for Chapter 3 robustness validation.

Model	Configuration
IINN	Propagation-informed architecture with explicit operator layers and trainable scalar coefficients; RMSE loss; Adam optimiser, learning rate 10^{-3} ; batch size 32; maximum 150 epochs; early stopping on validation loss with patience 20.
DNN	128–64–32 ReLU layers; batch normalisation; dropout 0.2; L2 regularisation 10^{-5} ; linear output; RMSE loss; Adam optimiser, learning rate 10^{-4} ; batch size 64; maximum 300 epochs; early stopping patience 30.
RF	400 trees; unrestricted depth; minimum leaf size 2.
ET	400 randomised trees; unrestricted depth; minimum leaf size 2.
HistGB	400 boosting iterations; learning rate 0.04; maximum 31 leaf nodes.
XGBoost	600 boosted trees; learning rate 0.03; maximum depth 5; subsampling ratio 0.9; column-sampling ratio 0.9; squared-error objective.
LightGBM	800 boosted trees; learning rate 0.03; 31 leaves; subsampling ratio 0.9; column-sampling ratio 0.9.

The experimental setup considers realistic deployment shifts across cities, frequencies, and training-data regimes. The simulated deployment consists of five macro sites with three sectors per site. For Chicago, transmitter heights range from 20–35 m with sector azimuths of $\{0^\circ, 120^\circ, 240^\circ\}$ and downtilt values between 0° – 3° at 3.3 and 3.5 GHz. For Brussels, similar sectorization is used, with transmitter heights ranging from 20–37 m and carrier frequencies of 3.5 and 3.7 GHz. These differences introduce both city-level and frequency-level distribution shifts.

All models use the same training and evaluation splits. Neural-network models are implemented in TensorFlow/Keras, whereas the ensemble baselines use their standard implementations. IINN is trained using RMSE loss, the Adam optimiser with learning rate 10^{-3} , batch size 32, maximum 150 epochs, and early stopping based on validation loss with patience 20. The stabilised DNN uses a 128–64–32 ReLU architecture with batch normalisation, dropout 0.2, L2 regularisation 10^{-5} , a linear output layer, RMSE loss, Adam with learning rate 10^{-4} , batch size 64, maximum 300 epochs, and early stopping patience 30. Five random seeds are used for the reported reruns: 42, 7, 123, 99, and 2024.

The trainable PLO coefficients are initialized from the corresponding analytical coefficients when such coefficients exist. This means that the model begins from a propagation-informed parameterisation rather than a purely random one. Trainable calibration terms that do not have direct analytical counterparts are initialized close to identity mappings, while final output calibration terms are learned from data. This initialization strategy preserves the physical role of the analytical terms at the start of training and lets the optimiser refine them using the observed data.

Table 4.3: Key parameters used for analytical equation development, dataset creation, and IINN training.

Sr.	Parameter	Symbol
1	gNodeB antenna tilt	ϕ_{bs}
2	gNodeB antenna azimuth	θ_{bs}
3	gNodeB height	h_{bs}
4	UE antenna tilt	ϕ_{ue}
5	UE antenna azimuth	θ_{ue}
6	Propagation distance	d
7	Horizontal antenna beamwidth	B_h
8	Vertical antenna beamwidth	B_v
9	Building height	h
10	UE height	h_{ue}
11	Horizontal antenna attenuation	A_h
12	Vertical antenna attenuation	A_v
13	Operating frequency	f_c
14	Transmit power	P_t
15	Street width	W
16	Antenna efficiency	ζ

Direct experimental comparison with some recent frameworks is not always feasible due to implementation availability and dataset differences. However, their underlying design principles are relevant and are therefore qualitatively discussed to better position the proposed IINN framework relative to existing structured learning approaches.

4.3.2 Input Features and Evaluation Metrics

The input feature set consists of UE downtilt, transmitter downtilt, vertical and horizontal beamwidths, vertical and horizontal attenuation, UE and transmitter azimuth, antenna efficiency, 3D distance, carrier frequency, street width, building height, UE height, transmitter height, LoS/N-LoS constants, and transmit power. These parameters are also listed in Table 4.3.

The evaluation metrics include root mean square error (RMSE), mean absolute error (MAE), normalized root mean square error (NRMSE), and out-of-distribution (OOD) degradation. For N test samples, ground-truth received power y_i , and prediction $\hat{y}_i$, RMSE and MAE are defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}, \quad (4.1)$$

and

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|. \quad (4.2)$$

NRMSE normalizes RMSE by the range of the target variable:

$$\text{NRMSE} = \frac{\text{RMSE}}{y_{\max} - y_{\min}}. \quad (4.3)$$

OOD degradation quantifies the relative increase in prediction error under shifted test conditions and is defined as

$$\text{OOD Degradation} = \frac{\text{RMSE}_{\text{OOD}} - \text{RMSE}_{\text{ID}}}{\text{RMSE}_{\text{ID}}} \times 100. \quad (4.4)$$

Here, RMSE_{ID} denotes the in-distribution RMSE, while RMSE_{OOD} denotes the out-of-distribution RMSE.

4.3.3 OOD Robustness Across Deployment Shifts

This subsection follows the latest IINN robustness-evaluation structure. Four complementary settings are considered: cross-city deployment shift, joint city–frequency deployment shift, deployment shift under reduced training data, and propagation-label shift. Experiments I–III use the relative OOD degradation metric in Equation (4.4); Experiment IV uses absolute OOD RMSE because the source and target labels are generated by different propagation-label mechanisms.

4.3.3.1 Experiment I: Cross-City Deployment Shift

Experiment I evaluates cross-city deployment shift from Chicago to Brussels at 3.3 GHz. Since the carrier frequency is fixed, the distribution shift is mainly due to differences in urban morphology, site geometry, transmitter-height range, clutter distribution, and antenna downtilt configuration. As shown in Fig. 4.9, IINN achieves the lowest relative OOD degradation, $(17.18 \pm 2.73)\%$. Among the baseline models, Extra Trees gives the lowest relative degradation, $(37.50 \pm 1.16)\%$, while the remaining baselines range from 39.04% to 60.00%.

The opposite transfer direction, Brussels to Chicago, produces the same qualitative outcome. IINN again achieves the lowest relative OOD degradation, $(37.54 \pm 1.78)\%$, compared with $(46.37 \pm 1.72)\%$ for the strongest baseline. This direction check indicates that the observed robustness is not specific to one source–target ordering.




Figure 4.9: Experiment I: relative OOD degradation under cross-city deployment shift. Results are reported over five random-seed runs.

4.3.3.2 Experiment II: Joint City–Frequency Deployment Shift

Experiment II evaluates a more difficult deployment shift in which both the city and carrier frequency change. The model is trained on Chicago at 3.3 GHz and tested on Brussels at 3.7 GHz. As shown in Fig. 4.10, IINN achieves the lowest relative OOD degradation, $(22.13 \pm 2.91)\%$, whereas the strongest baseline reaches $(45.32 \pm 1.01)\%$. The remaining baselines exhibit larger degradation, ranging from 48.12% for XGBoost to 58.04% for LightGBM.

This joint shift is more challenging than city shift alone because it changes both deployment morphology and system-level frequency conditions. The result indicates that the IINN architecture learns propagation-relevant structure that transfers more reliably when multiple deployment factors change simultaneously.

4.3.3.3 Experiment III: Deployment Shift Under Limited Training Data

Experiment III evaluates whether robustness is preserved when the amount of labelled source-domain data is reduced. The source domain is Chicago at 3.3 GHz and the OOD target domain is Brussels at 3.3 GHz. The Brussels test set remains fixed, while the Chicago source-domain training data are reduced to 100%, 50%, 25%, 10%, and 5% of the available training samples.




Figure 4.10: Experiment II: relative OOD degradation under joint city–frequency deployment shift. Results are reported over five random-seed runs.

As shown in Fig. 4.11, IINN maintains the lowest relative OOD degradation across the training fractions. Although all models are affected by reduced labelled data, the IINN curve remains consistently separated from the black-box baselines. The averaged result in Fig. 4.12 confirms this trend: IINN achieves the lowest average relative OOD degradation, $(17.70 \pm 3.38)\%$, compared with $(36.53 \pm 1.68)\%$ for the strongest baseline. The remaining baselines range from 42.40% to 52.57%.

This is the few-shot and limited-data result most relevant to deployment: IINN still needs data to calibrate trainable components, but the embedded propagation structure reduces the amount of data needed to maintain stable OOD behaviour.

4.3.3.4 Experiment IV: Propagation-Label Shift

Experiment IV evaluates propagation-label shift. In this setting, the deployment geometry remains unchanged but the label-generation mechanism changes from the 3GPP propagation model to the Standard Propagation Model (SPM). This is different from Experiments I–III: it tests whether the learned IINN structure is overly tied to one analytical label source. Because the source and target errors are measured against different label-generation mechanisms, absolute OOD RMSE is the appropriate metric.




Figure 4.11: Experiment III: relative OOD degradation across source-domain training fractions under cross-city deployment shift.




Figure 4.12: Experiment III: average relative OOD degradation across data-scarcity settings.

As shown in Fig. 4.13, IINN achieves the lowest absolute OOD RMSE among the compared models, although the margin relative to the strongest baseline is smaller than in the deployment-shift experiments. This behaviour is expected because the underlying geography remains fixed and only the label rule changes. The result supports the revised thesis claim: IINN provides structural traceability and competitive propagation-label-shift performance when the analytical labelling rule changes.




Figure 4.13: Experiment IV: absolute OOD RMSE under propagation-label shift from 3GPP labels to SPM labels.

Overall Winner Matrix

Experiment	Evaluation setting	Best RMSE	Best MAE	Best NRMSE	Best OOD	Overall
Experiment 1	Full-data city shift	DNN	DNN	✓ IINN/DNN/XGBoost	✓ IINN	✓ IINN
Experiment 2	City shift + data scarcity	✓ IINN	–	–	✓ IINN	✓ IINN
Experiment 3	Joint city + frequency shift	✓ IINN	✓ IINN	✓ IINN	✓ IINN	✓ IINN

Figure 4.14: Comparative performance of IINN across cross-city and cross-frequency evaluation experiments.

4.3.4 Overall Comparative Analysis

The three experiments show that standard black-box models can achieve strong in-distribution accuracy but are more sensitive to domain shift. In contrast, IINN may sacrifice marginal source-domain accuracy in some cases, but provides significantly stronger transferability across cities, frequencies, and limited-data settings. These findings support the claim that embedding physical structure into neural architectures improves robustness in realistic wireless planning scenarios.

Fig. 4.14 provides a consolidated comparison of the three experiments. IINN emerges as the strongest overall performer across cross-city, cross-frequency, and limited-data scenarios, achieving the best robustness in all experiments and the best metric balance under the most challenging distribution-shift settings.

4.3.5 Spatial Consistency Analysis

Fig. 4.15 compares ground-truth RSRP maps with the corresponding IINN predictions under city-level OOD, frequency-level OOD, and data-scarcity scenarios. In all cases, the predicted maps preserve the dominant spatial propagation structure, including coverage sectors, attenuation gradients, and high-power regions. This indicates that IINN learns physically meaningful geometric relationships governed by antenna orientation, transmitter–receiver distance, and environmental layout.




Figure 4.15: Spatial comparison between ground truth and IINN predictions under (a)–(b) city-level OOD, (c)–(d) frequency-level OOD, and (e)–(f) in-domain data scarcity conditions.

