

**Application of Machine Learning for Power  
System Frequency Stability: Security  
Constrained Optimisation and Adaptive  
Emergency Control**

PhD Thesis

Alinane Brown Kilembe

Power Systems Stability and Dynamics

Institute for Energy and the Environment

Department of Electronic and Electrical Engineering

University of Strathclyde, Glasgow

September 23, 2025

# Dedication

*Dedicated to my dad, who journeyed ahead into our future in March of 2023 —I am  
sure this makes you proud.*

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination, which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by the University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in or derived from this thesis.

Signed: 

Date: September 23, 2025

# Abstract

This thesis develops a novel set of Machine Learning (ML)-based techniques for frequency stability management in modern power systems with high penetration of Converter-Interfaced Generation (CIG). Specifically, the techniques aim to enhance situational awareness by capturing detailed frequency dynamics, which conventional approaches struggle to do, and ensure frequency stability through preventive actions —by introducing ML-based constraints into optimisation models to account for the detailed dynamics —and corrective actions —by implementing agentic ML-based adaptive load shedding, focusing on locational aspects and optimising *where*, *when* and *how much* load to shed. Adapted for both online and offline applications, these techniques equip operators with advanced decision-making tools.

Traditionally, system operators use detailed analytical expressions to manage and ensure secure operation. This approach captures and represents the frequency dynamics of the system by solving the Differential Algebraic Equations (DAEs) of the network. Furthermore, operators usually perform offline or time domain simulations (TDS) analyses on a selected set of scenarios, including expected worst-case scenarios. However, these approaches have significant computational requirements, making them typically suitable only for offline applications with a limited number of scenarios. Even when evaluating such a limited set of scenarios, identifying the most critical ones a priori has become increasingly difficult. As power systems become increasingly complex and uncertain, the challenge intensifies, with the growing number of candidate scenarios creating a significant computational hindrance. Consequently, operators often need to over-secure the system, which comes at a cost and hinders widespread adoption of clean energy to achieve net-zero targets. Moreover, numerical approaches —simplified alter-

native models for fast screening —are increasingly facing accuracy issues due to the increasing complexity of the system, potentially introducing significant errors, which can lead to unforeseen instability.

In response to these challenges, this research proposes an ML-based approach to ensure frequency stability in modern low-inertia networks. ML models are capable of accurately learning complex system relationships using simulation or observational data used during training. As a result, they can predict the system response almost instantly, because solving the network’s DAEs is no longer required to make predictions. Such an approach is especially suitable for real-time or near-real-time applications, providing quick scenario screenings of detailed dynamics, where TDS would be computationally prohibitive. Ultimately, operators are provided with advanced decision-making tools, enabling them to more accurately and efficiently manage frequency stability while imposing minimal computational overhead.

# Contents

<b>Dedication</b>	<b>i</b>
<b>List of Figures</b>	<b>ix</b>
<b>List of Tables</b>	<b>xii</b>
<b>List of Abbreviations</b>	<b>xv</b>
<b>Acknowledgements</b>	<b>xvii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background and Motivation . . . . .	1
1.2 Contributions . . . . .	5
1.3 List of Publications . . . . .	7
1.3.1 Published or Accepted Work . . . . .	8
1.3.2 Submitted and Under Review . . . . .	8
1.3.3 Work in Progress . . . . .	9
1.3.4 Presentations . . . . .	9
<b>2 Literature Review</b>	<b>10</b>
2.1 Power System Frequency Dynamics . . . . .	10
2.2 Understanding and Representing Complex Frequency Dynamics of Power Systems . . . . .	12
2.3 Security Constrained Optimisation for Preventive Frequency Stability .	18

## Contents

2.4	Corrective Power System Frequency Stability Management Using Advanced Frequency Control Techniques . . . . .	21
2.5	Discussion of Current Challenges . . . . .	23
<b>3</b>	<b>Enhancing Power Systems Situational Awareness for Frequency Stability</b>	<b>27</b>
3.1	Frequency Response Dynamics and Locational Aspects in Power Systems	28
3.1.1	Representing Locational Frequency Dynamics in Power Systems	28
3.2	Machine Learning Approach for Capturing and Monitoring of Locational Frequency Dynamics . . . . .	30
3.2.1	Machine Learning Models for Locational Frequency Dynamics Prediction . . . . .	31
3.3	Methodology for Capturing and Representing Locational Frequency Dynamics . . . . .	35
3.3.1	Dataset Generation and Dynamic Simulations . . . . .	35
3.3.2	Data Pre-processing and Training . . . . .	36
3.3.3	Model Evaluation . . . . .	37
3.4	Case Studies and Results . . . . .	38
3.4.1	Test Network and Case study . . . . .	39
3.4.2	Locational Frequency Response (RoCoF) Variations . . . . .	39
3.4.3	Quantification of Locational Frequency Response (RoCoF) Variations . . . . .	41
3.4.4	Predictive Accuracy of the Machine Learning Models . . . . .	43
3.4.5	Computational Aspects . . . . .	46
3.5	Conclusion . . . . .	46
<b>4</b>	<b>Enhancing Understanding of Frequency Dynamics and Preventive Stability Management</b>	<b>48</b>
4.1	Explainable Machine Learning: A SHAP Value-Based Approach to Locational Frequency Stability . . . . .	49

## Contents

4.2	ML Explainability for Understanding Locational Frequency Dynamics through SHapley Additive exPlanations (SHAP) . . . . .	50
4.3	Methodology for Enhancing Understanding of Frequency Dynamics . . .	52
4.3.1	Proposed ML-based approach to understanding Local Frequency Characteristics . . . . .	52
4.3.2	SHAP-Informed System Optimisation for Secure System Dispatch	53
4.3.3	Caveats of Using SHAP Explanations for Extracting Stability Constraints . . . . .	56
4.3.4	Estimating Locational Frequency Dynamics with Machine Learning	56
4.4	Case Studies and Results . . . . .	57
4.4.1	Test Network and Case Study . . . . .	58
4.4.2	Accuracy of the Locational ML Models for Predicting Frequency Stability Metrics . . . . .	59
4.4.3	Understanding the System-wide Minimum RoCoF with SHAP Values . . . . .	60
4.4.4	SHAP-Informed Optimisation - enhancing frequency stability through actionable SHAP insights . . . . .	65
4.4.5	A Discussion on SHAP-Derived Rules for Frequency Stability Enhancement . . . . .	69
4.4.6	Computational Aspects . . . . .	71
4.5	Conclusion . . . . .	72
<b>5</b>	<b>Security Constrained Optimisation for Preventive Frequency Stability</b>	<b>74</b>
5.1	Artificial Neural Network Constrained Optimal Power Flow for Locational Frequency Stability . . . . .	75
5.2	Stability Constrained Optimal Power Flow Problem . . . . .	76
5.3	Artificial Neural Networks (ANNs) for Locational Frequency Stability .	79
5.4	Methodology for Deep Learning Constrained Optimal Power Flow . . .	80
5.4.1	Overview of the Methodology . . . . .	80
5.4.2	Dataset Generation and ML Model Training . . . . .	81

## Contents

5.4.3	Incorporating Stability Constraints Within Optimisation Models	83
5.4.4	Triangle Relaxation of the ReLU Activation Function for Generating Frequency Stable Optimisation Solutions	84
5.4.5	Selective Activation Pruning OPF (SAP- $\Delta$ -OPF)	87
5.4.6	Model Evaluation and Baseline Methods	88
5.5	Case Studies and Results	89
5.5.1	Observed Locational Frequency Deviations in the System Due to High CIG Penetration	90
5.5.2	Case I: Fixed Disturbance Location With Single Area CIG Penetration	92
5.5.3	Case II: Examining the Effect of Disturbances Across all Areas of the System	94
5.5.4	Neural Network Pre-activation Statistics and the Necessity of SAP- $\Delta$ -OPF	97
5.5.5	Computational Considerations	101
5.6	Conclusion	102
<b>6</b>	<b>Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning</b>	<b>104</b>
6.1	Emergency Frequency Control in Power Systems	105
6.1.1	Challenges of Existing Alternatives to the UFLS Scheme	106
6.1.2	Enhancing the Effectiveness of Existing AUFLS RL Methods	107
6.2	Deep Reinforcement Learning Implementation	107
6.2.1	Problem Formulation	109
6.3	Methodology for the Physics-Informed RL-based AUFLS for Frequency Stability	113
6.3.1	Capturing Locational Frequency Dynamics with Artificial Neural Network	115
6.3.2	RoCoF-Based Adaptive System Coherence Detection	116
6.3.3	Physics Shield and Rewards Integration of the RL Agent	118

Contents

6.3.4	Modelling of UFLS Relays and Implementation of Conventional UFLS Scheme . . . . .	122
6.3.5	Battery Energy Storage Systems (BESS) . . . . .	123
6.4	Case Studies and Results . . . . .	125
6.4.1	The Modified IEEE 39-Bus Test Case . . . . .	127
6.4.2	The Texas 2000-Bus Synthetic Grid . . . . .	135
6.4.3	Locational Aspect Considerations and the Effectiveness of Shedding Strategies . . . . .	137
6.4.4	Computational Considerations . . . . .	140
6.5	Conclusion . . . . .	141
<b>7</b>	<b>Conclusions and Future Work</b>	<b>143</b>
7.1	Future Work . . . . .	145
7.1.1	Power System Stability Situational Awareness and Stability Management . . . . .	145
7.1.2	Preventive Frequency Stability Management in Security-Constrained Optimisation . . . . .	147
7.1.3	Adaptive Frequency Control Mechanisms for Emergencies . . . . .	148
	<b>Bibliography</b>	<b>149</b>

# List of Figures

1.1	World primary energy consumption from 2000 to 2022 [1] . . . . .	2
1.2	Basic structure of a converter interfaced generation with a stationary reference frame control [2] . . . . .	3
2.1	Active power balance and frequency stability in power systems . . . . .	11
2.2	Sample frequency oscillations following the loss of a generator in the Texas 2000-bus system . . . . .	13
2.3	Sample frequency oscillations following the loss of a generator in the IEEE 39-bus network with a nominal frequency of 60 Hz . . . . .	15
3.1	Artificial Neural Network (ANN) architecture . . . . .	32
3.2	The Hybrid integrated COI-MLP Model architecture . . . . .	34
3.3	The modified IEEE 39-Bus network highlighting Area 2 where CIG is connected at Bus 16 and location of the disturbance SG-6 (G06) . . . . .	40
3.4	Locational RoCoFs at selected buses in the network compared with COI-based predictions . . . . .	42
3.5	Maximum locational RoCoF variations following the outage of SG-6 . . . . .	43
3.6	A distribution of locational RoCoF prediction errors by different models . . . . .	44
3.7	A distribution of locational nadirs prediction errors by different models . . . . .	45
4.1	A SHAP values-based understanding of system dynamics with concise security rules . . . . .	53