A consistent difference is observed in signal magnitude, where the predicted maps show slight underestimation and a smoother dynamic range than the ground truth. This effect is more visible in the city-shift and frequency-shift cases, suggesting that domain changes primarily affect calibration of absolute signal levels rather than spatial consistency. Under frequency OOD conditions, the model maintains similar coverage geometry with only modest amplitude distortion, demonstrating stable transfer across carrier frequencies. In the low-data scenario, IINN still produces smooth and physically meaningful coverage maps, although some fine local fluctuations are suppressed due to limited supervision.

Table 4.4: Summary of the evaluated OOD shift settings and their corresponding IINN and best-baseline results. Values are reported as mean $\pm$ standard deviation over five random-seed runs. Experiments I–III report relative OOD degradation; Experiment IV reports absolute OOD RMSE because the label-generation mechanism changes.

Exp.	Train city	Test city	Freq. shift	Train frac.	Train samples	Test samples	IINN result	Best baseline result
I	Chicago	Brussels	No	100%	14,969	17,437	$17.18 \pm 2.73\%$	ET: $37.50 \pm 1.16\%$
I [†]	Brussels	Chicago	No	100%	17,437	14,969	$37.54 \pm 1.78\%$	ET: $46.37 \pm 1.72\%$
II	Chicago	Brussels	Yes	100%	14,969	17,978	$22.13 \pm 2.91\%$	ET: $45.32 \pm 1.01\%$
III	Chicago	Brussels	No	100%	14,969	17,437	$15.14 \pm 3.46\%$	ET: $36.99 \pm 1.03\%$
III	Chicago	Brussels	No	50%	7,484	17,437	$17.35 \pm 2.89\%$	ET: $37.68 \pm 2.14\%$
III	Chicago	Brussels	No	25%	3,742	17,437	$15.99 \pm 3.88\%$	ET: $38.42 \pm 2.17\%$
III	Chicago	Brussels	No	10%	1,497	17,437	$20.73 \pm 3.78\%$	ET: $35.37 \pm 3.78\%$
III	Chicago	Brussels	No	5%	748	17,437	$19.29 \pm 6.69\%$	ET: $34.20 \pm 2.25\%$
IV	Brussels	Brussels	No	100%	12,205	21,901	10.52 ± 0.11 dB	HistGB: 10.57 ± 0.05 dB

[†]Opposite-direction cross-city transfer, Brussels to Chicago, is included as a direction check. For Experiment III, each row reports the corresponding training-fraction result. Experiment IV evaluates propagation-label shift from 3GPP labels to Standard Propagation Model labels; no paired five-seed Wilcoxon test is reported for this setting.

4.4 Structural Traceability and Training Stability

Table 4.4 consolidates the latest IINN-paper results across the evaluated robustness settings. Experiments I–III use relative OOD degradation as the robustness metric because the source and target labels follow the same propagation-label mechanism. Experiment IV uses absolute OOD RMSE because the target labels are generated using a different propagation model. This distinction keeps the interpretation consistent with the latest IINN paper: the first three settings test deployment-shift robustness, while the fourth setting tests propagation-label shift.

4.4.1 Training Stabilization and OOD Optimization

After constructing IINN from explicit propagation components, it is important to verify that the model remains a trainable neural model rather than a static analytical mapping. Therefore, the training and validation RMSE curves are inspected during optimization. Since the first epoch produces a large initial correction, the curve is shown after the first epoch to better visualize the convergence behaviour.

As shown in Fig. 4.16, IINN exhibits stable convergence. The training RMSE decreases rapidly during the early epochs and then gradually stabilizes, while the validation RMSE converges near 0.70 on the standardized RSRP scale. The small gap between training and validation curves and the absence of validation divergence indicate that the model is learning a stable propagation mapping rather than overfitting the training data.




Figure 4.16: Training and validation RMSE curves of IINN. The curves are shown on the standardized RSRP target scale after removing the first epoch for visual clarity. Both curves converge smoothly, with the best validation RMSE of approximately 0.697 at epoch 111, indicating stable data-driven learning of the trainable propagation-aligned components.

An initially flexible training run can become numerically unstable under shifted test conditions if feature ranges and residual corrections are unconstrained. The final training setup therefore uses clipped normalized features, bounded residual and antenna-related branches, standardized RSRP targets, a three-seed ensemble, and OOD feature weighting based on target-domain inputs only. These stabilization steps preserve the component-level traceability of IINN while making the learned mapping numerically reliable under deployment shift.

4.4.2 Forward Structural Traceability

Most explainability methods operate after a prediction has been produced. They estimate the relative importance of input variables but do not expose the sequence of internal computations that generated the prediction. For example, SHAP has been applied to wireless propagation modelling to rank influential variables such as propagation distance, antenna orientation, and antenna height [71]. Although such feature attribution is informative, it remains external to the predictive model. It cannot directly reveal which propagation branch was activated, how intermediate analytical terms were computed, or how individual physical components combined to produce the final received-power estimate.

Table 4.5: 3GPP equation coefficients and corresponding IINN learned coefficients after training with all available source-domain samples and with only 5% of the source-domain samples.

Term	3GPP coefficients	Learned coeff. 100% train data	Learned coeff. 5% train data
LoS path-loss branch			
$w_1 \log_{10}(d)$	22.00	22.06	21.44
$w_2 \log_{10}(f_c)$	20.00	21.66	19.85
$w_3 C_{\text{LoS}}$	32.40	14.08	27.95
NLoS path-loss branch			
$w_4 \log_{10}(f_c)$	20.00	21.08	20.94
$w_5 \log_{10}(W)$	7.10	5.04	6.09
$w_6 \log_{10}(h)$	14.90	16.96	15.91
$w_7 \log_{10}(h_{bs})$	18.76	15.98	17.56
$w_8 h_{ue}$	0.60	0.76	0.77
$w_9 \log_{10}(d) \log_{10}(h_{bs})$	3.10	3.36	2.92
$w_{10} \log_{10}(d)$	43.42	42.30	43.51
$w_{11} C_{\text{NLoS}}$	13.54	21.96	24.89
Antenna-gain branch			
$w_{12} \log_{10}(\zeta)$	10.00	13.21	10.61
$w_{13} \log_{10}\left(\frac{4\pi}{B_h B_v}\right)$	10.00	8.59	9.49
Directional antenna-attenuation branch			
$w_{14} \left(\frac{\phi_{ue} - \phi_{bs}}{B_v}\right)^2$	12.00	12.68	12.00
$w_{15} \min(\cdot, A_v)$	–	1.88	1.32
$w_{16} \left(\frac{\theta_{ue} - \theta_{bs}}{B_h}\right)^2$	12.00	12.00	12.00
$w_{17} \min(\cdot, A_h)$	–	1.88	1.32
$w_{18} A_{\text{att}}$	–	1.88	1.32
Output calibration			
$w_{19} R_x + b$	–	0.71, –1.53	0.72, –0.54

In contrast, IINN provides structural traceability through its architecture. Because every analytical operation is represented explicitly within the computational graph, the complete forward computation remains observable during inference. During inference, each intermediate node can export its input, output, and, where applicable, the learned analytical coefficient associated with the corresponding propagation term. Thus, IINN exposes both the intermediate analytical quantities and the learned coefficient corrections, rather than only the final received-power estimate.

Since every PLO is attached to a named analytical quantity rather than an abstract hidden feature, the learned analytical coefficients themselves remain physically interpretable. Table 4.5 compares the analytical coefficients defined by the 3GPP propagation model with the corresponding values learned by IINN under both full-data and limited-data training. Each learned parameter therefore represents either a correction to an analytical coefficient or a multiplier applied to an analytical constant while preserving the original analytical computation.

The learned coefficients illustrate how IINN combines analytical propagation knowledge with data-driven adaptation. Propagation terms that are already well represented by the analytical model require little correction; for example, the LoS distance coefficient changes only from 22.00 to 22.06 after training with all available data, indicating that the analytical distance dependence remains valid for the simulated environment. In contrast, environment-dependent terms are selectively recalibrated when supported by the data. For example, the NLoS street-width coefficient decreases from 7.10 to 5.04, indicating that the simulated environment exhibits a weaker street-width dependence than assumed by the analytical model, whereas the building-height coefficient increases from 14.90 to 16.96, suggesting a stronger influence of urban morphology on NLoS propagation. Even when trained using only 5% of the available training data, many coefficients remain close to their analytical initialization, demonstrating that the embedded propagation knowledge stabilizes coefficient learning under limited-data conditions. Together, these results show that IINN preserves established propagation knowledge while adapting only those analytical terms required by the observed data, providing parameter-level interpretability in addition to computational traceability.

Table 4.6: Section-wise coefficient trace for one Experiment I prediction. The table reports the 3GPP coefficients, the corresponding IINN learned coefficients, and the forward calculation for the selected test sample. Displayed values are rounded; section outputs are computed from the saved model values.

Term	3GPP coeff.	Learned coeff. 100% train data	Learned coeff. 5% train data
LoS path-loss branch			
$w_1 \log_{10}(d)$	22.00	22.06	21.44
$w_2 \log_{10}(f_c)$	20.00	21.66	19.85
$w_3 C_{\text{LoS}}$	32.40	14.08	27.95
$\widetilde{PL}_{\text{LoS}} = w_1 \log_{10}(d) + w_2 \log_{10}(f_c) + w_3 C_{\text{LoS}}$ $= 22.06 \log_{10}(130.47) + 21.66 \log_{10}(3300) + 0.43(32.40) = 136.97 \text{ dB.}$			
NLoS path-loss branch			
$w_4 \log_{10}(f_c)$	20.00	21.08	20.94
$w_5 \log_{10}(W)$	7.10	5.04	6.09
$w_6 \log_{10}(h)$	14.90	16.96	15.91
$w_7 \log_{10}(h_{bs})$	18.76	15.98	17.56
$w_8 h_{ue}$	0.60	0.76	0.77
$w_9 \log_{10}(d) \log_{10}(h_{bs})$	3.10	3.36	2.92
$w_{10} \log_{10}(d)$	43.42	42.30	43.51
$w_{11} C_{\text{NLoS}}$	13.54	21.96	24.89
$\widetilde{PL}_{\text{NLoS}} = w_4 \log_{10}(f_c) - w_5 \log_{10}(W) + w_6 \log_{10}(h) - w_7 \log_{10}(h_{bs})$ $- w_8 h_{ue} - w_9 \log_{10}(d) \log_{10}(h_{bs}) + w_{10} \log_{10}(d) + w_{11} C_{\text{NLoS}}$ $= 74.16 - 6.56 + 22.06 - 20.79 - 14.76 - 9.25 + 89.49 + 21.96 = 156.31 \text{ dB.}$			
Antenna-gain branch			
$w_{12} \log_{10}(\zeta)$	10.00	13.21	10.61
$w_{13} \log_{10}\left(\frac{4\pi}{B_{h,r} B_{v,r}}\right)$	10.00	8.59	9.49
$\widetilde{G} = w_{12} \log_{10}(\zeta) + w_{13} \log_{10}\left(\frac{4\pi}{B_{h,r} B_{v,r}}\right)$ $= 13.21 \log_{10}(0.70) + 8.59 \log_{10}\left(\frac{4\pi}{1.13 \times 0.11}\right) = 15.22 \text{ dB.}$			
Here $B_{h,r}$ and $B_{v,r}$ are beamwidths converted from degrees to radians for the gain calculation.			
Directional antenna-attenuation branch			
$w_{14} \left(\frac{\phi_{ue} - \phi_{bs}}{B_v}\right)^2$	12.00	12.68	12.00
$w_{15} \min(\cdot, A_v)$	–	1.88	1.32
$w_{16} \left(\frac{\theta_{ue} - \theta_{bs}}{B_h}\right)^2$	12.00	12.00	12.00
$w_{17} \min(\cdot, A_h)$	–	1.88	1.32
$w_{18} A_{\text{att}}$	–	1.88	1.32
$\widetilde{A}_{\text{att}} = w_{18} \left[w_{15} \min\left(w_{14} \left(\frac{\phi_{ue} - \phi_{bs}}{B_v}\right)^2, A_v\right) \right.$ $\left. + w_{17} \min\left(w_{16} \left(\frac{\Delta\theta}{B_h}\right)^2, A_h\right) \right],$ $= 1.88 [1.88 \min(4.04, 0.07) + 1.88 \min(66.84, 0.07)] = 0.49 \text{ dB.}$			
For this sample, $\Delta\theta$ is the wrapped azimuth difference.			
Output calibration			
$w_{19} R_{\text{raw}} + b$	–	0.71, –1.53	0.72, –0.54
$\widetilde{PL} = \max(136.97, 156.31) = 156.31 \text{ dB}$ $R_{\text{raw}} = P_t + \widetilde{G} - \widetilde{PL} - \widetilde{A}_{\text{att}} = 50.00 + 15.22 - 156.31 - 0.49 = -91.59 \text{ dB}$ $\widehat{R}_x = w_{19} R_{\text{raw}} + b = 0.714(-91.59) - 1.532 = -66.96 \text{ dB.}$			
Measured RSRP for this sample is –66.88 dB, giving an error of approximately –0.08 dB.			