## List of Figures

4.2	Local SHAP explanation of a randomly sampled scenario for the system-wide minimum RoCoF. SHAP values (horizontal axis) are plotted against system features (vertical axis) . . . . .	61
4.3	Global SHAP explanation for the system-wide minimum RoCoF. Colour represents the magnitude of the feature . . . . .	62
4.4	System-wide minimum RoCoF Dependence plot for CIG MW penetration. Colour represents G 04 dispatch in MW . . . . .	64
4.5	Global SHAP explanation for localised RoCoF at Bus 34. Colour represents the magnitude of the feature . . . . .	66
4.6	Case Study 1: RMS-TDS results of the RoCoF at Bus 34. The <i>Base Case</i> is where the standard AC-OPF includes no SHAP insights as constraints . . . . .	67
4.7	Case Study 2: RMS-TDS results of the system-wide minimum RoCoF. The <i>Base Case</i> shows a scenario where the standard AC-OPF includes no SHAP insights as constraints . . . . .	68
4.8	Case Study 3: RMS-TDS results of the system-wide minimum and Bus 34 RoCoF. The ( <i>Base Case</i> ) shows a scenario where the standard AC-OPF includes no SHAP insights as constraints . . . . .	69
4.9	Sample scenarios showing the RoCoF at Bus 34 of <i>Base Case</i> , and following the incorporation of SHAP-informed constraints within the AC-OPF . . . . .	70
5.1	A three-stage methodology for the development of a Neural Network Constrained Optimal Power Flow (OPF) . . . . .	81
5.2	Triangle relaxation (red) of the ReLU activation function (brown) with bounded inputs $z_v \in [\text{LB}, \text{UB}]$ for layer $v$ neuron $l$ . . . . .	85
5.3	The modified IEEE 39-bus network highlighting the three network areas with CIG penetrations at buses 5, 16 and 26 [3] . . . . .	89
5.4	Impact of CIG penetration on the maximum locational RoCoF variations (in Hz/s) between any two locations in the modified IEEE 39 Bus network . . . . .	91
5.5	Locational Frequency Nadir Violations by four different models, including the proposed $\Delta$ – OPF, across different CIG levels (up to 40% of system generation) . . . . .	95

## List of Figures

5.6	Layer-wise preactivation mean distribution of a neural network with three hidden layers . . . . .	99
5.7	Neuron-wise minimum-maximum preactivation distribution of a neural network with three hidden layers and 400 neurons per layer . . . . .	100
5.8	Full neuron preactivation trace for neuron 249 in Layer 2 (mostly inactive)	101
6.1	Architecture of the DDPG algorithm illustrating the actor and critic networks and the simulation environment (TDS) . . . . .	108
6.2	Overview of the NN-Physics-shielded RL agent implementation for adaptive emergency frequency control . . . . .	112
6.3	Battery Energy Storage System (BESS) Model [4] . . . . .	124
6.4	The modified IEEE 39-bus network highlighting Area 2, which is a low inertia region with the CIG location at Bus 16 and three BESS units at Buses 14, 18, and 28 . . . . .	128
6.5	AUFLS training progress of the RL physics-informed agent in the Modified IEEE 39-Bus Network . . . . .	131
6.6	Stacked plot of the load shedding schemes' performance by different methods for emergency frequency control in the Modified IEEE 39-bus network . . . . .	133
6.7	One line diagram of the synthetic Texas2000 case. Orange lines indicate 500 kV, pink lines indicate 230 kV, green lines indicate 161 kV, and black lines indicate 115 kV [5] . . . . .	134
6.8	Load shedding behaviour of different frequency control models focusing on shedding contribution (in MW) from participating loads in the Texas 2000-bus network. . . . .	138
6.9	Efficiency of the emergency load shedding schemes in the Modified IEEE 39-Bus and the Texas 2000-Bus Networks . . . . .	138

# List of Tables

3.1	Cost function coefficient values for Synchronous Generator (SG) polynomial cost-function [3,6] . . . . .	41
3.2	Mean Prediction Errors for RoCoF and Frequency Nadir . . . . .	46
3.3	Comparison of Computational Requirements for Different Models . . . .	47
4.1	Frequency Stability Metric (RoCoF (Hz/s) and Nadir (Hz)) Estimation Performance by Different Models . . . . .	59
4.2	Performance of the SHAP-informed Optimisation Model for three Case Studies . . . . .	71
5.1	Accuracy of the Neural Network Model for locational frequency stability metrics prediction for Case I . . . . .	93
5.2	Model Optimisation Results for Locational Frequency Stability (Nadir) for Case I . . . . .	93
5.3	Accuracy of the Neural Network Model for Locational Frequency Stability Metrics Prediction for Case II . . . . .	96
5.4	Model Optimisation Results for Locational Frequency Stability (Nadir) for Case II . . . . .	97
5.5	Model Optimisation Results for $\Delta$ -OPF and SAP- $\Delta$ -OPF for Case II . .	98
6.1	NN and RL Algorithms Architecture and Hyperparameter Settings . . .	125
6.2	Performance Evaluation of ML Model for RoCoF Prediction . . . . .	127
6.3	Standard 9-Stages Under-Frequency Load Shedding Scheme for the IEEE 39-Bus Network . . . . .	129

List of Tables

6.4	Performance Evaluation of ML Model for RoCoF Prediction in the IEEE 39-Bus Network . . . . .	132
6.5	Mean Load Shedding (MW) for Emergency Frequency Control by Different Methods in the IEEE 39-Bus Network . . . . .	134
6.6	The 6-Stages Under-Frequency Load Shedding Scheme for the Texas 2000-Bus Network [7] . . . . .	135
6.7	Performance Evaluation of ML Model for RoCoF Prediction in the Texas 2000-Bus Network . . . . .	136
6.8	Mean Load Shedding (MW) for Emergency Frequency Control by Different Methods in the Texas 2000-Bus Network . . . . .	137
6.9	Comparison of load shedding (MW) and frequency nadir (Hz) between four load shedding schemes . . . . .	141

# List of Abbreviations

<b>AC</b>	Alternating Current
<b>ANN</b>	Artificial Neural Network
<b>AUFLS</b>	Adaptive Under-Frequency Load Shedding
<b>BESS</b>	Battery Energy Storage System
<b>CCT</b>	Critical Clearing Time
<b>CIG</b>	Converter Interfaced Generation
<b>COI</b>	Centre of Inertia
<b>DAE</b>	Differential-Algebraic Equations
<b>DDPG</b>	Deep Deterministic Policy Gradient
<b>EMT</b>	Electro-Magnetic Transients
<b>FFC</b>	Frequency Formation Coefficient
<b>HIL</b>	Hardware In the Loop
<b>IML</b>	Interpretable Machine Learning
<b>MAE</b>	Mean Absolute Error
<b>MIPS</b>	MATLAB Interior Point Solver
<b>ML</b>	Machine Learning
<b>MLP</b>	Multilayer Perceptron
<b>NN</b>	Neural Network
<b>OC</b>	Operating Condition

## List of Abbreviations

<b>OPF</b>	Optimal Power Flow
<b>PFI</b>	Permutation Feature Importance
<b>PINN</b>	Physics Informed Neural Network
<b>PMU</b>	Phasor Measurement Unit
<b>PV</b>	Photovoltaic
<b>PWL</b>	Piecewise Linear
<b>ReLU</b>	Rectified Linear Unit
<b>RL</b>	Reinforcement Learning
<b>RMS</b>	Root Mean Square
<b>RMSE</b>	Root Mean Square Error
<b>RoCoF</b>	Rate of Change of Frequency
<b>SCADA</b>	Supervisory Control and Data Acquisition
<b>SG</b>	Synchronous Generator
<b>SHAP</b>	SHapley Additive exPlanations
<b>SPC</b>	Synchronizing Power Coefficient
<b>SCUC</b>	Security-Constrained Unit Commitment
<b>TDS</b>	Time Domain Simulation
<b>UFLS</b>	Under-Frequency Load Shedding
<b>WECC</b>	Western Electricity Coordinating Council
<b>WTG</b>	Wind Turbine Generator

# Acknowledgements

First and foremost, I would like to thank the Almighty God and our Lord Jesus Christ for the gift of life, good health, blessings, and opportunities to undertake this PhD successfully at the prestigious University of Strathclyde, Glasgow, United Kingdom. May His name be praised forevermore. Amen.

Secondly, I would like to acknowledge my primary supervisor, Dr Panagiotis Papadopoulos, for his guidance and insights throughout my PhD journey. His support has kept me motivated and enabled me to devise innovative solutions to the challenges associated with modern grids.

I would also like to thank my secondary supervisor, Dr Waqqas Bukhsh, whose insights helped me develop a broader perspective on the challenges I am addressing. His input has been invaluable. I am also grateful to my research group for their ongoing support and assistance throughout my studies.

As a Commonwealth Scholar, I would like to express my gratitude to the Commonwealth Scholarship Commission (CSC) for their financial support, which made this PhD possible under a fully funded program.

Finally, I would like to acknowledge the unwavering support and encouragement of my family and friends throughout my PhD journey—for their patience in listening to explanations they never asked for, and for so graciously pretending to understand them!

May God bless you all and keep you in good health.

*And the earth was without form, and void; and darkness was upon the face of the deep. And the Spirit of God moved upon the face of the waters.*

— Genesis 1:2

# Chapter 1

## Introduction

### 1.1 Background and Motivation

The data on the world's primary energy consumption shows that energy consumption from renewables has steadily increased since the early 2000s [1]. This shift is driven by the need to fight climate change and achieve net-zero goals by reducing greenhouse gas emissions from electricity generation. As a result, fossil fuel-based generation has decreased significantly, while renewable energy technologies, mainly wind and solar PV, have expanded. Fig. 1.1 illustrates the world's primary energy consumption from 2000 to 2022, showing a sharp increase in energy consumption from renewables. Over this period, the share of renewables in the global primary energy mix rose from nearly 0% to 8%, whereas the consumption from fossil fuel sources generally declined. Notably, oil's contribution dropped from approximately 40% to 30%. The UK is among the most ambitious countries pursuing this transition, with a net-zero target set for 2050 [8]. As of July 2024, 47% of the electricity generated in the country came from renewable sources [9]. Recently, the UK government accelerated its net-zero ambitions to achieve full decarbonisation of the electricity system by 2035. However, in December 2024, this target was further revised to 2030 by introducing the Clean Power 2030 Action Plan [10]. Additionally, the transition will be supported by the electrification of heating and transportation, which will significantly increase electricity demand and, in turn, necessitate higher levels of renewable generation capacity and grid infrastructure.