4.5 Discussion

4.5.1 Key Findings and Insights

Training data scarcity, data imbalance, label noise, and distributional shifts between training and deployment scenarios are well-known limitations of conventional data-driven AI models. These limitations are particularly important in wireless communication systems, where the data-generating process may vary across cities, carrier frequencies, urban morphologies, clutter classes, and radio-planning tools. Motivated by these challenges, this chapter evaluated IINN as a domain-informed NN framework whose internal computation is aligned with physically meaningful propagation relationships.

The experimental results show that IINN provides a useful alternative to purely data-driven models when robustness and interpretability are required in addition to predictive accuracy. Across city-shift, frequency-shift, and limited-data settings, IINN generally achieves lower relative OOD degradation than conventional black-box baselines. This indicates that embedding propagation-aware structure into the architecture helps the model learn transferable physical relationships rather than relying only on city-specific statistical correlations. Although unconstrained models may sometimes achieve marginally lower RMSE under specific in-distribution or full-data settings, their higher OOD degradation suggests weaker stability under deployment shifts.

The propagation-label-shift experiment should be interpreted separately. In this setting, the label-generation mechanism changes from 3GPP to SPM, so absolute OOD RMSE is the more direct metric. IINN achieves the lowest absolute OOD RMSE among the compared models, indicating that its learned propagation structure remains competitive when the analytical labelling rule changes.

Although the experiments are presented using representative 1 km² urban analysis tiles, the cross-city results show that the proposed framework is not tied to a single local environment. Consistent behaviour across Brussels and Chicago suggests that the methodology can transfer across distinct urban morphologies when relevant domain variables are available.

4.5.2 Comparison with Existing Approaches

IINN is designed to optimize a broader objective than raw predictive accuracy alone. Its main advantages include transparent internal reasoning, compact architecture, and stronger behaviour under limited-data or shifted-environment conditions. The predictive performance of IINN is also influenced by the analytical model used to construct the architecture. Therefore, the observed performance should be interpreted with respect to the present model instantiation rather than treated as an unconditional claim about all possible IINN designs.

Beyond higher robustness in sparse-data and shifted-deployment conditions, IINN introduces methodological advantages that distinguish it from conventional DL models. Interpretability is achieved by decomposing the prediction into named and inspectable propagation components. Each component can be exported and visualized for every prediction, providing structural transparency without relying only on post-hoc explanation tools such as SHAP or LIME. This enables auditing, debugging, and deployment-time diagnosis because systematic errors can be traced to specific physical components rather than to an opaque hidden representation.

Compared with conventional black-box models, IINN restricts the hypothesis space using domain structure rather than relying solely on large-scale data fitting. This provides a useful inductive bias under limited-data and shifted-environment conditions. Compared with post-hoc explainability approaches, IINN provides interpretability by design: the internal computation itself is physically meaningful, rather than being explained externally after training.

4.5.3 Scope and Practical Considerations

IINN is beneficial when approximate analytical models are available and labelled data are limited or costly to obtain. In such settings, embedding domain structure into the network improves data efficiency, robustness, and interpretability. When very large and diverse datasets are available, larger unconstrained models may become competitive due to higher flexibility. Moreover, the performance of IINN depends on the quality of the chosen analytical scaffold and may be limited if that scaffold poorly matches the target environment. Therefore, IINN should be viewed as a targeted learning framework that is particularly advantageous in structured low-data settings rather than as a universal replacement for all neural architectures.

Although this chapter focused on wireless propagation, the broader IINN principle is applicable to other engineering and scientific domains where partial physical knowledge exists but labelled data are expensive. The key lesson is that interpretable NN systems should embed domain knowledge as reusable computational structure: components are constrained, named, and inspectable, while their precise coefficients and residual corrections remain adaptable through data.

4.6 Chapter Summary

This chapter has provided the experimental evidence for the IINN methodology introduced in Chapter 3. The simulation and data discussion established the wireless propagation setting, while the benchmark experiments showed that IINN reduces relative OOD degradation under cross-city, joint city-frequency, and limited-data settings. The propagation-label-shift experiment used absolute OOD RMSE and showed that IINN remains competitive when the label-generation mechanism changes. The structural-traceability analysis then demonstrated that the trained IINN exposes intermediate propagation terms, learned analytical coefficients, and section-wise prediction calculations during inference.

The next chapter uses the same IINN principle for a different technical contribution: domain-aware transfer learning. Instead of testing propagation robustness directly, it asks whether the physical meaning of IINN components can guide which layers should be frozen or fine-tuned when a model is adapted to a new deployment scenario.

Chapter 5

IINN-Based Transfer Learning for Wireless Propagation Modelling

Chapter 3 introduced the construction methodology of Innately Intelligent Neural Networks (IINNs), and Chapter 4 evaluated IINN for wireless propagation prediction. This chapter uses the same IINN design principle for a different purpose: transfer learning (TL). The focus is therefore not to introduce a new propagation model, but to test whether the component-level interpretability of IINN can guide adaptation from a source deployment to a target deployment using limited target-domain data [5], [13], [28].

This distinction also clarifies the role of Chapter 2. The Chapter 2 model is a convolutional classifier trained on received constellation images for spectrum awareness and modulation recognition. Its inputs, labels, task objective, and experimental setting are different from the IINN propagation-modelling chapters. Chapter 2 motivates the thesis transition from capable but opaque black-box learning to structurally traceable wireless AI; Chapters 3–5 then develop IINN construction, evaluate IINN for RSRP prediction, and use IINN for transfer learning [4], [11], [12].

5.1 Introduction

Conventional transfer learning usually adapts a pre-trained neural network by freezing a selected group of layers and fine-tuning the remaining layers on target-domain data. In generic DNNs, however, the choice of which layers to freeze is often heuristic. It commonly requires repeated experiments over different freezing depths, learning rates, and target-data fractions [28], [89]. This is inefficient in wireless deployments, where target measurements may be limited and where adaptation must be explainable to network engineers.

IINN changes the transfer-learning problem because its layers and branches correspond to meaningful physical or operational components. For propagation modelling, these components include distance-frequency path loss, antenna alignment, site-height terms, sector orientation, and bounded residual corrections. If a deployment shift is known to involve base-station height, azimuth, or downtilt, the corresponding IINN components can be selected for fine-tuning directly. This replaces trial-and-error layer search with a domain-aware adaptation rule [8], [12], [13].

The central question addressed in this chapter is:

Can the structural traceability of IINN be used to decide which parts of a wireless propagation model should be adapted during transfer learning?

5.1.1 Related Work

Existing TL methods improve adaptation by reusing a source-domain model, but most methods do not connect layer selection to domain semantics. Standard fine-tuning freezes early layers and updates later layers, assuming that generic features appear early and task-specific features appear late. This assumption is not always appropriate for wireless propagation, where the variable factor may be a specific physical term rather than a particular depth of the network [13], [28], [29].

More recent TL approaches include adaptive routing, adapter modules, parameter-efficient tuning, co-tuning, and distillation-based transfer [29], [30], [90], [91]. These methods can reduce the number of trainable parameters, but they still typically decide what to adapt from data-driven or architectural heuristics rather than from explicit wireless-domain structure. Physics-informed neural networks incorporate constraints through loss functions [10], but their internal layers are not necessarily aligned with inspectable propagation components.

In contrast, this chapter uses the IINN architecture introduced in Chapter 3 as an interpretable transfer backbone. The contribution is the fine-tuning strategy: the source model is adapted by updating only those branches that correspond to deployment variables expected to change [8], [12], [13].




Figure 5.1: IINN-based transfer-learning framework. The source IINN is constructed from domain-informed propagation components, pre-trained on source-domain data, and selectively fine-tuned on target-domain data by updating only the branches associated with changed deployment variables.

5.1.2 Contributions

This chapter makes three contributions.

- It formulates a domain-aware TL procedure for IINN in which source-domain knowledge is retained in invariant branches and target-domain adaptation is restricted to physically variable branches.
- It defines the layer-selection logic for base-station height, sector, azimuth, and downtilt changes, showing how IINN structural traceability supports principled fine-tuning.
- It evaluates the method on a simulated LTE propagation dataset and shows improved transfer gains relative to conventional DNN fine-tuning, especially when target base-station height or sector configuration differs from the source.

The rest of the chapter is organized as follows. Section 5.2 presents the transfer-learning framework, Section 5.3 describes the simulation setup and performance metrics, Section 5.4 discusses the results, and Section 5.5 concludes the chapter.

5.2 IINN-Based Transfer Learning Framework

5.2.1 Architecture-to-Parameter Mapping

The IINN construction process follows the Chapter 3 principle: variables, coefficients, and mathematical operations from an analytical propagation model are mapped into explicit neural components. Let the input variables be

$$x_1, x_2, \dots, x_n \in \mathbb{R}, \quad (5.1)$$

where each x_i represents a distinct wireless variable such as distance, carrier frequency, base-station height, antenna tilt, azimuth, or user height. Analytical coefficients are denoted by

$$\psi = \{\psi_1, \psi_2, \dots, \psi_m\}. \quad (5.2)$$

When reliable coefficient values are available, they can be used to initialize the corresponding trainable parameters:

$$w_i, b_i \leftarrow \psi_i. \quad (5.3)$$

Each operation $\mathcal{O}_i$ is implemented as a dedicated computational block. For example, an expression of the form

$$y = \min \left(\left(\frac{x_1 - x_2}{x_3} \right)^2, x_4 \right) \quad (5.4)$$

is decomposed into subtraction, division, square, and minimum operations, each of which remains visible in the computational graph. This explicit mapping is what allows transfer learning to be guided by physical meaning rather than by network depth alone.

For the propagation case study, the received-power expression used to define the source IINN includes distance, frequency, antenna-gain, base-station height, street-width, building-height, azimuth, downtilt, and beamwidth terms:




Figure 5.2: Identification of variables and mathematical operators from a propagation equation. Variables define IINN input features, while operators define the structure of the corresponding network components.