## Chapter 1. Introduction

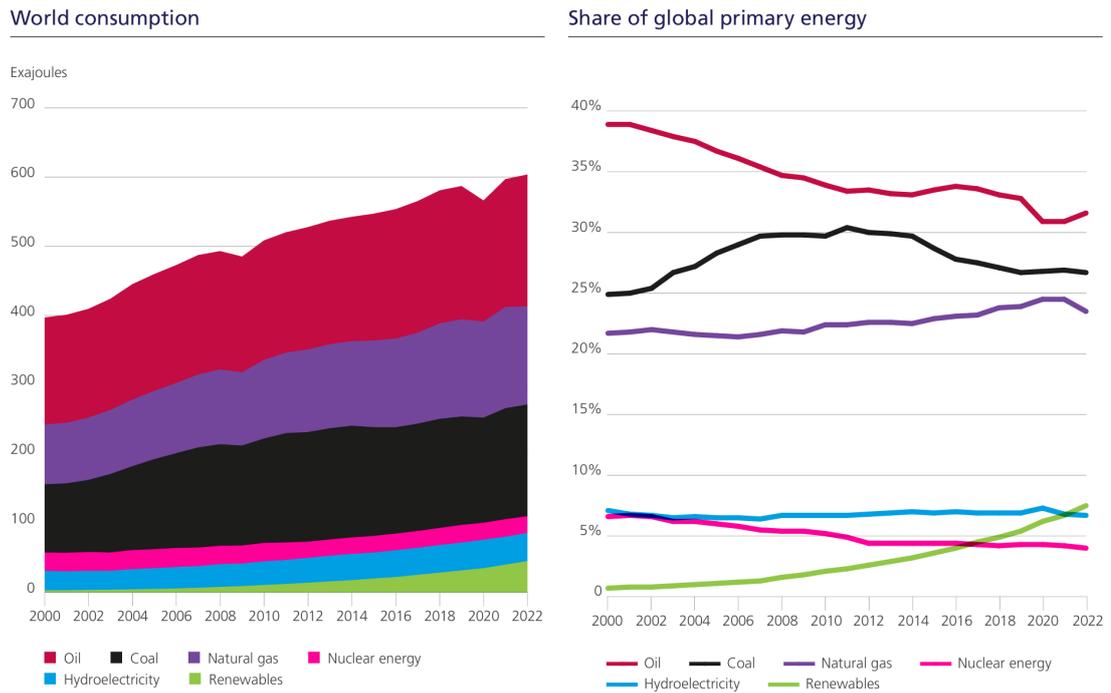


Figure 1.1: World primary energy consumption from 2000 to 2022 [1]

However, while these ambitions are noble, renewable energy technologies, also known as Converter Interfaced Generation (CIG), are different from conventional SG sources concerning their operational characteristics, introducing uncertainties and complexities into the system. Specifically, power systems rely on the availability of sufficient inertia, primarily supplied by the SG's rotating masses, to withstand disturbances such as loss of load or generation. However, renewables cannot inherently provide this inertia, as they are coupled to the grid through power electronic converters. Despite the ability of CIGs to provide synthetic inertia, which is achieved through fast control actions given a continuous energy supply, this inertia differs from the response of SGs' rotating masses, leading to increased system sensitivity and vulnerability to disturbances. This, along with other technical differences, increases the system's vulnerability to disturbances [2]. As a result, this shift presents new challenges for operators to ensure the safe operation of the power system. One of the key security concerns for system operators is frequency stability, that is, the system's ability to maintain a stable frequency by balancing generation and demand. This task becomes increasingly difficult as CIGs

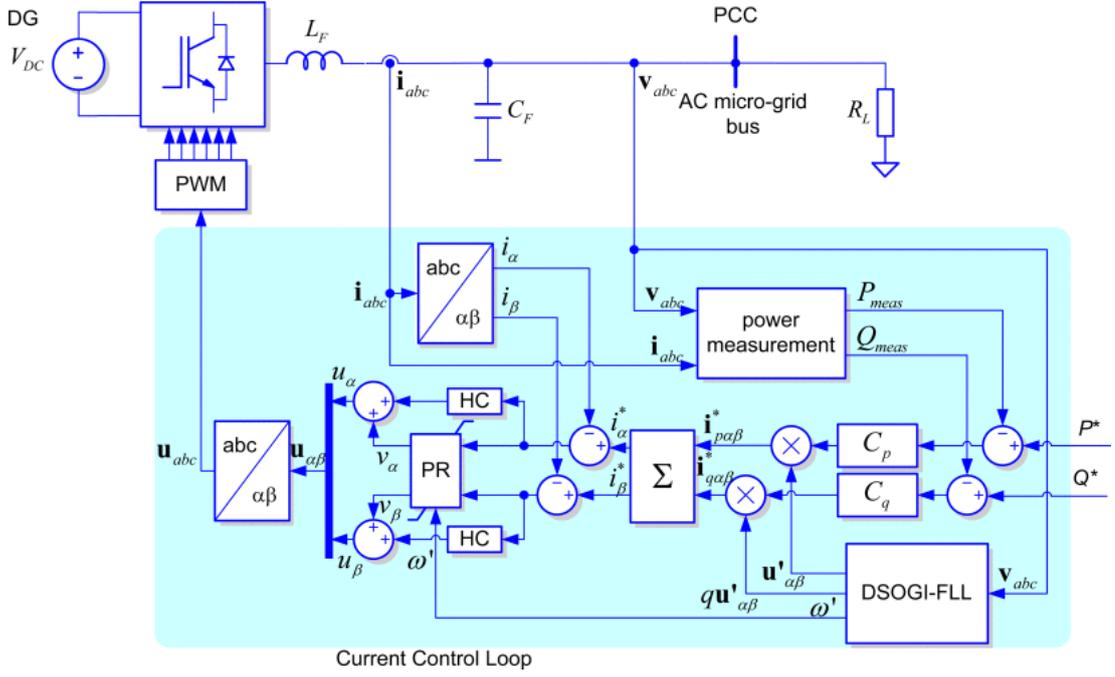


Figure 1.2: Basic structure of a converter interfaced generation with a stationary reference frame control [2]

are often weather-dependent with unpredictable generation.

The frequency behaviour of a power network is typically modelled using the Centre of Inertia (COI) approach, where all online SGs are combined into a single rotating mass coupled to an infinite bus. This method enables a global perspective on the frequency response, informing various critical security measures such as the allocation of frequency control resources. For example, the Electric Reliability Council of Texas (ERCOT) used the COI to determine the inertia floor for its network to be greater than 100 GW.s [11]. Other systems worldwide have set inertial floors as well, using this technique, with 135 GW.s for Great Britain, 23 GW.s for Ireland, 125 GW.s for the Nordic interconnect, and 6.2 GW.s for Australia [12, 13]. Furthermore, in preventive security-constrained optimisation models such as Unit Commitment (UC) and Optimal Power Flow (OPF), security constraints are typically founded on this concept [12]. Similarly, in real-time control techniques, the COI is utilised to determine the control strategy under emergencies. Yet, while this approach is reliable in grids dominated by SGs where the frequency response is more centralised, high integration of CIG poses a

new challenge as the frequency dynamics are increasingly becoming localised [14–16]. As a result, even if the system-wide frequency response, based on the COI model, appears stable, locational frequency response violations can still occur. This is due to reduced dynamic coupling between areas with high local CIG penetration and the rest of the system, which limits the inertia support from other regions. Consequently, this can lead to unforeseen blackouts because of frequency-related events, such as Rate of Change of Frequency (RoCoF) and under-frequency relay activations [12, 14, 15].

To account for the frequency dynamics in detail and avoid these challenges, time domain simulations (TDS) are required. This approach involves the use of a full, detailed phasor domain power system simulator, such as PowerFactory [17], to simulate the power system. However, TDS is a computationally intensive process due to the solving of the network’s Differential Algebraic Equations (DAEs). This issue becomes particularly problematic in modern systems where there are growing CIG-related uncertainties, leading to an increased number of scenarios that need to be assessed. This, in turn, renders such applications computationally prohibitive for applications in real-time or close-to-real-time settings, that is, during operational timescales. Analytical and/or numerical approaches have been proposed to address these challenges, offering improved accuracy than the traditional COI-based methods. However, these approaches often encounter accuracy limitations due to significant complexities being introduced, and as more details are considered, they become computationally intensive. This highlights the need for more reliable methods capable of accurately accounting for these complexities without imposing a significant computational overhead.

Data-driven or Machine Learning (ML) techniques have gained traction in various power systems applications [3, 18, 19]. These methods can establish complex relationships with high accuracy based on simulation or observed data during training. They do not rely on directly solving the DAEs to make a prediction. As a result, they are computationally efficient and suitable for real and close-to-time applications. Therefore, in response to the challenges faced by conventional methods in the face of the increasing displacement of SGs with CIGs to achieve net zero ambitions, this thesis takes a different approach by developing a novel set of ML-based models to efficiently

and accurately account for the detailed frequency dynamics necessary to ensure the frequency stability of modern power networks.

## 1.2 Contributions

This thesis tackles the increasing complexity of modern power systems, driven by the high integration of CIGs, using an ML-based approach. Specifically, it addresses frequency stability challenges related to situational awareness, as well as preventive and corrective frequency stability management. The proposed approach aims to improve both accuracy and computational efficiency in frequency stability applications, challenging conventional approaches. The key contributions of this thesis are summarised as follows:

- i. *Accurately capturing and representing detailed locational frequency dynamics of the network:* We propose an ML-based method to predict the detailed locational frequency dynamics of the network, which would otherwise be hard and computationally intensive to represent using existing analytical methods — which relies on the explicit formulation of power system equations — and where TDS would be computationally prohibitive. The proposed method is a data-driven technique capable of learning the complex relationships of the system based on observational data. Unlike the conventional analytical approaches, this eliminates the need to solve DAEs to predict the dynamics. As a result, the proposed approach can predict almost instantaneously (under 100 milliseconds), while maintaining a high enough accuracy similar to that of TDS.
- ii. *Enhancing the understanding of frequency dynamics in the network amid the growing complexity and numerous interacting devices:* As systems become more complex and grow in scale with numerous interacting devices, understanding various dynamics at play becomes crucial for maintaining stability. We propose an explainable ML-based approach — a technique that enhances the understanding of ML algorithms by explaining the logic behind their predictions — to generate concise system stability rules, providing insights into the system’s dynamics. We

implement a game-theoretic approach using additive feature-attribution methods to explain the behaviour of highly accurate black-box ML models used to capture detailed dynamics of the system. This allows planners and operators to pinpoint important system variables shaping the system’s stability boundary, without entirely relying on the computationally intensive TDS for insights. Such an approach could also inform system optimisation models to enhance secure system operation, by leveraging such insights as optimisation rules or preventive interventions in stability-aware generator dispatch or unit commitment.

- iii. *Accounting for the system’s increasing complexity by directly representing detailed frequency dynamics within security-constrained optimisation models:* Deriving accurate analytical expressions of detailed locational frequency responses in systems with high CIG integration is a challenging task. These expressions, later integrated into security-constrained models, i.e., optimisation models with stability-oriented constraints, such as in Security-Constrained Unit Commitment (SCUC), are necessary to prevent locational frequency violations. However, they can incur significant computational budgets, especially since they are part of an optimisation routine, and typically require solving algebraic equations numerically equivalent to DAEs. As a result, necessary simplifications, such as those offered by COI-based methods, are often employed. We propose to overcome these limitations by directly representing the detailed frequency dynamics within security-constrained optimisation models using a reformulated ML model as constraint(s). Since ML models are trained to capture such detailed dynamics, the optimisation model no longer requires solving DAEs to converge to a frequency-stable solution. As a result, the optimisation process becomes more efficient, achieving faster convergence, thereby facilitating near-real-time applications. Furthermore, in analytical-based constraints such as the Piecewise Linear (PWL) approximation of differential equations, the number of constraints significantly increases with each introduction of new devices such as SGs, thereby increasing complexity and further prolonging solver convergence due to the corresponding DAEs introduced. In contrast, the number of constraints imposed by the proposed approach

remains constant regardless of the introduced changes.