Figure 5.3: Segment of an IINN architecture constructed by mapping analytical propagation-model terms into neural components. The same component alignment is later used to decide which branches should be fine-tuned.

$$\begin{aligned}
& P_r(d, f_c, W, h, h_{ue}, h_{bs}, B_h, B_v, \phi_{ue}, \phi_{bs}, \theta_{ue}, \theta_{bs}, \lambda_h, \lambda_v, A_h, A_v) = \\
& P_t - \max[(22 \log_{10}(d) + 20 \log_{10}(f_c) + 28), (20 \log_{10}(f_c) - 7.1 \log_{10}(W) + 14.9 \log_{10}(h) \\
& \quad - 18.76 \log_{10}(h_{bs}) - 0.6 h_{ue} - 3.1 \log_{10}(d) \log_{10}(h_{bs}) + 43.42 \log_{10}(d) + 31.69)] \\
& + 10 \log_{10}(\zeta) + 10 \log_{10} \left(\frac{4\pi}{B_h B_v} \right) + \lambda_v \min \left[12 \left(\frac{\phi_{ue} - \phi_{bs}}{B_v} \right)^2, A_v \right] + \lambda_h \min \left[12 \left(\frac{\theta_{ue} - \theta_{bs}}{B_h} \right)^2, A_h \right].
\end{aligned} \tag{5.5}$$

Table 5.1: Domain-aware layer selection for IINN transfer learning.

Target-domain change	IINN components updated	Components retained
Base-station height change	Height, site-geometry, and height-dependent path-loss branches	Frequency, transmit-power, and unchanged antenna terms
Sector or azimuth change	Horizontal antenna-alignment and sector-specific offset branches	Distance-frequency path-loss and height branches when unchanged
Downtilt change	Vertical antenna-alignment and downtilt-related attenuation branches	Transmit-power, carrier-frequency, and unchanged site terms
Small calibration shift	Output intercept or bounded residual calibration branch	Main propagation and antenna branches

5.2.2 Domain-Aware Fine-Tuning Strategy

Let the pre-trained IINN be represented as

$$\hat{y}_T = F(\mathbf{x}; \varphi_i, \varphi_v), \tag{5.6}$$

where φ_i denotes invariant parameters retained from the source domain, and φ_v denotes variable parameters associated with domain factors expected to change in the target domain. Fine-tuning solves

$$\varphi_v^T = \arg \min_{\varphi_v} \mathcal{L}_T(F(\mathbf{x}; \varphi_i, \varphi_v), y_T), \tag{5.7}$$

where the target-domain loss is the RMSE over the available target samples:

$$\mathcal{L}_T = \sqrt{\frac{1}{N_T} \sum_{j=1}^{N_T} (\hat{y}_{T,j} - y_{T,j})^2}. \tag{5.8}$$

The key distinction from conventional DNN fine-tuning is the selection rule. A DNN typically requires a search over freezing configurations. In IINN, the trainable subset is chosen from the known deployment change:

This targeted procedure has a direct computational implication. If the full model has p trainable parameters and only p_v parameters are associated with the changed target-domain components, IINN fine-tuning scales with the smaller adaptation set. In contrast, exhaustive DNN freezing searches require training multiple candidate configurations. The practical benefit is not only fewer updated parameters, but also fewer experimental trials and a lower risk of overfitting small target datasets.

5.3 Simulation Setup and Performance Evaluation

The transfer-learning evaluation uses a dataset generated in Atoll, a commercial network-planning tool, for a 4G LTE macro-cell scenario in Brussels, Belgium. The simulated area covers 3.8 km^2 and contains 10 macro-cell eNodeBs. Users are distributed according to a Poisson process, and RSRP is aggregated over spatial bins. The simulation parameters are summarized in Table 5.2.

The baseline is a conventional DNN implemented as a Keras sequential model with four fully connected layers of 32, 64, 64, and 16 neurons, ReLU activation functions, and the Adam optimiser. Both the DNN and IINN models are pre-trained on source-domain data and fine-tuned using target-domain data. The data split is 80% training and 20% testing in each experiment.

Performance is reported using transfer gain. Let E_{ref} denote the target-domain prediction error before transfer or under the reference adaptation setting, and let E_m denote the target-domain error after applying model m . The percentage transfer gain is

$$G_m = 100 \times \frac{E_{\text{ref}} - E_m}{E_{\text{ref}}}. \quad (5.9)$$

Table 5.2: Network scenario setting for the transfer-learning study.

Category	Parameters and Values
System	Air technology: 4G LTE; 10 macro-cell sites; 3 sectors per eNodeB; Carrier frequency: 2100 MHz; channel bandwidth: 5 MHz
Geographic	Propagation model: Aster ray tracing; location: Brussels, Belgium; Area: 3.8 km^2 ; UE distribution: Poisson; Propagation-matrix resolution: 10 m; average building height: 20 m; average street width: 20 m
eNodeB	Antenna model: Kathrein directional model 742265; Height: 28–37 m; tilt: 0° – 6° ; azimuth: 360° ; Gain: 18 dBi; beamwidth (H/V): $65^\circ/9^\circ$; transmit power: 43 dBm

Table 5.3: Summary of transfer-learning experimental setup.

Parameter	Details
Base-station heights	29 m, 32 m, 35 m
Sectors per base station	Sector A (0°–120°), Sector B (120°–240°), Sector C (240°–360°)
Source-domain data	Train: 2600; test: 650
Target-domain data	Train: 1200; test: 300

Positive gain therefore indicates reduced target-domain error. The difference between the proposed IINN-based approach (PA) and the baseline DNN approach (BL) is

$$\Delta G = G_{\text{PA}} - G_{\text{BL}}. \quad (5.10)$$

Positive ΔG means that IINN-based selective fine-tuning improves transfer more than conventional DNN fine-tuning.

5.4 Transfer-Learning Results

5.4.1 Base-Station Height Transfer

Fig. 5.4 compares transfer gains across base-station height changes. When source and target heights are the same, the benefit of transfer is limited, as expected. When the target height differs from the source height, IINN-based fine-tuning produces larger gains than the DNN baseline, with improvements of up to approximately 10.3% in the evaluated configurations. This result is consistent with the layer-selection logic in Table 5.1: height-related changes are handled by adapting the height and site-geometry components rather than by retraining unrelated branches.

5.4.2 Sector-Based Transfer

Fig. 5.5 shows transfer gains across sectors A, B, and C. Sector changes affect azimuth and antenna-orientation relationships, so the relevant IINN components are the horizontal antenna-alignment and sector-offset branches. The IINN approach reaches improvements of up to 16.4% in sector-based transfer, indicating that the model benefits most when the target shift corresponds to a clearly identifiable physical component.




Figure 5.4: Transfer-gain comparison for source and target base-station heights. Off-diagonal entries show adaptation when the target base-station height differs from the source height.

5.4.3 Combined Height and Sector Configurations

Fig. 5.6 summarizes the combined source-target height and sector experiments. The strongest positive differences occur when the source and target domains differ in the variables that IINN exposes as separate components. Small or negative differences can occur when the source and target settings are already similar, because the baseline DNN has less adaptation burden and the additional selectivity of IINN provides less marginal benefit.

The results therefore support a conditional interpretation rather than a universal claim. IINN-based TL is most useful when the target-domain change is physically meaningful, visible in the model structure, and represented by available input variables. It is less likely to help when the source and target domains are nearly identical, when the relevant target-domain factor is missing from the feature set, or when the label-generation mechanism changes in a way that is not represented by the selected IINN components.

5.4.4 When the Pre-Trained IINN Should and Should Not Transfer

The pre-trained IINN is expected to transfer well when source and target domains share the same prediction task, the same feature semantics, and broadly similar propagation mechanisms, while differing in identifiable parameters such as base-station height, sector orientation, antenna tilt, transmit power, or frequency. In these cases, the source model retains useful propagation knowledge and only the affected branches need adaptation.




Figure 5.5: Sector-based transfer-gain comparison averaged across base-station heights. Positive PA–BL values indicate that IINN-based selective fine-tuning outperforms conventional DNN fine-tuning.




Figure 5.6: Transfer-gain difference (PA–BL) for combined source-target base-station height and sector configurations.

Transfer is less reliable when the target environment changes in ways not represented in the model. Examples include missing clutter or building metadata, a different label-generation tool with incompatible assumptions, a change from sub-6 GHz macro-cell propagation to a substantially different propagation regime, or a target task whose output no longer represents the same received-power quantity. Under large target datasets, flexible black-box models may also become competitive because they can learn the target distribution directly.

5.5 Conclusion

This chapter extended IINN from interpretable propagation prediction to domain-aware transfer learning. Because IINN components correspond to meaningful wireless variables and analytical operations, the model provides a principled basis for deciding which parameters should be retained and which should be adapted when moving from source to target deployments. The experiments show that selective IINN fine-tuning improves transfer over a conventional DNN baseline, with gains of up to 16.4% in sector-based transfer and approximately 10.3% in base-station height transfer.

5.6 Chapter Summary

This chapter has shown that IINN interpretability is not only explanatory, but operational. The same structural traceability that makes a prediction auditable can guide targeted model adaptation. The next chapter synthesizes the findings from the preceding technical chapters, connecting conventional deep learning, IINN-based structure-aware modelling, and IINN-guided transfer learning into a unified thesis argument.

Chapter 6

Synthesis, Scope, Limitations, and Future Research Directions

This final chapter synthesizes the six-chapter thesis argument and consolidates the scope, limitations, and future research directions. It replaces a separate synthesis chapter and conclusion chapter so that the revised thesis contains exactly six chapters: background and literature review; four technical contribution chapters; and this final synthesis and future-work chapter [11], [12], [13], [16].

The preceding chapters present four technical contribution chapters: a deep learning spectrum-awareness framework for vehicular communication, the IINN construction methodology, the IINN propagation experiments and ablation study, and an IINN-based transfer learning strategy for domain-aware adaptation. Read separately, these chapters could appear as related wireless AI studies. Read together, they form a single argument: wireless learning systems become more defensible when the model architecture reflects the structure of the domain rather than relying entirely on statistical correlations learned from data [4], [5], [7], [12].

This chapter synthesizes the thesis as a unified research programme. It does not repeat the experimental details of Chapters 2–5. Instead, it explains how the chapters fit together, what the combined evidence supports, what it does not support, and how the proposed IINN framework should be positioned relative to black-box learning, post-hoc explainability, physics-informed learning, and transfer learning [7], [10], [13], [28].

6.1 The Research Arc

The thesis progresses through three stages.

Stage 1: Correlation-driven capability. Chapter 2 demonstrates that deep learning can provide useful wireless intelligence. The CNN-based spectrum-awareness framework classifies modulation orders under varying vehicular channel conditions and supports a CQI-oriented modulation adaptation strategy. This is an important practical result because it shows that data-driven learning can extract useful structure from wireless signals. However, the model remains a black-box classifier: the internal features responsible for the decision are not aligned with identifiable channel mechanisms [4], [7], [11].

Stage 2: Architecture-level interpretability. Chapter 3 addresses the limitation exposed by Chapter 2 by moving the intervention point from output explanation to model construction. The IINN framework maps analytical domain relationships into neural structure before training. In the wireless propagation case study, received-power prediction is decomposed into physically meaningful components such as path loss, antenna alignment, site/environment contribution, clutter offset, and bounded residual correction. This provides interpretability by construction rather than by post-hoc approximation [8], [9], [12].

Stage 3: Interpretability as an adaptation mechanism. Chapter 5 shows that structural interpretability is not only valuable for human inspection. It can also guide transfer learning. Because IINN layers correspond to physical terms, the model can distinguish between invariant and variable components when moving from one deployment condition to another. This enables targeted fine-tuning without exhaustive layer-search heuristics [13], [28].

The resulting thesis arc is therefore:

black-box wireless AI works → black-box wireless AI is not enough → domain structure can be embedded into the architecture → architectural structure enables more principled adaptation.

6.2 Levels of Intervention in Wireless AI

A useful way to position the thesis is to ask where domain knowledge enters the learning pipeline. Existing approaches generally intervene at the data level, output level, loss level, or training-procedure level. IINN intervenes at the architecture level.

Table 6.1: Different intervention levels for introducing reliability, interpretability, or domain knowledge into wireless learning systems.

Intervention level	Representative idea	Main benefit	Remaining limitation
Data level	Simulation, digital twins, augmentation, synthetic data generation	Reduces measurement burden and expands training coverage	The model may still learn opaque correlations from generated labels.
Output level	SHAP, LIME, confidence scores, credibility estimation	Improves explanation or reliability assessment after training	Does not change how the model internally reasons.
Loss level	PINNs and physics-regularized objectives	Encourages physical consistency during optimization	Internal hidden representations may remain generic and difficult to inspect.
Training-procedure level	Transfer learning, fine-tuning, domain adaptation	Reuses source-domain knowledge for target domains	Layer selection is often heuristic in black-box architectures.
Architecture level	IINN equation-to-architecture mapping	Makes computation structurally aligned with physical terms before training	Requires an approximate analytical or empirical model.

This distinction is central to the thesis. The claim is not that architecture-level intervention replaces every other method. Rather, it addresses a different and earlier design question: what should the model know before it sees data? In physics-governed wireless tasks, the answer should not always be “nothing”.

6.3 What the Evidence Establishes

The experimental evidence supports four bounded claims.

6.3.1 Deep Learning Is Useful But Structurally Opaque

Chapter 2 validates the usefulness of deep learning in a wireless classification task. The proposed spectrum-awareness pipeline uses a CNN classifier to support modulation-order decisions under changing channel variance, speed, SNR, and rural/urban conditions. This contribution establishes that data-driven learning is a viable component of wireless intelligence [4], [11].

However, the chapter also clarifies the limitation that motivates IINN. The CNN can classify effectively, but its internal computation cannot be decomposed into physically meaningful channel terms. For engineering use cases requiring debugging, auditing, or causal diagnosis, this is a fundamental limitation. Thus, Chapter 2 is not merely an isolated application of deep learning; it is the diagnostic starting point of the thesis [7], [8], [12].

6.3.2 Structure-Aware Architecture Improves Interpretability and Robustness

Chapter 3 demonstrates that domain knowledge can be translated into a neural architecture. The IINN framework does this by decomposing analytical equations into scalar inputs, parametric linear operations, fixed mathematical operators, and physically named components. This design makes intermediate computations inspectable [8], [9], [12].

The empirical results show that IINN is particularly useful under conditions where black-box DNNs are vulnerable: limited data and shifted environments. In cross-city and cross-frequency experiments, IINN shows lower OOD degradation than representative black-box baselines. The structural-traceability analysis further shows that IINN exposes intermediate propagation terms and learned analytical coefficients rather than only a final RSRP prediction [5], [12], [21].

This distinction matters. The propagation-label-shift experiment is interpreted using absolute OOD RMSE because the label-generation mechanism changes from 3GPP to SPM. Under this setting, IINN remains competitive with the best baseline while preserving an inspectable propagation pathway, which is the central value of the IINN contribution [5], [12], [22].

6.3.3 Interpretability Enables Principled Transfer

Chapter 5 demonstrates a downstream benefit of architectural interpretability. In conventional DNN transfer learning, the decision about which layers to freeze and which to fine-tune is often empirical. In IINN, the decision can be made from domain knowledge because each layer has a physical meaning [13], [28].

The results show gains of up to 16.4% in sector-based transfer and approximately 10.3% in base-station-height adaptation, with average improvements in the 10–15% range over a DNN transfer baseline. These results support a stronger claim than “IINN is explainable.” They show that interpretability can be operationally useful: it reduces the search space for adaptation and makes fine-tuning more controlled [13].

6.3.4 The Thesis Supports Conditional Superiority, Not Universal Superiority

The evidence does not imply that IINN is universally better than all ML models. In high-data, in-distribution settings, flexible ensemble methods such as LightGBM and XGBoost can achieve lower raw error than IINN. This is not a contradiction; it reflects a difference in objective. IINN optimizes for a broader engineering target: data efficiency, interpretability, bounded physical reasoning, and adaptation under shift.

6.4 Distinction Matrix

Table 6.2 summarizes this distinction.

Table 6.2: Defensibility of the main thesis claims.

Claim	Supporting evidence	Boundary condition
Deep learning can support wireless decision-making.	Chapter 2 CNN spectrum-awareness framework under varying vehicular channel conditions.	Does not provide intrinsic interpretability.
IINN provides structural interpretability.	Chapter 3 equation-to-architecture mapping and component-level structurally traceable IINN.	Requires physically meaningful input variables and an approximate analytical model.
IINN improves robustness under shift.	Chapter 4 city, frequency, data-scarcity, and cross-propagation ablation experiments.	Strongest evidence is simulation-based; live measurement validation remains needed.
Propagation-label shift requires separate evaluation.	Chapter 4 uses absolute OOD RMSE for 3GPP-to-SPM label shift and reports IINN as competitive with the strongest baseline.	Results are simulation-based and should be validated against measured labels.
Architectural interpretability improves transfer learning.	Chapter 5 domain-aware fine-tuning improves over DNN transfer baselines.	Evaluated on selected source-target configurations; broader cross-city and cross-band transfer remains future work.

6.5 Accuracy, Interpretability, and the Role of LightGBM

A critical thesis must address the fact that some black-box or weakly interpretable baselines can achieve very strong raw accuracy. Gradient-boosting models such as LightGBM often perform well on tabular propagation datasets because they capture nonlinear feature interactions efficiently and require less feature scaling than DNNs. In some settings considered in this thesis, such models achieve lower absolute RMSE than IINN.

This does not invalidate the IINN contribution, but it does define its correct use case. If the only objective is minimizing RMSE on a large representative dataset, a high-capacity ensemble model may be preferable. If the objective includes physical traceability, low-data learning, controlled adaptation, and diagnosis of why a prediction changes, then a structurally traceable architecture has value that a black-box ensemble does not provide.

This is an important shift in evaluation philosophy. For operational wireless AI, a model should not be judged only by whether it wins a leaderboard on one metric. It should also be judged by whether its predictions can be audited, whether it can adapt safely, and whether its errors can be related to physical causes.

6.6 Propagation-Label Shift and Structural Traceability

The strongest critical insight from the propagation experiments is that deployment shift and propagation-label shift should not be evaluated using the same metric. When the source and target labels follow the same propagation-label mechanism, relative OOD degradation captures how much the error grows under deployment shift. When the label-generation mechanism changes, absolute OOD RMSE is more appropriate.

Chapter 4 therefore reports the 3GPP-to-SPM setting separately from the city and frequency deployment-shift experiments. This framing prevents the robustness claim from being overstated: IINN shows its clearest advantage under city, frequency, and limited-data shifts, while the propagation-label-shift result shows competitive absolute OOD RMSE under a changed analytical labelling rule.

The structural-traceability evidence is the defensible interpretability claim. IINN exposes intermediate propagation terms, learned analytical coefficients, and section-wise calculation pathways, allowing the prediction process to be inspected rather than inferred only through post-hoc feature attribution.

6.7 Implications for AI-Native Wireless Networks

The thesis has implications for AI-native network design. Future RAN automation will require models that can operate under limited measurement data, transfer across deployment scenarios, and provide explanations meaningful to engineers. The IINN framework contributes to this direction in three ways.

First, it reduces the data burden by starting from physically meaningful structure. Second, it exposes internal components that can be inspected during deployment. Third, it enables adaptation strategies based on known changes in network configuration rather than blind fine-tuning.

The most promising practical deployment model may not be a pure IINN-only system. A future network could combine a high-accuracy black-box predictor with an IINN-based structurally traceable diagnostic layer, or use IINN to initialize, regularize, calibrate, or audit more flexible models. This hybrid direction is consistent with the thesis: the goal is not to reject data-driven learning, but to prevent learning systems from ignoring structure that engineers already know.

6.8 Synthesis Summary

The synthesis above positions the technical findings within a broader argument for structure-aware wireless AI. The evidence from the preceding chapters shows that conventional deep learning is useful but opaque, that domain-embedded architecture can improve interpretability and robustness, and that interpretable components can support more principled transfer learning. The remainder of this final chapter consolidates the thesis answer, contributions, limitations, scope, and future research directions.

6.9 Overall Conclusion and Future Work

This thesis investigated whether wireless learning systems can move from correlation-driven modelling toward structure-aware neural design. The motivating problem was not simply that conventional deep learning models sometimes fail under limited data or distribution shift. The deeper problem was that their internal computations are usually not aligned with the physical mechanisms that govern wireless systems. This makes them difficult to audit, difficult to adapt, and difficult to trust in operational settings where engineers need to understand why a prediction changes.

The thesis addressed this problem through four technical stages. It first demonstrated the practical value and limits of conventional deep learning in spectrum awareness for vehicular communication. It then introduced the IINN construction methodology, evaluated IINN through propagation-modelling experiments and ablation studies, and finally showed that IINN interpretability can guide transfer learning in a principled way.

6.10 Answer to the Central Research Question

The central research question of the thesis was:

Can wireless learning systems be designed so that they are not only accurate, but also interpretable, data efficient, robust under deployment shift, and adaptable with limited target-domain data?

The answer supported by this thesis is *yes, under clearly defined conditions*. If the domain provides an approximate analytical or empirical model, then parts of that knowledge can be embedded into the neural architecture before training. This reduces the burden on data, constrains the hypothesis space to physically meaningful relationships, and exposes intermediate computations that can be inspected. However, this approach is not universally superior. It is most appropriate when interpretability, data efficiency, and robustness under shift matter alongside predictive accuracy.

6.11 Consolidated Thesis Contributions

The thesis makes four consolidated contributions [11], [12], [13], [16].

- C1: Deep learning for wireless spectrum awareness and its structural limitation.** Chapter 2 developed a lightweight deep learning-based modulation classification and CQI-oriented spectrum-awareness framework for vehicular communication. This contribution showed that correlation-driven learning can provide useful wireless decision-making capability, while also motivating the need for models whose internal computations are physically traceable [4], [11].
- C2: IINN framework for architecture-level domain embedding.** Chapter 3 introduced a methodology for converting domain-informed analytical relationships into neural computational structure. Instead of selecting neural architecture heuristically, IINN decomposes physical equations into scalar inputs, trainable parametric operations, fixed mathematical operators, and physically named model components [8], [9], [12].
- C3: Robust and structurally traceable propagation modelling under deployment shift.** Chapter 4 showed that IINN improves robustness under cross-city, joint city–frequency, and limited-data settings, while remaining competitive under propagation-label shift when evaluated using absolute OOD RMSE. The chapter also demonstrated structural traceability through exported propagation terms, learned analytical coefficients, and section-wise prediction traces [5], [12], [21].
- C4: Domain-aware transfer learning enabled by interpretability.** Chapter 5 demonstrated that IINN interpretability can be used operationally. By identifying which layers correspond to variable physical parameters, the model can fine-tune only the components expected to change between source and target domains, reducing trial-and-error adaptation [13], [28].

6.12 Main Technical Findings

The thesis leads to the following technical findings.