- iv. *Developing an efficiently adaptive frequency control scheme for emergencies:* Real-time frequency control actions are integral to maintaining the integrity of power systems. The conventional under-frequency load-shedding scheme (UFLS) is inefficient as it operates based on the violation of pre-set frequency thresholds, triggering predefined shedding of a certain load percentage at various steps. We propose an adaptive reinforcement learning (RL) agent-based emergency frequency control approach, capable of determining an efficient load-shedding strategy in real-time by adaptively optimising *when*, *where* and *how much* load to shed to maintain stability. The proposed approach employs a two-layered protection scheme, retaining the conventional emergency control technique as the last resort to safeguard the agent’s actions for robustness. Furthermore, we introduce the physics governing the dynamic system within the RL agent’s training process, enabling more efficient and optimal performance over conventional purely data-driven agents that are trained through excessive sampling. The benefits offered by this approach address the common challenges faced by conventional RL methods in modern large-scale power systems as follows: i) reduced reliance on extensive training data, ii) enhanced sampling and solution optimality by eliminating non-essential actions, through a Physics Shield (PS) —post-action filtering through the swing equation to constrain the search space, iii) faster convergence by limiting RL agents’ non-beneficial exploration, iv) better extrapolation capabilities resulting from the system’s physics accounted for during the model’s training and, iv) less vulnerability to dimensionality through a physics-informed dimensionality reduction technique.

### 1.3 List of Publications

During the course of this research, the following works were published:

### 1.3.1 Published or Accepted Work

[**Pub A**] **A. B. Kilembe**, P. N. Papadopoulos and R. I., Hamilton (Accepted/In press), “Data-driven approach to capturing wide-area frequency response dynamics,” 59th International Universities Power Engineering Conference (UPEC), Cardiff, Wales, 2024. (**1st Prize for Best Paper, 59th UPEC**).

[**Pub B**] **A. B. Kilembe**, P. N. Papadopoulos and R. I., Hamilton, “Understanding Local Frequency Characteristics Using Explainable Machine Learning: A SHAP-Based Approach,” International Journal of Electrical Power and Energy Systems, 170, Article 110885.

[**Pub C**] **A. B. Kilembe** and P. N. Papadopoulos, “Regulation of Disturbance Magnitude for Locational Frequency Stability Using Machine Learning,” 2023 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), Glasgow, United Kingdom, 2023, pp. 1-6.

[**Pub E**] **A. B. Kilembe** and P. N. Papadopoulos (Accepted/In press), “Deep reinforcement learning adaptive under-frequency load shedding for frequency control under extreme events,” 14th Mediterranean Conference on Power Generation, Transmission, Distribution and Energy Conversion (MEDPOWER 2024), Athens, Greece.

[**Pub H**] **A. B. Kilembe**, W. Bukhsh and P. N. Papadopoulos, “Learning of Wide-Area Dynamics in Power Systems with Physics Informed Neural Networks,” 2025 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), Toronto, Canada.

### 1.3.2 Submitted and Under Review

[**Pub D**] **A. B. Kilembe**, P. N. Papadopoulos and W. Bukhsh, “Convex Relaxation of Neural Network with Selective Activation Pruning for Locational Frequency

## Chapter 1. Introduction

Stability,” (IEEE Transactions on Sustainable Energy —as of August 2025).

[**Pub F**] **A. B. Kilembe** and P. N. Papadopoulos, “Physics-Shielded Deep Reinforcement Learning for Adaptive Emergency Frequency Control,” (IEEE Transactions on Power Systems —as of August 2025).

[**Pub G**] **A. B. Kilembe**, P. N. Papadopoulos and W. Bukhsh, “Linear Constraints with Critical Region Refinement for Frequency Stability-Constrained Optimisation,” (Conference paper).

### 1.3.3 Work in Progress

[**Pub I**] C. McGarry, J. Dickson, O. Smith, P. Fennell and **A. B. Kilembe**, “Coordination of heat pumps and electric vehicle charging demand incorporating building physics modelling,” (Journal paper —under preparation: The University of Strathclyde and The University College London).

### 1.3.4 Presentations

**A. B. Kilembe** and P. N. Papadopoulos, “Regulation of Disturbance Magnitude for Locational Frequency Stability Using Machine Learning,” presented at the Manchester Energy and Electrical Power Systems (MEEPS), Manchester, UK, Nov 2024.

**A. B. Kilembe**, W. Bukhsh and P. N. Papadopoulos, “Physics-Informed Deep Reinforcement Learning for Adaptive Load Shedding Scheme in Power Systems with Heterogeneous Resources,” presented at the Manchester Energy and Electrical Power Systems (MEEPS), Manchester, UK, Nov 2023.

## Chapter 2

# Literature Review

In this chapter, we present the literature review that motivates the research conducted in this thesis. We provide a technical introduction to the frequency dynamics of the power system, the main focus of this thesis. Then, we subdivide the literature discussion into three key aspects of frequency stability management as follows: First, we focus on understanding and representing the detailed frequency dynamics in systems with high CIG integration; second, we explore preventive techniques aimed at accurately addressing frequency stability issues in near-real time and planning timescales; and third, we move into corrective techniques for real-time applications, focusing on advanced adaptive frequency control during emergencies.

### 2.1 Power System Frequency Dynamics

Frequency stability refers to the ability of the system to maintain a steady and acceptable frequency after a disturbance. An illustration of this concept is shown in Fig. 2.1. To maintain nominal frequency,  $f_{\text{nom}}$ , the system generation and demand must remain balanced. If the system generation exceeds demand, the system exhibits over-frequency, and vice versa. If the demand exceeds generation, under-frequency occurs. The available inertial response in the system can be quantified as the available immediate energy,  $E_{\text{sys}}$ , that is injected into the power system before the activation of reserves [12, 20]. This is the total contribution from  $N$  online SGs, stated as:

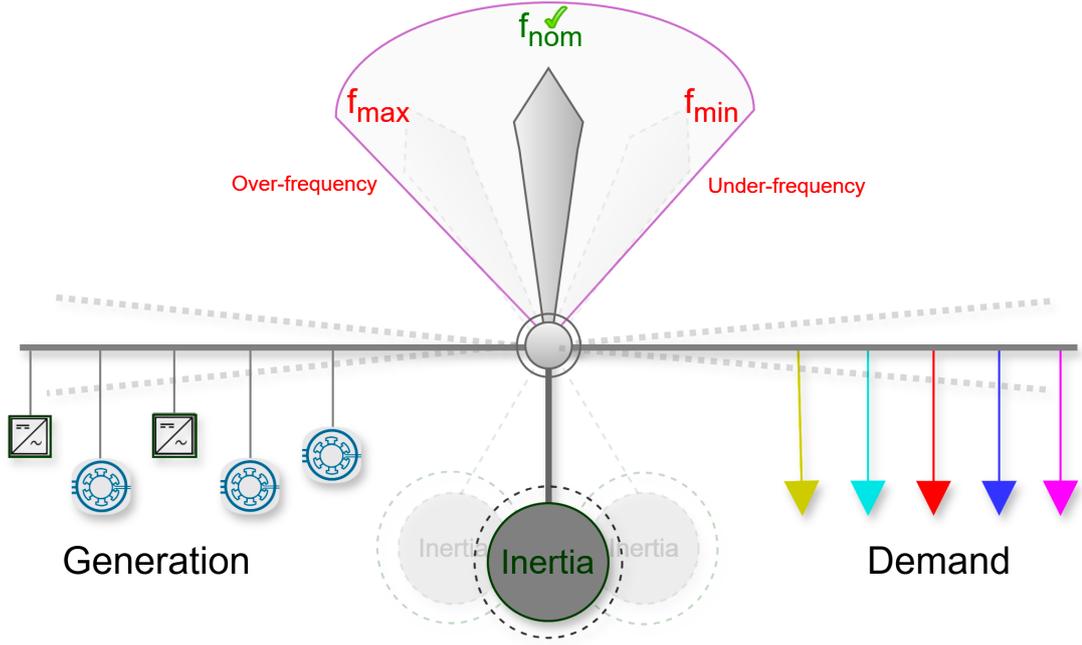


Figure 2.1: Active power balance and frequency stability in power systems

$$E_{\text{sys}} = \sum_{i=1}^N \frac{1}{2} J_i \omega_i^2 \quad (2.1)$$

where  $J_i$  is the moment of inertia of the shaft in  $kg.m^2.s$  and  $\omega_i$  is the angular velocity in  $rad/s$ . Power systems maintain a balance between generation and demand (including losses) for a steady frequency of the system. Immediately following a disturbance, such as a loss of generation, the RoCoF of the system frequency can be estimated according to an equivalent swing equation [21, 22] of a single machine (or an aggregated system response) stated as follows:

$$\frac{df(t)}{dt} = \frac{P_m(t) - P_e(t) - D\Delta f(t)}{M_H} \quad (2.2)$$

where  $P_e(t)$  represents the electric power demanded by the load (including losses) in the system, while  $P_m(t)$  represents the mechanical power of the generator(s).  $M_H$  is the inertia coefficient of the system, while  $D$  is the damping coefficient. At the onset of the disturbance, where  $\Delta P_m = 0$  and damping coefficient component,  $D\Delta f(t) = 0$ , the

system response is due to the disturbance magnitude,  $\Delta P_e(t)$ , and the available inertia coefficient,  $M_H$ .

The overall imbalance in the system following the disturbance can be represented as  $\Delta P$ . To capture the system's frequency nadir following the disturbance, the energy lost must be equal to or less than the energy injected into the system; otherwise, the frequency decline is perpetual [14]. This relationship can be stated as follows:

$$\int_{t=0}^{t_{\text{Nadir}}} \Delta P(t) dt \leq \left( \frac{2H_t \Delta f_{\text{max}}}{f_{\text{nom}}} \right) + D \int_{t=0}^{t_{\text{Nadir}}} \Delta f(t) dt + \int_{t_{\text{Activ.}}}^{t_{\text{Activ.}}+t'} FR(\tau) d\tau \quad (2.3)$$

where  $t_{\text{Nadir}}$  is the time the frequency nadir occurs,  $t = 0$  is the time of disturbance,  $f_{\text{nom}}$  is the nominal frequency,  $t_{\text{Activ.}}$  is the response activation time,  $f_{\text{max}}$  is the maximum permissible deviation,  $H_t$  is the aggregated system synchronous inertia in MWs, and  $t'$  is the delivery period ( $t' \leq t_{\text{Nadir}} - t_{\text{Activ.}}$ ) of the frequency response service (in sec.). The first part corresponds to the energy lost as a result of the disturbance. The first term on the right-hand side of the equation corresponds to the energy injection due to inertia; the second term is the energy 'contribution' from damping, while the last term,  $FR$ , is the energy contribution from frequency control resources due to governor response, fast frequency control, and/or energy imports.

## 2.2 Understanding and Representing Complex Frequency Dynamics of Power Systems

The proliferation of CIGs introduces new challenges to the safe and reliable operation of the power system. Unlike SGs, which provide inertia —the kinetic energy stored in their rotating mass —CIGs are coupled to the grid through power electronics, preventing their ability to directly contribute to system inertia [21, 23]. Whenever the system experiences an imbalance resulting from a disturbance, such as a loss of generation, inertia acts as the immediately available energy released into the power network. This,

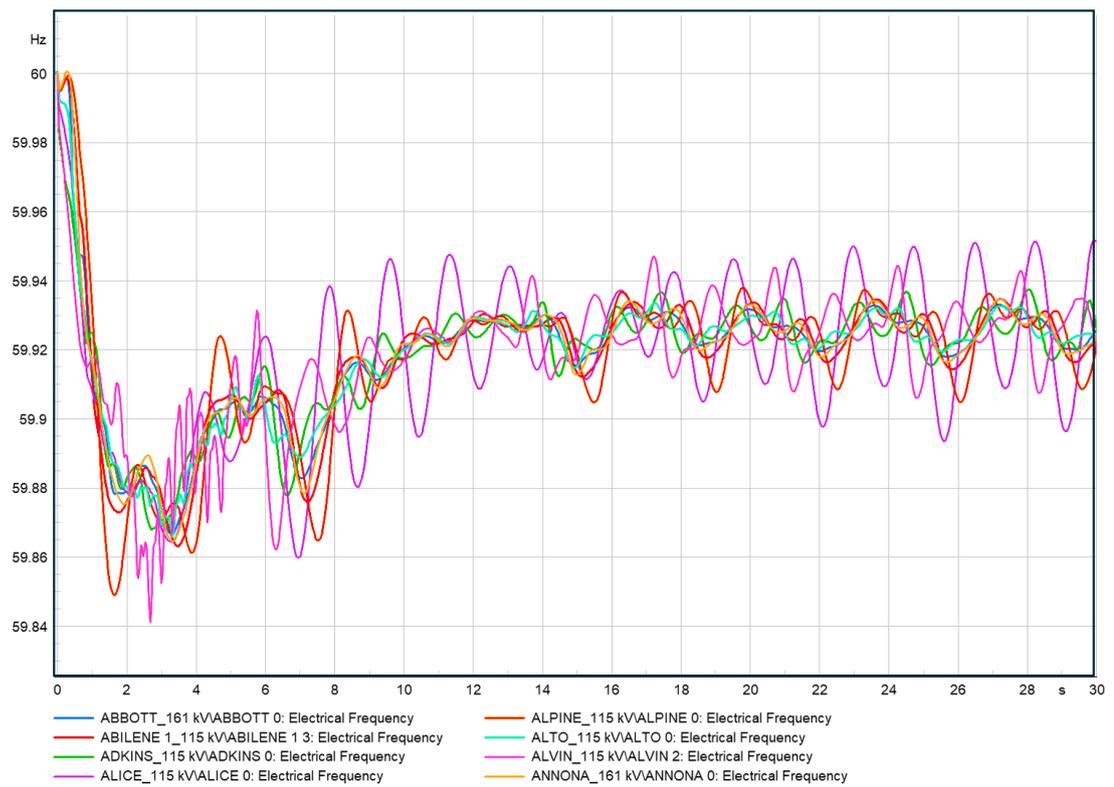


Figure 2.2: Sample frequency oscillations following the loss of a generator in the Texas 2000-bus system

in turn, helps the system to resist rapid frequency declines until the frequency response and control systems restore balance [12].

Typically, in frequency response studies and the allocation of minimum frequency control resources, the COI-based approach is used —where all SGs in a system are combined into a single rotating mass coupled to an infinite bus, thus modelling the network frequency response as a global phenomenon [12, 24]. In a network with  $N$  buses and  $n$  SGs, this function is expressed as,  $\forall i \in \{1, \dots, N\}$ :

$$f_{\text{COI}}(t) = \frac{\sum_{i=1}^n H_i f_i^g(t)}{\sum_{i=1}^n H_i} \quad (2.4)$$

where  $f_{\text{COI}}(t)$  is the average frequency response of the system, which is weighted by the inertia contribution,  $H_i$ , of  $n$  online SGs and  $f_i^g(t)$  is the rotor speed of the  $i^{\text{th}}$  generator. Employing COI-based methods in frequency studies is a widespread approach; for instance, the Electric Reliability Council of Texas determined its minimum network inertia as 100 GW.s, while Great Britain determined 135 GW.s for its network through this approach. However, systems with high CIG integration are experiencing increasing locational frequency variations (similar to that shown in Fig. 2.3) compared to SG-dominated systems [12, 14]. Consequently, the conventional approach of representing the frequency dynamics of the network using COI-based methods is becoming increasingly unreliable, as it puts the network at risk of unforeseen protection relay activation, capable of causing large-scale blackouts [15]. Moreover, the cost of generation increases when the system is overly secured to mitigate such risks, which undermines sustainability efforts to make clean energy more accessible.

The frequency divider (FD) is an analytical method proposed in [25, 26] to address the limitation of conventional COI-based approaches. It uses an augmented system admittance matrix to estimate the local bus frequency, including systems with CIG integration. Given a steady-state model where SG frequencies determine the boundary conditions, the FD operates on the principle that the frequency at a particular bus is affected by power flows —by defining the quota of active power that contributes to frequency variations of buses —and the dynamics of SGs connected to it, weighted by electrical distances. [12] also demonstrated this phenomenon on the Texas7000 model [5, 27]

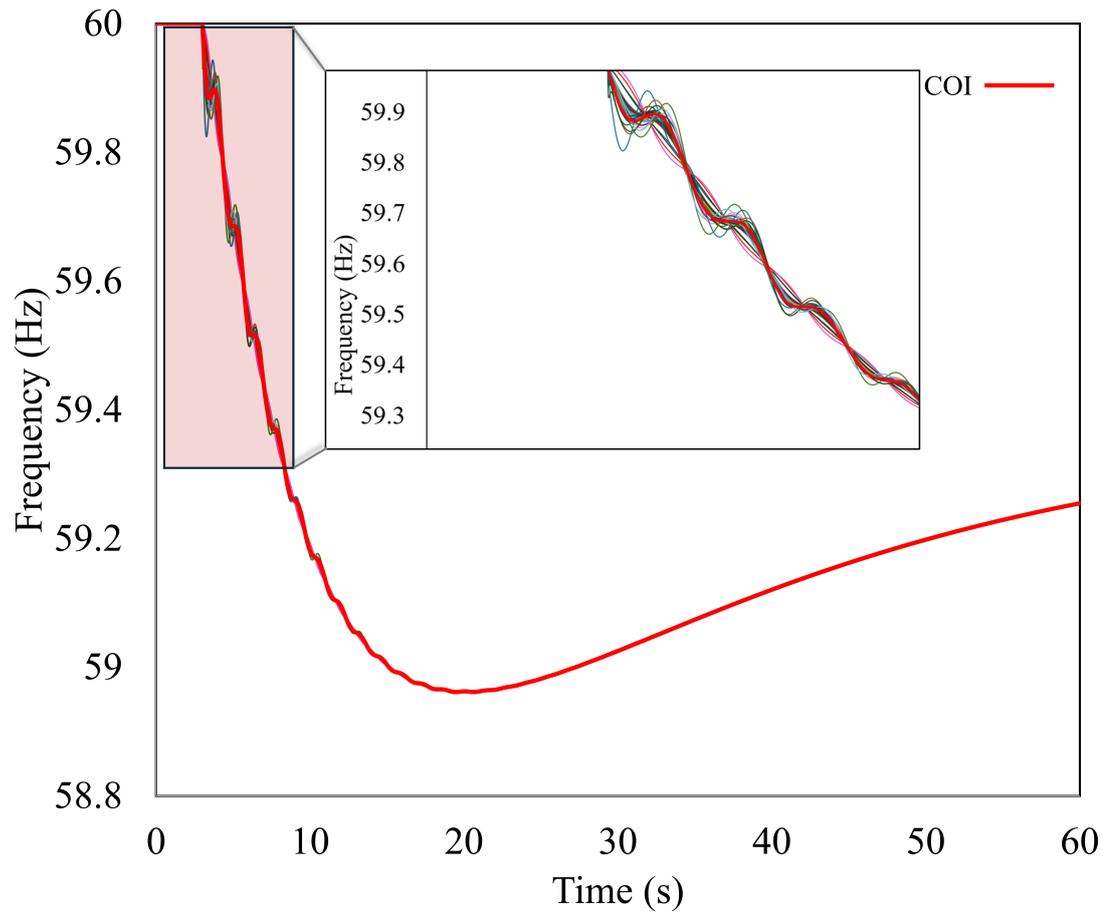


Figure 2.3: Sample frequency oscillations following the loss of a generator in the IEEE 39-bus network with a nominal frequency of 60 Hz

by examining locational frequency response variations under different operating conditions (OCs), showing that various parts of the system crossing the stability boundary (such as RoCoF and nadir) may remain undetected by COI-based methods, potentially leading to unforeseen relay activation. To address this challenge, the paper proposes utilising the numerical derivative of the voltage phase angle to model locational frequency dynamics. [14, 15] demonstrated that in a CIG-dominated power system composed of two areas with significantly different inertia levels, the lower inertia area is the most sensitive to disturbances. This can result in instances of locational frequency violations when using a single COI-based response to represent the entire system. Analytical approaches, complemented by conservative regression —whose weights are calculated through a linear regression on training samples generated by RMS-TDS —were proposed in the paper [14] to accurately determine area-specific-COI-based responses to minimise such risks of locational violations. However, effectively defining such areas in a network is challenging, as they are OC-specific (in terms of the number of areas and composition), rather than static (two areas considered), as done in this study.

Despite the potential demonstrated by analytical methods [12, 16, 25, 28] in addressing the challenges associated with the increasingly locational frequency dynamics of modern networks, these methods, including the frequency divider [26] and the derivative of voltage phase angles [12], continue to face accuracy issues due to the growing system complexity, and can be computationally intensive in high-dimensional settings involving detailed solving of the network’s DAEs —making them less suitable where fast assessment is necessary. Measurement or data-driven approaches have been proposed in the literature to address these limitations. Data-driven techniques use Machine Learning (ML) models to make predictions based on the learned relationships from the training data —leading to fast estimation time and the learning of ‘hard-to-model’ relationships. The Kalman filter was employed for online estimation of the COI frequency response in the IEEE 39 bus network by [29]. A model decoupling strategy was applied, utilising generator active power outputs as the swing equation input for adaptive COI frequency estimation in real-time applications. [30] employs locational Nonlinear Auto-Regressive (NAR) models based on an Artificial Neural Network (ANN) to predict the

locational frequency nadir and the time the nadir occurs in the Nordic 32 Test System, using post-disturbance bus frequency samples as inputs. [29, 31] applied data-driven techniques for a better online estimation of the COI frequency response using Phasor Measurement Units (PMUs) data. Similarly, [24, 32] used data-driven techniques to enhance the accuracy of the COI-based model in power systems with increased complexity due to interconnections and CIG penetration. However, despite the numerous data-driven applications in the literature, existing methods focus only on improving the accuracy of the COI model, without addressing the emerging challenge of locational frequency dynamics, driven by the high integration of CIGs. Moreover, such ML-based methods consider predictive accuracy more than model interpretability, hindering the necessary trust required for adoption among practitioners in the safety-critical power system infrastructure [33].

In [3, 34, 35], the authors proposed using Interpretable ML (IML) techniques to enhance understanding of the entire transient stability boundary. IML aims to reveal the decision-making processes of ML models by making them understandable, rather than simply focusing on the ML model’s predictive accuracy. Both methods propose locational ML models trained to predict the Critical Clearing Time (CCT)—the maximum allowable time after a fault occurs within which the fault must be cleared to ensure that the power system remains stable—at each busbar in the network. Thereafter, IML is used to uncover underlying system relationships and variables influencing such stability boundaries at each fault location. These insights could inform various interventions to enhance the system’s transient stability. [3] achieves this using Decision Trees (DTs) and Permutation Feature Importance (PFI). However, DTs may not always be able to achieve high accuracy in more complex scenarios, and PFI is only able to provide information regarding feature importance and not feature effect. [35] goes further to address this limitation by using SHapley Additive exPlanations (SHAP), which can extract both the importance of features and the effect on high-accuracy black-box models like ANNs. SHAP values are therefore used to derive a detailed location-specific understanding of what affects the transient stability of the system. Similarly, in [36], a SHAP-value-based approach is proposed to uncover the underlying frequency stability

relationships in the system. Insights are derived using SHAP values to explain the effects of various features on the system frequency stability. By combining these insights with domain knowledge, the approach aims to explain essential aspects of power system operation by identifying the main drivers and stability risks. Based on the literature above, three main challenges are observed: i) explicit representation of the emerging locational frequency dynamics to prevent unforeseen local frequency violations remains underexplored, ii) current IML and explainable methods often focus on getting insights related to dynamics by explaining the behaviour of the black-box models, and iii) while the methods aim to uncover the underlying system relationships, they do not sufficiently go beyond statistically explaining the insights into verifying them by defining rules such as through preventive actions, to quantify their effectiveness. In our work, we go beyond gaining insights by formulating SHAP-informed interventions in the form of constraints to enhance system stability through targeted preventive actions, with their final effectiveness validated via RMS-TDS —making their potential in practical application uncertain.