6.12.1 Architecture Is a Form of Prior Knowledge

In conventional DNNs, prior knowledge is usually introduced through features, data augmentation, regularization, or post-hoc explanation. IINN shows that architecture itself can encode prior knowledge. When the network topology follows the structure of an analytical model, the learning problem changes: the model no longer needs to discover all relationships from data, but instead learns how to calibrate and refine physically meaningful components.

6.12.2 Interpretability and Robustness Can Have a Shared Cause

The thesis shows that interpretability and robustness need not be separate objectives. In IINN, both arise from the same design choice: structural alignment with domain relationships. The model is interpretable because its components correspond to physical terms, and it is more stable under certain deployment shifts because those physical relationships remain meaningful beyond a single training distribution.

6.12.3 Structural Traceability Requires Inspectable Computation

One of the most important lessons is that interpretability should mean more than exposing input-feature importance. In IINN, structural traceability means that the computation itself can be inspected through named propagation branches, learned analytical coefficients, and section-wise prediction traces. This makes the model auditable during inference rather than only explainable after prediction [9], [12], [22].

6.12.4 Interpretability Can Improve Learning Efficiency

Chapter 5 shows that interpretability is not merely a property that helps humans understand a trained model. It can also improve learning efficiency. When layers have physical meaning, fine-tuning can be guided by knowledge of which physical parameters change between source and target domains. This transforms transfer learning from an empirical search problem into a domain-informed adaptation procedure [13], [28].

6.13 Limitations

The conclusions of this thesis are bounded by the following limitations.

- L1: **Simulation-based evaluation.** The propagation modelling experiments are primarily based on simulation or planning-tool generated data. This enables controlled testing but does not fully resolve performance under live network measurements.
- L2: **Single primary physical domain.** The IINN framework is demonstrated mainly for wireless propagation modelling. The broader claim that equation-native architectures can be used in other domains remains a plausible extension rather than a completed validation.
- L3: **Dependence on analytical model quality.** IINN requires at least an approximate physical or empirical model. If the starting model is strongly misspecified, the structural prior can become a constraint rather than an advantage.
- L4: **Raw accuracy trade-off.** IINN is not always the lowest-RMSE model in high-data regimes. Gradient-boosting models and other flexible black-box methods may outperform it when the only objective is prediction accuracy on representative data.
- L5: **Limited transfer learning scope.** The transfer learning study demonstrates the principle of domain-aware fine-tuning, but broader evaluation across cities, frequency bands, radio technologies, and deployment types remains necessary.
- L6: **Operational validation.** The thesis argues that structurally traceable components can support fault diagnosis, auditing, and engineering decisions, but these operational workflows are not yet tested with network operators in live decision-making loops.

6.14 Future Research Directions

The thesis opens several concrete research directions.

6.14.1 Real-Measurement Validation

The highest-priority future direction is validation on live-network measurements. A measurement campaign should include multiple sites, clutter classes, antenna configurations, and frequency settings. This would directly test whether the revised IINN retains its robustness when the label-generating process is decoupled from simulation. It would also quantify the remaining gap between simulation-trained and measurement-calibrated performance [3], [5], [6].

6.14.2 Systematic Model-Misspecification Analysis

Future work should intentionally perturb the analytical model used to construct the IINN architecture. This would reveal how sensitive the framework is to incorrect path-loss coefficients, missing environmental terms, wrong antenna assumptions, or incomplete clutter modelling. Such an analysis would make the framework more defensible by identifying when structural priors help and when they hurt [9], [12], [22].

6.14.3 Cross-Domain IINN Validation

The IINN design principle should be tested beyond urban macrocell propagation. Candidate domains include indoor propagation, millimetre-wave and sub-terahertz channel modelling, OFDM channel estimation, localization, mobility prediction, energy systems, biomedical signals, and structural health monitoring. Each domain would require its own equation decomposition and validation protocol [3], [6], [8].

6.14.4 Automated Equation-to-Architecture Generation

At present, IINN construction requires manual decomposition of analytical equations. Future work should develop tools that parse symbolic equations and generate neural architecture specifications automatically. Such tooling would reduce implementation errors, improve reproducibility, and make the framework easier to apply across domains [9], [52], [53].

6.14.5 Hybrid Traceable and Black-Box Systems

IINN should not be viewed only as a competitor to black-box models. A promising direction is hybrid modelling: use flexible black-box models for high-accuracy prediction and IINN components for physical auditing, calibration, uncertainty detection, or explanation. Another direction is to use IINN as a structured backbone with a bounded residual branch that captures effects not represented in the analytical model [12], [24], [58], [59].

6.14.6 IINN for AI-Native RAN Digital Twins

Digital twins require models that are fast, interpretable, and adaptable. IINN is well suited to this setting because its components can represent network-controllable variables such as antenna tilt, azimuth, height, frequency, and transmit power. Future work should integrate IINN into radio-planning digital twins and evaluate whether structurally traceable components improve network optimization workflows [2], [6], [16].

6.14.7 IINN-LLM Explanation Interfaces

Large language models (LLMs) can improve how model behaviour is communicated to engineers, but they do not by themselves make a model physically grounded. A useful future direction is to combine IINN and LLMs: IINN provides physically meaningful intermediate outputs, while an LLM translates those outputs into operator-facing explanations. This would combine structural interpretability with accessible communication [7], [12], [60].

6.14.8 Transfer Learning Across Deployment Shifts

The transfer learning framework should be extended to cross-city, cross-frequency, cross-technology, and cross-propagation-model scenarios. Future work should also quantify parameter counts, wall-clock fine-tuning cost, and negative-transfer cases, identifying when source and target domains differ so strongly that selective fine-tuning is insufficient [13], [21], [28].

6.15 Chapter Summary

The main conclusion of this thesis is that wireless learning systems do not have to begin as blank slates. When reliable domain structure exists, it can be embedded into the model architecture before training. This does not remove the need for data, nor does it guarantee universal superiority over flexible black-box models. It changes the role of data: instead of forcing data to discover every relationship from scratch, data refine a model that already contains a physically meaningful scaffold [8], [9], [12].

The thesis therefore contributes more than a particular neural architecture. It contributes an engineering principle: in physics-governed domains, model design should ask what the system already knows before learning begins. IINN is one implementation of that principle. Its results show that architecture-level domain embedding can improve interpretability, data efficiency, robustness under shift, and transfer learning. Its limitations show that this principle must be applied carefully, tested against independent measurements, and extended beyond a single domain [5], [12], [13], [22].

Within these boundaries, the thesis supports a clear position: for wireless systems moving toward AI-native operation, trustworthy learning will require more than larger datasets and larger models. It will require models whose structure reflects the systems they are intended to control [2], [7], [12].

Appendix A

Supplementary Material

A.1 Chapter 2 Supplementary Dissemination Material

The poster shown in Fig. A.1 is retained as supplementary evidence of dissemination for the spectrum-awareness work in Chapter 2. It is moved to the appendix so that the main chapter remains focused on the technical methodology and evaluation.

Abstract

- This paper proposes a deep learning-based modulation classification technique to optimise the modulation order according to different channel parameter settings, including, speed, variance, and signal-to-noise ratio in urban and rural areas for a 3D stochastic vehicular channel.
- We introduce a channel quality indicator block in which the classifier's accuracy of detection is used as a forward indicator for the choice of the best modulation type for transmission.

Contributions

- For effective observation of unoccupied channels, we apply the spectrum aggregation-based ultra-wideband spectrum monitoring method.
- By using deep learning, this work develops a lightweight feature-based modulation classification approach that can detect modulation order with high probability of detection for an acceptable Pe.
- Based on a 3-dimension (3D) stochastic vehicular channel, the proposed approach was evaluated at various vehicle speeds, signal-to-noise ratios (SNRs), and channel variations in urban and rural areas

The proposed Scheme

- In this paper, the classification accuracy is used as a forward channel quality indicator (CQI) engine to optimise the modulation order, as shown in Fig. 1(a).




Fig. 1(a) Block diagram of the modulation classification.

- High detection accuracy refers to better channel quality, thereby using the same or higher classified modulation order at the transmission and vice versa, as shown in Fig. 1(b).




Fig. 1(b) CQI optimization engine.

- Fig.2 depicts an OFDM communication system at $f_c = 5.9$ GHz for the DSRC, with 256 subcarriers, 64 cyclic-prefix, and 125 subcarriers holding the transmitted data.

The proposed Scheme Cont'd




Fig. 2 Implementation block diagram.

- We acquired 150 training images with dimensions 273×328 pixels for each constellation type. Training samples for different constellations are presented in Fig.3.




Fig. 3 Training samples of different constellations

Simulation Results

- Fig. 4 shows the received spectrum, highlighting the USRP X310 transmitter channel at 5.9 GHz.




Fig. 4 Training samples of different constellations

- Table I lists channel parameter settings.
- We evaluated the classifier's performance at different SNRs, vehicle speeds, and channel variations in urban and rural areas.

Table I: Channel parameter settings

Channel Parameters	Value
Number of multipath components (L)	Urban: 16, Rural: 5
Speed of Tx/Rx	10, 20, 30 m/s
Speed of the scatterer	30 m/s
Azimuth departure/arrival angles ($\alpha_{A(B),l}$)	$U[-\pi, \pi]$
Departure/arrivals' elevation angles ($\beta_{A(B),l}$)	$U[-\pi, \pi]$
Incident reflected scatterers' angles ($\alpha_{l(2),l}$)	$U[0, \pi/3]$
Weibull distribution - scale coefficient (ρ)	2.985
Weibull distribution - shape coefficient (a)	0.428

Simulation Results Cont'd

- Fig. 5 presents the flowchart of the proposed image-based deep learning model.




Fig. 5 The proposed deep learning model architecture

Conclusions

- This paper introduces an efficient deep learning-based modulation optimisation order that saves significant communication overhead compared to channel probing-based approaches.
- By designing a CQI engine block, we optimise the modulation order for V2V channel.
- The evaluation process discussed the effects of the channel variations, vehicle speeds, and SNR values on the classifier's detection probability in urban and rural areas.
- Based on the experimental results, the proposed classifier has sufficient detection probability for an acceptable false detection $\leq 20\%$.
- In future work, we will explore the possibility of testing the classifier on a realistic vehicular wireless channel at varying terminal speeds.

References

- M. A. Al-Shareeda, M. Anbar, and I. Hasbullah, "Survey of Authentication and Privacy Schemes in Vehicular Ad Hoc Networks", IEEE Sensors Journal, vol. 21, no. 2, Jan. 2021.
- M. A. Shawky, Q. H. Abbasi, M. A. Imran, S. Ansari, and A. Taha, "Cross-Layer Authentication based on Physical-Layer Signatures for Secure Vehicular Communication", IEEE Intelligent Vehicles Symposium (IV), pp. 1315-1320, 2022.
- G. Wu, and H. Chu, "Spectrum Sharing with Vehicular Communication in Cognitive Small-Cell Networks", International Journal of Antennas and Propagation, vol. 2020, Mar. 2020.