### **2.3 Security Constrained Optimisation for Preventive Frequency Stability**

Security-constrained optimisation involves incorporating stability-oriented constraints within the model to ensure secure system operation [16,37]. Traditionally, these methods rely on static constraints, such as minimum system inertia [12], to ensure frequency stability. This approach, which often employs COI-based constraints, is insufficient for application in modern power systems with high CIG penetration. This is because in such systems, the frequency dynamics are increasingly localised, and relying on the COI to represent these responses poses risks of undetectable local frequency violations, potentially causing blackouts [12,14–16].

Addressing the limitations of COI-based frequency constraints within security-constrained optimisation models has been investigated in established works [14–16,37]. A Piece-Wise Linearisation (PWL) technique in [22] is applied to linearise a modified

frequency response model in multi-machine systems dominated by conventional synchronous machine generators. While this approach was demonstrated to be effective in synchronous machine-dominated systems, systems with high CIG penetration, where the frequency dynamics are more complex, were not tested. A linearisation method while considering CIGs was implemented by [37–39]. The method uses the linearised and modified COI as a constraint to consider the influence of CIGs in the network. In [14, 15], it was demonstrated that the use of the COI, instead of the actual locational frequency responses, is not always reliable and may, at times, lead to economically inefficient optimisation decisions. However, capturing in detail the locational frequency responses of the network involves solving high-order differential equations (such as employing the 6th-order swing equation) [14, 40], which are analytically hard to derive. Consequently, numerical approaches (conservative regression) were proposed to describe the frequency responses in an  $N$  region network, using the COI plus an additional term to describe regional frequency differences —two regions demonstrated in this paper. However, while this is better than the system-wide COI approach, the definition of such regions remains a challenge, as they are operating condition (OC)-specific [41]. Furthermore, using conservative regression to model locational responses may lead to tighter frequency-security regions, especially if the security boundary is highly nonlinear, potentially leading to convergence issues or over-securing the system.

In [40], the detailed frequency behaviour of each SG in the system is captured using the Trapezoidal rule, a numerical integration method applied to convert the DAEs into numerically equivalent algebraic equations. These algebraic equations are then implemented within the optimisation model to capture the locational frequency characteristics of the system. However, solving DAEs in an optimisation model substantially increases the computational requirements, particularly due to the large set of algebraic equations introduced as constraints. As each additional SG increases the number of DAEs proportionally, so does the computational time, potentially leading to scalability issues in real-world, large-scale power systems [40]. A similar approach is proposed by [16] in a security-constrained unit commitment problem. The work introduces a function that models locational frequency dynamics in the network, parameterized by

the Frequency Forming Coefficient (FFC) and Synchronising Power Coefficient (SPC). This function is expressed as the sum of the product of FFC and SPC, divided by the generator's inertia contribution. While both coefficients depend on the system's topology and operational variables, FFC primarily reflects the contribution of individual generator dynamics, i.e., rotor speeds to the frequency at a specific bus. In contrast, SPC is calculated based on the electrical distances between the SGs and buses.

To address the challenges associated with analytical approaches, ML techniques are applied for frequency stability assessment and control in [24]. A neural network is integrated with model-based COI frequency estimations to enhance accuracy. The neural network is employed as an error correction for the COI-based frequency estimations to account for the growing complexities of frequency dynamics. In [42], an ML-based approach, without relying on COI-based frequency estimations, is also applied for adaptive estimation of the maximum frequency deviation, steady-state frequency and time at maximum frequency deviation. A model inheritance technique is adopted for real-time updates of the prediction model. The proposed approach provides accurate transient frequency predictions, enabling the timely application of control measures. Furthermore, embedding ML models as constraints within optimisation models in a preventive security-constrained optimal dispatch has also been investigated [33, 43, 44]. This approach is more computationally efficient compared to optimisation models constrained by detailed analytical formulations of the network, and more accurate when compared to COI-based optimisation models. Thus, the ML approach strikes a balance between these two methods. An iterative approach to incorporate neural network-based stability constraints within the optimisation model is proposed in [43]. In [44], security-constrained unit commitment using neural networks to ensure frequency stability is proposed. A linearised neural network, incorporated as optimisation constraints, is used for transient stability in [33]. However, explicitly accounting for detailed frequency dynamics in security-constrained optimisation models remains limited; as a result, the challenge of securing against local frequency stability in converter-dominated systems has not been fully addressed.

## 2.4 Corrective Power System Frequency Stability Management Using Advanced Frequency Control Techniques

The unprecedented transition of modern power systems from being SG-dominated to heterogeneous systems with high CIG integration introduces significant challenges in power balancing, which is essential for maintaining a steady system frequency. Since CIGs are often weather-dependent and non-dispatchable, balancing tasks have become increasingly complex.

The Under-Frequency Load Shedding (UFLS) scheme is a widespread emergency frequency control strategy that involves automated load shedding to maintain power balance and frequency stability. Its operation is based on the violation of pre-set frequency thresholds, triggering pre-defined shedding of a certain load percentage at various steps [45]. However, the assumptions and operational principles underlying this technique are increasingly being challenged in many ways amidst the increasing adoption of CIGs. For instance, relying on pre-set thresholds assumes relatively predictable load and generation characteristics, which is no longer the case considering that most CIGs are stochastic, whose generation capacity cannot be determined with full certainty. Moreover, the rapidly varying load further introduces more uncertainties, challenging the reliability of such pre-set load-shedding percentages. Additionally, bi-directional power flows in CIG-dominated systems can lead to unintentional shedding of the downstream generation, exacerbating power mismatches. Furthermore, the operation of UFLS neglects to account for the possibility of frequency services offered by controllable loads [46, 47].

Adaptive UFLS (AUFLS) has been proposed in the literature to address the limitations of conventional UFLS [45, 48–50]. This technique is designed to use real-time system data to provide a load-shedding strategy or scheme in the network, rather than relying on pre-set static settings, which might otherwise be detrimental towards stability under previously unconsidered scenarios. An AUFLS scheme is designed using a Koopman Model Predictive Control (KMPC) to handle local frequency variations by

availing the benefits of collective balancing from large clusters of flexible loads in the network by [48]. KMPC is a data-driven control strategy that uses Koopman operator theory to transform nonlinear system dynamics into a linear representation in a higher-dimensional space. This enables the use of efficient linear MPC techniques on systems that are inherently nonlinear, hence preserving convexity in optimisation. By arming each controllable load with a local frequency measurement device, power modulation of the load is achieved, ensuring the system's frequency remains within acceptable limits. The method, however, requires knowledge of the minimum amount of load at which the system can safely operate following any extreme event. In [49], an AUFLS is implemented by numerically developing a more accurate power deficit derivation model consisting of two main parts: i) the power deficit using the COI swing equation plus effects of system voltage-dependent loads, and ii) considering the effect of synthetic inertia, and contribution from inverter-based resources. While model-driven approaches perform better than UFLS due to their capability to offer real-time calculations of the load shedding strategy, their accuracy is increasingly challenged by the complexity and uncertainty of modern power systems, including the dynamic CIG responses —leading to excessive simplifications [51]. Moreover, the complexity of weather-dependent CIGs complicates the uncertainties of these approaches, and accurately modelling all the dynamics at play in real-time would require significant computational resources. Consequently, accurate analytical AUFLS becomes difficult to implement, particularly where real-time or close-to-real-time applications are concerned.

To overcome these limitations, data-driven approaches, mainly Reinforcement Learning (RL)-based methods, have demonstrated potential for AUFLS applications. Unlike analytical approaches, RL does not rely on solving the network's DAEs to establish an optimal policy from observed or simulation data during training. Instead, through continuous interactions with the environment, i.e., the power system simulator, RL agents can capture even complex frequency dynamics, without relying on the explicit formulation of system equations. An RL approach with continuous load-shedding action space for voltage control using the Deep Deterministic Policy Gradient algorithm (DDPG) is proposed in [52]. [53] adopts an emergency frequency control using the same

algorithm to restore system frequency in the event-driven load-shedding mismatch scenario, considering multiple faults. A double-Q learning approach is applied by [54] in an islanded micro-grid for recovering frequency based on ON and OFF actions of different categories of loads while considering load importance. Despite achieving high accuracy in these applications, training such models requires excessive sampling through a trial-and-error learning process, resulting in inefficient training. Consequently, this introduces significant computational inefficiency and suboptimal performance. This is particularly a challenge for applications in large-scale systems with high dimensionality. [55] attempts to address this by adopting a transfer learning approach —where an emulator Multi-Layer Perceptron (MLP) model was pre-trained on a frequency control database based on a fine-tuned Proportional Integral Derivative (PID) controller —resulting in faster and more stable training. Additionally, [56] addresses the problem using a multi-agent RL-based approach governed by data-enabled predictive control (DeePC) guided policy search. In this technique, the control actions of the RL agent are guided by the historical control data of the dynamical system, and there is no explicit system modelling. The actions combine the RL policy and the data-enabled predictive control algorithm, adjusted by some weighting parameters, which enhances the sampling efficiency and solution optimality of traditional RL. However, in all these methods, it is observed that the physics governing the dynamic system is not considered, making them heavily reliant on extensive datasets. In addition, without guidance from the system’s underlying physics, the agent is prone to converging on suboptimal solutions and requires longer training times, as it may spend a considerable amount of time exploring non-essential actions.

## 2.5 Discussion of Current Challenges

From the literature review presented above, it is noted that the challenges associated with conventional approaches stem from the increased complexity of the system and the high computational requirements of highly accurate methods, such as TDS (a phasor domain simulation environment) or detailed analytical methods (mathematical equations to estimate system behaviour). The growing complexity of frequency dynamics

in modern power systems, driven by increased CIG integration, remains a challenge yet to be fully addressed. This is because most of the proposed approaches are analytical and rely on solving DAEs to make their predictions. This makes it difficult to balance complexity and computational efficiency, hindering their practical applications. This thesis addresses these challenges by taking a different direction, employing ML-based techniques as a suitable alternative. Unlike the existing analytical approaches, ML models are data-driven techniques that rely on training datasets, rather than explicit solutions from system equations, to establish highly accurate complex mapping functions between inputs to outputs. These models do not require solving DAE to make predictions, making them more computationally efficient, regardless of the system's scale. In contrast, as the system scale increases, such as through the introduction of an SG, the number of corresponding DAEs to be solved also increases significantly, thereby intensifying computational demands. This explains why ML methods have gained traction in power systems by capturing the level of detail needed while being computationally fast, making them suitable for close-to-real-time applications. Accordingly, this thesis proposes an ML-based approach to enhance situational awareness in power systems and improve the efficiency of security-constrained optimisation models, all while introducing minimal additional complexity —an area where traditional methods fall short.

It is worth noting that despite the notable contributions in the literature applying ML-based methods, they mostly focus on improving the accuracy of the COI model and not on capturing the emerging locational dynamics in the network. Moreover, despite the reported high accuracies, they normally employ black-box models, that is, models whose internal workings are not transparent to their users, without provisions to explain them. This makes it hard for their internal workings to be understood by potential users, hindering the necessary trust required for adoption. This thesis, therefore, addresses this dual problem by explicitly focusing on the locational frequency dynamics and implementing interpretability and/or explainability tools to explain the workings of such high-accuracy black-box models. This offers significant potential to build the necessary trust in ML models and foster adoption among power system prac-

titioners. An added advantage is that this allows operators to uncover the underlying relationships in the network —where there is a pool of numerous interacting devices —helping them pinpoint critical system variables or aspects influencing the stability boundary. As a data-driven approach, once the ML model(s) undergo a one-time offline training, they can generate predictions and explanations almost instantaneously. This significantly enhances accurate real-time system visibility and equips operators with advanced decision-making tools, enabling more timely and informed decision-making as well as control.

Furthermore, considering that real-time frequency control is an integral aspect of maintaining system stability, existing methods still present opportunities for improvement. First, analytical adaptive load-shedding schemes cannot accurately estimate the amount of load shedding required to maintain stability as the complexity grows, leading to inefficient actions. Second, while ML-based methods, specifically AUFLS RL, have been demonstrated as promising alternatives to fill this gap, they heavily rely on extensive datasets and often struggle in high-dimensional applications due to their trial-and-error training. This is especially a problem in modern power systems, which are growing in scale due to interconnections and/or the introduction of new devices. Consequently, the feasibility of such methods means that necessary assumptions are introduced. For instance, existing AUFLS RL often limits analysis to a few randomly selected loads or clusters an arbitrary number of loads to achieve collective benefits for the load-shedding strategy. This affects the quality and practical applicability of such RL policies developed. Moreover, to the best of our knowledge, no work has demonstrated how such AUFLS RL integrates with the traditional under-frequency load-shedding scheme. As such, this thesis addresses these challenges by developing a physics-informed AUFLS RL, trained to account for the physical laws and relationships governing the dynamic behaviour of the system —which traditional approaches ignore. Although it can be claimed that training datasets may capture some of these physics, their absence leads to reliance on extensive datasets, sub-optimal performance, computational inefficiency leading to long training time and difficulty in extrapolating beyond the training domain —potentially leading to detrimental actions in those sce-

## Chapter 2. Literature Review

narios. Additionally, the proposed AUFLS RL is integrated with the conventional UFLS as a two-layered protection scheme, effectively making it more robust and reliable. Therefore, the research in this thesis enhances the effectiveness of these AUFLS RL techniques by improving the optimality of solutions while enhancing computational efficiency—a crucial aspect where practical applicability in modern high-dimensional power systems is concerned.

## Chapter 3

# Enhancing Power Systems Situational Awareness for Frequency Stability

In Chapter 3, we propose an ML-based method to enhance situational awareness of power systems, focusing on locational frequency dynamics. The proposed method has the advantage of offering preferable qualities of computational speed while exhibiting high accuracy —comparable to approaches relying on much more computationally demanding iterative TDS. This contribution was presented at the 59th International Universities Power Engineering Conference (UPEC), Wales, Cardiff, UK, based on the award-winning paper [Pub. A]. As an initial step, we highlight the complexity of the underlying problem and the need for a detailed representation of locational frequency dynamics. We then quantify the potential locational frequency response variations under varying CIG penetrations, demonstrating how COI-based methods may fail to capture these variations. While COI-based simplifications are necessary due to the prohibitive system complexity and significant computational overhead associated with modelling all the dynamics in detail, they could lead to unforeseen locational frequency violations. This arises because the COI model, being a simplified representation based on the average system response, is inherently unable to capture the full extent of spatially distributed dynamic behaviours. To address this limitation, we implement and

test several ML algorithms to rapidly predict locational frequency response dynamics —striking a balance in overcoming the drawbacks of the fast and yet overly simplistic COI-based approach and the accurate and yet computationally intensive RMS-TDS. The contributions of this work are as follows:

- i. Explicit representation of the emerging locational frequency dynamics, which are increasingly critical in modern power systems with high renewable penetration, demonstrating how COI-based methods may fail to model such locational violations.
- ii. Capturing and representing these locational dynamics through fast and yet accurate ML models, while providing quantitative performance comparisons with the traditional COI-based method —overcoming the drawbacks of the overly simplistic COI-based methods and the detailed but time-consuming TDS.

### **3.1 Frequency Response Dynamics and Locational Aspects in Power Systems**

In this chapter, the foundational understanding related to power frequency dynamics is provided, building up on the content introduced in Chapter 2.1, mostly to highlight the complexity of the underlying problem and the need for a detailed representation of locational frequency dynamics.

#### **3.1.1 Representing Locational Frequency Dynamics in Power Systems**

In systems with high penetration of CIGs, frequency dynamics are becoming increasingly localised (exhibiting locations similar to Fig. 2.2). This eventually reduces the reliability of COI-based methods (described by (2.4)) for representing the system’s frequency, due to unforeseen locational violations risks. As a result, more granularity in frequency studies to account for such locational variations is required. [25,26] proposed the Frequency Divider (FD) to estimate local bus frequency, using an augmented system admittance matrix. This method operates on the principle that local frequencies

are affected by power flows and the dynamics of nearby generators. It defines a quota of active power that contributes to frequency variations of buses, weighting it by electrical distances. [12] proposed using the numerical derivative of the voltage phase angle to model locational frequency dynamics. A similar approach is proposed by [16], where a function that models locational frequency dynamics in the network is introduced. The function is parameterised by two coefficients: the Frequency Forming Coefficient (FFC) and the Synchronising Power Coefficient (SPC). The frequency at any bus in the network is estimated by summing the product of FFC and SPC across all online SGs and dividing this sum by their total inertia contribution. Both coefficients are operation condition-specific and are influenced by the system's topology. However, FFC primarily captures the dynamic contribution of individual SGs at a given bus, while SPC reflects the electrical proximity between SGs and buses, derived from network impedance characteristics. However, despite the potential of these methods, they share a common drawback: they are computationally intensive, since they rely on solving the system's DAEs to make predictions. Furthermore, their accuracy is challenged as system complexity grows, increasing the scale of assumptions required to make them hold. Such limitations reduce their suitability for applications requiring fast or real-time frequency stability assessments. Accurately accounting for the locational frequency dynamics means that given a network with  $N$  buses and  $n$  SGs, the frequency response for bus  $i$  can be estimated as described in [14, 15],  $\forall i \in \{1, \dots, N\}$ :

$$\frac{df_i(t)}{dt} = \frac{df_{\text{COI}}(t)}{dt} + A_i\omega_i(t) \quad (3.1)$$

where  $f_{\text{COI}}(t)$  is the average frequency response of the system, which is weighted by the inertia contribution of  $n$  online SGs (described by (2.4)). The factor  $A_i\omega_i(t)$  represents locational oscillations or deviations from the COI of the frequency response at every  $i^{\text{th}}$  location of the network. This factor depends on the operating condition of the system, that is, a function parametrised by variables such as inertia distribution, disturbance size and location, reserves, etc. As such, high enough deviations, captured by  $A_i\omega_i(t)$ , can significantly deteriorate the accuracy of COI-based methods (Chapter 2.1) —potentially leading to the risk of unforeseen relay activations or the need to over-secure

to avoid this from happening. However, analytically deriving  $A_i\omega_i(t)$  is very hard, as it would require finding the global minimum of a nonconvex function —justifying the continued use of the simplified COI model. Alternatively, if more details are required, such as accurately accounting for the spatial distribution of the dynamics, RMS-TDS are used despite their intensive computational requirements. Similarly, for nadir, this would again involve deducing locational  $t_{\text{Nadir},i}, \forall i \in \{1, \dots, N\}$ , which is hard to do for the same reasons analytically proved in [14]. Therefore, to address this challenge, we propose using ML-based techniques to capture such  $N$  locational frequency stability dynamics associated with the factor  $A_i\omega_i(t)$ . The ML model is trained using a detailed RMS-TDS dataset and consequently captures the detailed local frequency variations. Once trained, the model can achieve accuracy levels comparable to using RMS-TDS while requiring only a fraction of the time.

### 3.2 Machine Learning Approach for Capturing and Monitoring of Locational Frequency Dynamics

ML models are universal approximators capable of establishing accurate mappings between given inputs and outputs. In the context of power system applications, the inputs may include system state variables, while the output can be a stability classification or regression target. Given that these are data-driven approaches, this effectively creates a fast and efficient surrogate model —a simplified approximation of the complex system. Moreover, such an approach becomes even more valuable for applications in large-scale grids where real-time TDS is impractical, or when numerous scenarios must be screened. To achieve this, we use RMS-TDS data generated offline for a one-time ML model training. The RMS-TDS dataset contains the system’s operational and characteristic variables, which act as the ML model’s input features. In addition, the dataset contains frequency stability variables as the ML model’s regression targets. During training, the ML model aims to accurately learn the complex function mapping the inputs (the characteristic features) to the outputs (regression targets or the frequency dynamic response of the system). Note that while Phasor Measurement Units

(PMUs) are sensors capable of providing locational frequency measurements across the system, they simply allow for real-time monitoring of the system's current state and postmortem analysis; as such, they are incapable of predicting the post-disturbance dynamic response of the system, like the proposed method. This limitation leaves no room for preventive actions. Moreover, their high cost often limits widespread deployment, reducing system visibility. In contrast, our models can effectively predict the system's responses to disturbances.

### **3.2.1 Machine Learning Models for Locational Frequency Dynamics Prediction**

ML models are proposed and used as surrogate models to predict locational frequency stability metrics in a power system, specifically, nadir and RoCoF. To this end, four regression algorithms, namely Decision Trees (DT), K-Nearest Neighbours (KNN), Multi-Layer Perceptron (MLP), and a Hybrid model, an MLP algorithm whose vector of input features consists of both the variables from the power system, and COI-based frequency response predictions, are trained and tested to provide a comparison against the COI-based method.

#### **i) Decision Trees (DT)**

Decision Trees (DT) are among the supervised learning classes used for classification (predicting labels) and regression (predicting continuous values). The tree is constructed after splitting the dataset, starting from the root node (source set) into subsets (known as successor children) based on the predefined rules derived from the available features [57,58]. This process is repeated recursively on each derived subset until further splitting no longer improves the predictive power. The leaves of the tree represent the final predicted labels or regression values, while the branches are the decision rules leading to these predictions [57].