Bibliography

- [1] C. Zhang, P. Patras and H. Haddadi, ‘Deep learning in mobile and wireless networking: A survey’, *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, 2019.
- [2] O-RAN Alliance. ‘O-RAN Alliance completed its specification release 5 (O-RAN-R005)’, Accessed: 11th Aug. 2026. [Online]. Available: <https://www.o-ran.org/blog/o-ran-alliance-completed-its-specification-release-5-o-ran-r005>.
- [3] T. Zemen et al., *Site-specific radio channel representation for 5g and 6g*, 2024. DOI: 10.48550/arXiv.2406.09025. arXiv: 2406.09025 [eess.SP]. [Online]. Available: <https://arxiv.org/abs/2406.09025>.
- [4] T. Huynh-The et al., ‘Automatic modulation classification: A deep architecture survey’, *IEEE Access*, vol. 9, pp. 142 950–142 971, 2021. DOI: 10.1109/ACCESS.2021.3120419. [Online]. Available: <https://doi.org/10.1109/ACCESS.2021.3120419>.
- [5] M. Vasudevan and M. Yuksel, ‘Machine learning for radio propagation modeling: A comprehensive survey’, *IEEE Open Journal of the Communications Society*, vol. 5, pp. 5123–5153, 2024. DOI: 10.1109/OJCOMS.2024.3446457. [Online]. Available: <https://ieeexplore.ieee.org/document/10640063>.
- [6] S. Bakirtzis, C. Yapar, M. Fiore, J. Zhang and I. Wassell, *Empowering wireless network applications with deep learning-based radio propagation models*, 2024. DOI: 10.48550/arXiv.2408.12193. arXiv: 2408.12193 [eess.SP]. [Online]. Available: <https://arxiv.org/abs/2408.12193>.
- [7] H. Sun et al., ‘Advancing 6g: Survey for explainable ai on communications and network slicing’, *IEEE Open Journal of the Communications Society*, vol. 6, pp. 1372–1412, 2025. DOI: 10.1109/OJCOMS.2025.3534626.
- [8] J. D. Willard, X. Jia, S. Xu, M. S. Steinbach and V. Kumar, ‘Integrating physics-based modeling with machine learning: A survey’, *ArXiv*, vol. abs/2003.04919, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:212657607>.

- [9] A. Karpatne, X. Jia and V. Kumar, *Knowledge-guided machine learning: Current trends and future prospects*, 2024. arXiv: 2403.15989 [cs.LG]. [Online]. Available: <https://arxiv.org/abs/2403.15989>.
- [10] M. Raissi, P. Perdikaris and G. E. Karniadakis, ‘Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations’, *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [11] S. B. A. Zaidi, M. A. Shawky, A. Taha, Q. H. Abbasi, M. A. Imran and S. Ansari, ‘An efficient deep learning-based spectrum awareness approach for vehicular communication’, in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2023, pp. 1–6. DOI: 10.1109/WCNC55385.2023.10118615.
- [12] S. B. A. Zaidi, H. N. Qureshi, W. Raza, M. A. Imran, S. Ansari and A. Imran, ‘IINN: Innately intelligent neural networks for applications in wireless networks and beyond’, 2025, Primary thesis contribution; manuscript under review. Code available at: <https://github.com/syedalibasit/IINN>.
- [13] S. B. A. Zaidi, W. Raza, M. U. B. Farooq, S. Ansari, A. Imran and M. A. Imran, ‘A novel innately-intelligent transfer learning framework for wireless networks and beyond’, in *Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2025, pp. 1–6. DOI: 10.1109/PIMRC62392.2025.11274900.
- [14] U. Masood, H. Farooq, A. Imran and A. Abu-Dayya, ‘Interpretable AI-based large-scale 3D pathloss prediction model for enabling emerging self-driving networks’, *IEEE Transactions on Mobile Computing*, vol. 22, no. 7, pp. 3967–3984, Jul. 2023.
- [15] C. Rudin and J. Radin, ‘Why are we using black box models in AI when we don’t need to? a lesson from an explainable AI competition’, *Harvard Data Science Review*, vol. 1, no. 2, 2019. DOI: 10.1162/99608f92.5a8a3a3d. [Online]. Available: <https://doi.org/10.1162/99608f92.5a8a3a3d>.
- [16] S. B. Ali Zaidi, W. Raza, H. N. Qureshi, M. A. Imran, A. Imran and S. Ansari, ‘A domain-aware framework for interpretable and resilient propagation models: Enabling digital twins for wireless networks’, in *2024 IEEE 99th Vehicular Technology Conference (VTC2024-Spring)*, 2024, pp. 1–6. DOI: 10.1109/VTC2024-Spring62846.2024.10683646.

- [17] P. Linardatos, V. Papastefanopoulos and S. Kotsiantis, ‘Explainable ai: A review of machine learning interpretability methods’, *Entropy*, vol. 23, p. 18, 2021. DOI: 10.3390/e23010018. [Online]. Available: <https://doi.org/10.3390/e23010018>.
- [18] H. N. Qureshi et al., ‘Toward addressing training data scarcity challenge in emerging radio access networks: A survey and framework’, *IEEE Communications Surveys & Tutorials*, vol. 25, no. 3, 2023.
- [19] S. Lundberg and S.-I. Lee, *A unified approach to interpreting model predictions*, 2017. arXiv: 1705.07874 [cs.AI]. [Online]. Available: <https://arxiv.org/abs/1705.07874>.
- [20] M. T. Ribeiro, S. Singh and C. Guestrin, *Why should i trust you?: Explaining the predictions of any classifier*, 2016. arXiv: 1602.04938 [cs.LG]. [Online]. Available: <https://arxiv.org/abs/1602.04938>.
- [21] H. Li, X. Wang, Z. Zhang and W. Zhu, ‘Out-of-distribution generalization on graphs: A survey’, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 11, Nov. 2025.
- [22] K. Li et al., ‘Probing out-of-distribution generalization in machine learning for materials’, *Communications Materials*, vol. 6, no. 9, 2025. DOI: 10.1038/s43246-024-00731-w. [Online]. Available: <https://doi.org/10.1038/s43246-024-00731-w>.
- [23] C. Guo, G. Pleiss, Y. Sun and K. Q. Weinberger, ‘On calibration of modern neural networks’, in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 1321–1330.
- [24] V. Kuleshov, N. Fenner and S. Ermon, ‘Accurate uncertainties for deep learning using calibrated regression’, in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 2796–2804.
- [25] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang and L. Yang, ‘Physics-informed machine learning’, *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [26] K. Gregor and Y. LeCun, ‘Learning fast approximations of sparse coding’, *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 2010.
- [27] J. R. Hershey, J. L. Roux and F. Weninger, ‘Deep unfolding: Model-based inspiration of novel deep architectures’, *arXiv preprint arXiv:1409.2574*, 2014.
- [28] B. Neyshabur, H. Sedghi and C. Zhang, ‘What is being transferred in transfer learning?’, *Advances in neural information processing systems*, vol. 33, pp. 512–523, 2020.

- [29] Y. Guo, H. Shi, A. Kumar, K. Grauman, T. Rosing and R. Feris, *Spottune: Transfer learning through adaptive fine-tuning*, 2018. arXiv: 1811.08737 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/1811.08737>.
- [30] N. Houlsby et al., ‘Parameter-efficient transfer learning for nlp’, in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, 2019, pp. 2790–2799.
- [31] Q. Shi and Y. Karasawa, ‘Automatic modulation identification based on the probability density function of signal phase’, *IEEE Trans. Commun.*, vol. 60, no. 4, pp. 1033–1044, 2012. DOI: 10.1109/TCOMM.2012.022612.110469.
- [32] J. Zheng and Y. Lv, ‘Likelihood-based automatic modulation classification in OFDM with index modulation’, *IEEE Trans. on Vehicular Tech.*, vol. 67, no. 9, pp. 8192–8204, 2018. DOI: 10.1109/TVT.2018.2846758.
- [33] J. H. Lee, K.-Y. Kim and Y. Shin, ‘Feature image-based automatic modulation classification method using CNN algorithm’, in *2019 International Conference on Artificial Intelligence in Information and Communication (ICAIIIC)*, 2019, pp. 1–4. DOI: 10.1109/ICAIIIC.2019.8662523.
- [34] X. Zhang, J. Sun and X. Zhang, ‘Automatic modulation classification based on novel feature extraction algorithms’, *IEEE Access*, vol. 8, pp. 16 362–16 371, 2020. DOI: 10.1109/ACCESS.2020.2965277.
- [35] N. Daldal, A. Sengur, K. Polat and Z. Cömert, ‘A novel demodulation system for base band digital modulation signals based on the deep long short-term memory model’, *Applied Acoustics*, vol. 166, no. 107346, 2020. DOI: 10.1016/j.apacoust.2020.107346.
- [36] Z. Wu, S. Zhou, Z. Yin, B. Ma and Z. Yang, ‘Robust automatic modulation classification under varying noise conditions’, *IEEE Access*, vol. 5, pp. 19 733–19 741, 2017. DOI: 10.1109/ACCESS.2017.2757329.
- [37] U. Masood, H. Farooq, A. Imran and A. Abu-Dayya, ‘Interpretable AI-based large-scale 3d pathloss prediction model for enabling emerging self-driving networks’, *IEEE Transactions on Mobile Computing*, pp. 1–1, 2022. DOI: 10.1109/tmc.2022.3147191. [Online]. Available: <https://doi.org/10.1109/tmc.2022.3147191>.
- [38] Y. Liu, J. Dong, W. Huangfu, J. Liu and K. Long, ‘3.5 ghz outdoor radio signal strength prediction with machine learning based on low-cost geographic features’, *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 6, pp. 4155–4170, 2022. DOI: 10.1109/TAP.2022.3175214.

- [39] S. Wu, B. Ma, T. Ye, J. Zhang, W. Shao and W. Zheng, ‘A machine learning based intelligent propagation model for rsrp prediction’, in *2022 International Seminar on Computer Science and Engineering Technology (SCSET)*, 2022, pp. 1–5. DOI: 10.1109/SCSET55041.2022.00010.
- [40] A. Gupta, J. Du, D. Chizhik, R. A. Valenzuela and M. Sellathurai, ‘Machine learning-based urban canyon path loss prediction using 28 GHz manhattan measurements’, *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 6, pp. 4096–4111, 2022. DOI: 10.1109/tap.2022.3152776. [Online]. Available: <https://doi.org/10.1109/2Ftap.2022.3152776>.
- [41] M. T. Waheed, Y. Fahmy and A. Khattab, ‘Deepchannel: Robust multimodal outdoor channel model prediction in lte networks using deep learning’, *IEEE Access*, vol. 10, pp. 79 289–79 300, 2022. DOI: 10.1109/ACCESS.2022.3194652.
- [42] R.-T. Juang, ‘Explainable deep-learning-based path loss prediction from path profiles in urban environments’, *Applied Sciences*, vol. 11, no. 15, 2021, ISSN: 2076-3417. DOI: 10.3390/app11156690. [Online]. Available: <https://www.mdpi.com/2076-3417/11/15/6690>.
- [43] R.-T. Juang, ‘Deep learning-based path loss model in urban environments using image-to-image translation’, *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 12, pp. 12 081–12 091, 2022. DOI: 10.1109/TAP.2022.3209229.
- [44] V. V. Ratnam et al., ‘Fadenet: Deep learning-based mm-wave large-scale channel fading prediction and its applications’, *IEEE Access*, vol. 9, pp. 3278–3290, 2021. DOI: 10.1109/ACCESS.2020.3048583.
- [45] M. T. Waheed, A. Khattab and Y. Fahmy, ‘Highly accurate multi-modal lte channel prediction via semantic segmentation of satellite images’, in *2022 10th International Japan-Africa Conference on Electronics, Communications, and Computations (JAC-ECC)*, 2022, pp. 90–93. DOI: 10.1109/JAC-ECC56395.2022.10043911.
- [46] U. S. Sani, O. A. Malik and D. T. C. Lai, ‘Improving path loss prediction using environmental feature extraction from satellite images: Hand-crafted vs. convolutional neural network’, *Applied Sciences*, vol. 12, no. 15, 2022, ISSN: 2076-3417. DOI: 10.3390/app12157685. [Online]. Available: <https://www.mdpi.com/2076-3417/12/15/7685>.

- [47] K. Simonyan, A. Vedaldi and A. Zisserman, ‘Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps’, *arXiv preprint arXiv:1312.6034*, 2014. [Online]. Available: <https://arxiv.org/abs/1312.6034>.
- [48] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, D. Pedreschi and F. Giannotti, *A survey of methods for explaining black box models*, 2018. arXiv: 1802.01933 [cs.CY].
- [49] C. Molnar, *Interpretable machine learning — interpretability*, <https://christophm.github.io/interpretable-ml-book/interpretability.html>, Accessed: 10 April 2025.
- [50] M. Raissi, P. Perdikaris and G. Karniadakis, ‘Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations’, *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019, ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2018.10.045>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- [51] T. R. Besold et al., ‘Neural-symbolic learning and reasoning: A survey and interpretation’, *arXiv preprint arXiv:1711.03902*, 2017.
- [52] S. Kim et al., ‘Integration of neural network-based symbolic regression in deep learning for scientific discovery’, *IEEE transactions on neural networks and learning systems*, vol. 32, no. 9, pp. 4166–4177, 2020.
- [53] R. Wang, W. Jiang, Y. Shen, Q. Yue and K.-L. Hsiung, ‘Detecting eavesdropping nodes in the power internet of things based on kolmogorov-arnold networks’, *PloS one*, vol. 20, no. 5, e0321179, 2025.
- [54] S. B. A. Zaidi and S. Ansari, ‘Designing interpretable ai-driven propagation models for advanced wireless networks’, in *Proceedings of the IEEE AP-S/INC-USNC-URSI*, 2024, pp. 1847–1848. DOI: [10.1109/AP-S/INC-USNC-URSI52054.2024.10687134](https://doi.org/10.1109/AP-S/INC-USNC-URSI52054.2024.10687134).
- [55] W. Raza, F. A. Khan, H. N. Qureshi, U. Masood and A. Imran, ‘An ai-driven framework for enhancing resilience in propagation models to enable digital twin’, in *2024 IEEE 35th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2024, pp. 1–6. DOI: [10.1109/PIMRC59610.2024.10817403](https://doi.org/10.1109/PIMRC59610.2024.10817403).

- [56] W. Raza, S. B. A. Zaidi, M. U. B. Farooq, H. N. Qureshi and A. Imran, ‘Refining wireless propagation models using domain-informed gans amid data scarcity’, in *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, 2024, pp. 1–6. DOI: 10.1109/VTC2024-Fall163153.2024.10757539.
- [57] X. Li et al., ‘A disentangled representation learning framework for low-altitude network coverage prediction’, *IEEE Transactions on Mobile Computing*, 2025.
- [58] T. Chen and C. Guestrin, ‘Xgboost: A scalable tree boosting system’, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016.
- [59] G. Ke et al., ‘Lightgbm: A highly efficient gradient boosting decision tree’, *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [60] J. Shao et al., *Wirelessllm: Empowering large language models towards wireless intelligence*, 2024. arXiv: 2405.17053 [cs.NI]. [Online]. Available: <https://arxiv.org/abs/2405.17053>.
- [61] M. S. Gupta and K. Kumar, ‘Progression on spectrum sensing for cognitive radio networks: A survey, classification, challenges and future research issues’, *Journal of Network and Computer Applications*, vol. 143, pp. 47–76, 2019.
- [62] H. Robert, B. Paul, M. Simona and S. Annamaria, ‘Real time broadband electromagnetic spectrum monitoring system based on software defined radio technology’, in *9th International Conference on Modern Power Systems (MPS)*, IEEE, 2021, pp. 1–6.
- [63] T. Shiba, T. Furuichi, K. Akimoto, M. Motoyoshi, S. Kameda and N. Suematsu, ‘Real-time wideband spectrum monitor using multiple sampling frequency direct rf under-sampling for wireless iot’, in *2021 51st European Microwave Conference (EuMC)*, IEEE, 2022, pp. 725–728. DOI: 10.23919/EuMC48056.2021.9781597055325.
- [64] National Instruments, *Wideband spectrum monitoring with NI USRP*, <https://forums.ni.com/t5/Example-Code/Wideband-Spectrum-Monitoring-with-NI-USRP/tap/3533000>, Accessed 10 April 2023.
- [65] M. Al-Shareeda, M. Anbar, I. Hasbullah and S. Manickam, ‘Survey of authentication and privacy schemes in vehicular ad hoc networks’, *IEEE Sensors Journal*, vol. 21, pp. 2422–2433, Jan. 2021. DOI: 10.1109/JSEN.2020.3021731.

- [66] M. A. Shawky, Q. H. Abbasi, M. A. Imran, S. Ansari and A. Taha, ‘Cross-layer authentication based on physical-layer signatures for secure vehicular communication’, in *2022 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2022, pp. 1315–1320.
- [67] G. Wu and H. Chu, ‘Spectrum sharing with vehicular communication in cognitive small-cell networks’, *International Journal of Antennas and Propagation*, vol. 2020, pp. 1–10, Mar. 2020. DOI: 10.1155/2020/6897646.
- [68] C. Mehlfehrer, M. Wrulich, J. Colom Ikuno, D. Bosanska and M. Rupp, ‘Simulating the long term evolution physical layer’, *European Signal Processing Conference (EUSIPCO)*, vol. 27, Jan. 2009.
- [69] P. Karadimas and D. Matolak, ‘Generic stochastic modeling of vehicle-to-vehicle wireless channels’, *Vehicular Communications*, vol. 1, no. 4, pp. 153–167, 2014, ISSN: 2214-2096. DOI: <https://doi.org/10.1016/j.vehcom.2014.08.001>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214209614000400>.
- [70] NI, *Wideband spectrum monitoring with ni usrp*, [Online; accessed April 9, 2023], unknown. [Online]. Available: <https://forums.ni.com/t5/Example-Code/Wideband-Spectrum-Monitoring-with-NI-USRP/ta-p/3533000>.
- [71] U. Masood, H. Farooq, A. Imran and A. Abu-Dayya, ‘Interpretable ai-based large-scale 3d pathloss prediction model for enabling emerging self-driving networks’, *IEEE Transactions on Mobile Computing*, vol. 22, no. 7, pp. 3967–3984, 2023.
- [72] N. Khan, A. Abdallah, A. Celik, A. M. Eltawil and S. Coleri, ‘Explainable ai-aided feature selection and model reduction for drl-based v2x resource allocation’, *IEEE Transactions on Communications*, pp. 1–1, 2025. DOI: 10.1109/TCOMM.2025.3554655.
- [73] A. K. Gizzini, Y. Medjahdi, A. J. Ghandour and L. Clavier, ‘Towards explainable ai for channel estimation in wireless communications’, *IEEE Transactions on Vehicular Technology*, vol. 73, no. 5, pp. 7389–7394, 2024. DOI: 10.1109/TVT.2023.3345632.
- [74] A. Atmakuru, G. D. Fatta, G. Nicosia, A. Varzandian and A. Badii, ‘Sensitivity analysis for feature importance in predicting alzheimer’s disease’, in *LOD (2)*, 2023, pp. 449–465. [Online]. Available: https://doi.org/10.1007/978-3-031-53966-4_33.

- [75] C. Zhang, P.-J. Hoes, S. Wang and Y. Zhao, 'Intrinsically interpretable machine learning-based building energy load prediction method with high accuracy and strong interpretability', *Energy and Built Environment*, 2024, ISSN: 2666-1233. DOI: <https://doi.org/10.1016/j.enbenv.2024.08.006>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666123324000825>.
- [76] Y. Fan, Q. Li, H. Mao and F. Jiang, 'Magnetoencephalography decoding transfer approach: From deep learning models to intrinsically interpretable models', *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 5, pp. 2818–2829, 2024. DOI: 10.1109/JBHI.2024.3365051.
- [77] A. Zappone, M. Di Renzo and M. Debbah, 'Wireless networks design in the era of deep learning: Model-based, ai-based, or both?', *IEEE Transactions on Communications*, vol. 67, no. 10, pp. 7331–7376, 2019. DOI: 10.1109/TCOMM.2019.2924010.
- [78] F. Wang et al., 'Inherently interpretable physics-informed neural network for battery modeling and prognosis', *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 1, pp. 1145–1159, 2025. DOI: 10.1109/TNNLS.2023.3329368.
- [79] H. Song, M. Kim, D. Park, Y. Shin and J.-G. Lee, 'Learning from noisy labels with deep neural networks: A survey', *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 11, pp. 8135–8153, 2023. DOI: 10.1109/TNNLS.2022.3152527.
- [80] H. N. Qureshi et al., 'Toward addressing training data scarcity challenge in emerging radio access networks: A survey and framework', *IEEE Communications Surveys & Tutorials*, vol. 25, no. 3, pp. 1954–1990, 2023. DOI: 10.1109/COMST.2023.3271419.
- [81] C. Zhang, P. Patras and H. Haddadi, 'Deep learning in mobile and wireless networking: A survey', *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2224–2287, 2019. DOI: 10.1109/COMST.2019.2904897.
- [82] J. Xie, L. Sun and Y. F. Zhao, 'On the data quality and imbalance in machine learning-based design and manufacturing—a systematic review', *Engineering*, vol. 45, pp. 105–131, 2025, ISSN: 2095-8099. DOI: <https://doi.org/10.1016/j.eng.2024.04.024>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2095809924003734>.
- [83] Y. Sun et al., 'Accurate and interpretable log-based fault diagnosis using large language models', *IEEE Transactions on Services Computing*, 2025.

- [84] L. Liu, X. Chen, Z. Tang, M. Ma and W. Zhang, *Pinn and gnn-based rf map construction for wireless communication systems*, 2025. arXiv: 2507.22513 [eess.SP]. [Online]. Available: <https://arxiv.org/abs/2507.22513>.
- [85] B. Wan, G. Lei, Y. Guo and J. Zhu, ‘Physics-informed neural networks based on unsupervised learning for multidomain electromagnetic analysis’, *IET Electric Power Applications*, vol. 19, no. 1, e70083, 2025.
- [86] S. M. Lundberg and S.-I. Lee, ‘A unified approach to interpreting model predictions’, *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [87] ‘3GPP TR 38.900 V14.2.0: Study on channel model for frequencies from 0.5 to 100 GHz’, ETSI, Tech. Rep. TR 138 900 V14.2.0, 2017, https://www.etsi.org/deliver/etsi_tr/138900_138999/138900/14.02.00_60/tr_138900v140200p.pdf.
- [88] Forsk, *Atoll Overview*, Online, <https://www.forsk.com/atoll-overview>.
- [89] G. Vrbančič and V. Podgorelec, ‘Transfer learning with adaptive fine-tuning’, *IEEE Access*, vol. 8, pp. 196 197–196 211, 2020. DOI: 10.1109/ACCESS.2020.3034343.
- [90] K. You, Z. Kou, M. Long and J. Wang, ‘Co-tuning for transfer learning’, *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 236–17 246, 2020.
- [91] Y. Gao, F. Hou and R. Wang, ‘A novel two-step fine-tuning framework for transfer learning in low-resource neural machine translation’, in *Findings of the Association for Computational Linguistics: NAACL 2024*, K. Duh, H. Gomez and S. Bethard, Eds., Mexico City, Mexico: Association for Computational Linguistics, Jun. 2024, pp. 3214–3224. DOI: 10.18653/v1/2024.findings-naacl.203. [Online]. Available: <https://aclanthology.org/2024.findings-naacl.203/>.