#### **ii) K-Nearest Neighbours (KNN) Regression**

KNN algorithm makes a prediction using proximity, operating under the assumption that similar points are found near one another [57,59]. Given a query point, the model identifies a predefined number of nearest neighbours from the training dataset and

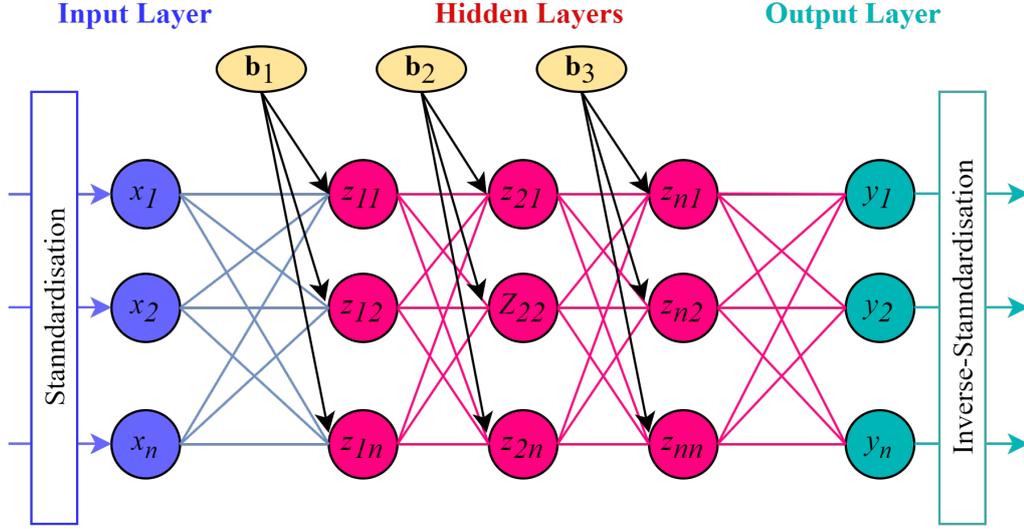


Figure 3.1: Artificial Neural Network (ANN) architecture

determines the predicted value based on their votes, weighted by a specific metric [58]. For example, in *uniform voting*, all points in each neighbourhood are weighted equally, whereas in *distance-weighted voting*, closer neighbours have a greater influence. We use the Euclidean distance as defined by (3.2) to determine proximity. Given two feature vectors  $\mathbf{x}_{\text{test}}$  (a vector from the testing dataset) and  $\mathbf{x}_{\text{train}}$  (a vector from the training dataset), the distance between a query point  $x_i$  and a training point  $x_k$  in a  $D$ -dimensional space is computed as follows:

$$d(x_i, x_k) \stackrel{\text{def}}{=} \sqrt{\sum_{j=1}^D (x_i^{(j)} - x_k^{(j)})^2} \quad (3.2)$$

### iii Multilayer Perceptron (MLP)

MLP is a class of feed-forward Artificial Neural Network (ANN) that contains at least three layers, that is, the input layer, the hidden layer with  $\sigma_v$  neurons each, and the output layer, as shown in Fig. 3.1. The input layer receives the input feature vector from the dataset, while the hidden layer performs computation to optimise the weights and biases necessary to minimise the model prediction errors. The output layer is the final layer that contains the model prediction, whose size depends on the number

of targets. Given an input feature vector  $\mathbf{x}$  and a  $V$  hidden layers MLP model with  $\mathbf{W}_v$  weights matrix of size  $(\sigma_v, \sigma_{v+1})$ , and  $\mathbf{b}_v$  bias vector of size  $(\sigma_v, 1)$ , the model's prediction is computed using (3.3-5.6) as follows:

$$\mathbf{z}_1 = \mathbf{W}_1^T \mathbf{x} + \mathbf{b}_1 \quad (3.3)$$

$$\mathbf{z}_{v+1} = \mathbf{W}_{v+1}^T \mathbf{z}_v + \mathbf{b}_{v+1}, \quad \forall v = [1, \dots, V-1] \quad (3.4)$$

where  $\mathbf{z}_1$  is the initial output vector of layer one without any transformations. In the subsequent  $v$  layers, activation functions,  $\Theta_v$ , are applied to introduce non-linearities, enabling the model to learn complex nonlinear relationships in the data. The transformed output of layer  $v$  is then passed to the subsequent layer  $v+n$ ,  $\forall n = [1, \dots, V-1]$ , that is  $\hat{\mathbf{z}}_v = \Theta(\mathbf{z}_{v+1})$ . The ReLU activation function expressed as  $\max(0, \mathbf{z}_v)$ , is adopted as it can overcome numerical issues associated with other activation functions. The predicted output vector of the model,  $\mathbf{y}$ , is then given as:

$$\mathbf{y} = \mathbf{W}_{V+1}^T \hat{\mathbf{z}}_V + \mathbf{b}_{V+1} \quad (3.5)$$

The process of optimising the model parameters involves the initial passing of the input data from the input layer to the output layer, with randomised weights and biases, a step known as forward propagation. Thereafter, back-propagation, which is the reverse process by passing the dataset from the output layer to the input layer, is applied to compute the error gradient,  $e$ , concerning the initially randomised parameters [57, 58, 60, 61]. The weights,  $w_i$ , and biases,  $b_i$ , are optimised or updated at a specified learning rate,  $\eta$ , until the overall error is minimised as follows (3.6)-(3.7):

$$w_i \leftarrow w_i - \eta \nabla e, \quad w_i \in \mathbf{W} \quad (3.6)$$

$$b_i \leftarrow b_i - \eta \nabla e, \quad b_i \in \mathbf{b} \quad (3.7)$$

Finally, these steps are repeated several times over epochs, that is, one complete

pass through the entire training dataset, to establish the best model parameters.

#### iv) Hybrid Model

The hybrid model, proposed in this work, is shown in Fig. 3.2. It integrates prediction based on COI-based and ML-based prediction methods to benefit from the strengths of both approaches. COI-based prediction, calculated using Equation (2.4), preserves key physical aspects governing the dynamic system, which a purely data-driven ML model may miss. Meanwhile, the ML model acts as an error corrector, accounting for the locational frequency dynamics to enhance overall accuracy. Specifically, this is achieved by transforming the COI-based prediction (the system's weighted frequency response) into  $N$ -locational frequency response predictions. Thus, the input vector of this model,  $\mathbf{x} = [x_1, \dots, x_n]^T$ , remains the same as the other ML models used in this work, with the addition of the COI-based predictions. A similar approach was adopted in [24, 32] using the Extreme Learning Machine (ELM) algorithm, where the data-driven model corrected the prediction errors of the COI-transfer function for control. However, unlike the previous work, our approach explicitly accounts for locational frequency dynamics, effectively enhancing the situational awareness of the system.

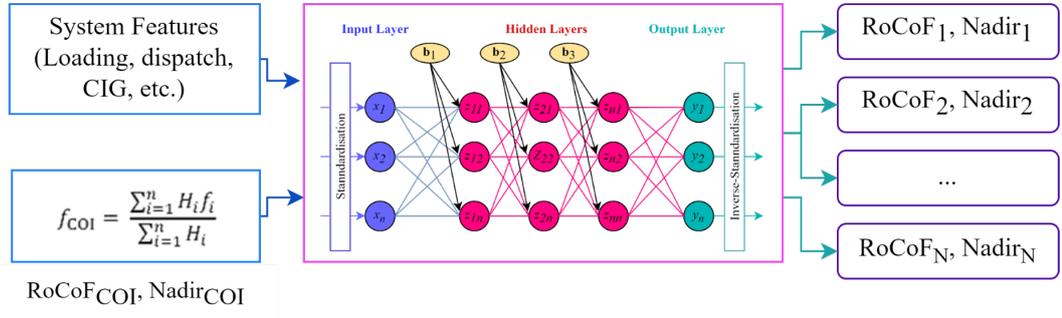


Figure 3.2: The Hybrid integrated COI-MLP Model architecture

Therefore, the final prediction of the Hybrid model, accounting for the locational frequency dynamics,  $\text{RoCoF}$ ,  $\hat{f}_{\text{inst},i}$ , and frequency nadir,  $\hat{f}_{\text{Nadir},i}$ , across  $N$  locations in the network is as follows:

$$\begin{bmatrix} \hat{f}_{\text{inst},i} \\ \hat{f}_{\text{Nadir},i} \end{bmatrix} = \begin{bmatrix} \dot{f}_{\text{COI,inst}} \\ f_{\text{COI,Nadir}} \end{bmatrix} + \begin{bmatrix} e_{1,i} \\ e_{2,i} \end{bmatrix} \quad (3.8)$$

where  $\dot{f}_{\text{COI,inst}}$  and  $f_{\text{COI,Nadir}}$  are the COI-based frequency RoCoF and nadir predictions respectively. The error(s) or individual locational deviations from the COI model,  $e_{1,i}$  and  $e_{2,i}$  for RoCoF and nadir, respectively, may be ascribed to ‘hard-to-model’ aspects of the system such as disturbance location, the impact of reactive power injections, and network topology, among others as described earlier in this chapter.

### 3.3 Methodology for Capturing and Representing Locational Frequency Dynamics

#### 3.3.1 Dataset Generation and Dynamic Simulations

The generation of scenarios or OCs for the ML frequency stability database involves a variation of three factors: the number of SGs connected, CIG output, and the system demand. Each SG is considered to comprise four equal-sized units,  $u \in (1, 2, 3, 4)$ , and its displacement by CIG output is in four stages. The new rating of the SG, i.e.,  $\text{SG}_{\text{MVA}_{\text{new}}}$ , is based on the number of remaining units and is then rated to  $\text{SG}_{\text{MVA}_{\text{old}}}$  as represented by (6.15). Accordingly, increasing generation capacity corresponds to bringing more units online rather than resizing individual units. This, in practice, may reflect the commitment of additional generators during operation. Similarly, the CIG output is scaled inversely based on  $u$  and  $\text{SG}_{\text{MVA}_{\text{old}}}$ . Variables  $r \in (1, 1.4)$  (representing low and high CIG penetration scenarios, respectively) and  $s \in (-0.05, 0.05)$ , which are sampled uniformly, are introduced to decouple the direct relationship between CIG penetration and SG displacement as follows [3].

$$SG_{MVA,new} = u(SG_{MVA,old}/4), \quad u = [1, \dots, 4] \quad (3.9)$$

$$CIG_{MVA} = r \left( \frac{(5-u)SG_{MVA,old}}{4} + s(SG_{MVA,old}) \right) \quad (3.10)$$

Consequently, based on the number of SGs in the system displaced by the CIGs, the total CIG penetration level is up to 40% of the overall system generation. The system demand ranges from 0.6 to 1.025 p.u. in steps of 0.025 p.u. All CIGs operate at fixed active power dispatch based on (6.16). The AC-OPF model is used for generation dispatch. This approach ensures that the generated scenarios are operationally feasible and consistent with the actual network's limits, including the voltages, which better reflects the practice by system operators. Moreover, the AC-OPF-based sampling helps us prevent including physically infeasible scenarios that could degrade the performance and generalisability of our trained ML model. Finally, OPF reflects market-based dispatch that minimises cost, leading to more realistic scenarios in this aspect too. To initialise the initial operating conditions of the power system in our RMS-TDS environment, the Newton-Raphson load flow is utilised. Thereafter, we apply  $N - 1$  contingency of the highest generation in the low inertia area (Area 2), SG-6 contributing between 20% and 50% of the total area generation. Note that our method is not contingency-specific; as such, any other contingency can be adopted without loss of generality. The system stability limits are considered to be 59.3 Hz and -0.5 Hz/s for nadir and RoCoF, respectively [12]. The simulation window is 60 seconds (see Fig. 2.2) to capture both the system's inertia response and primary responses to disturbances. RoCoF is calculated as an average frequency response change over a typical measurement window of 500 milliseconds post-disturbance to mitigate simulation or measurement noise errors. Frequency response predictions by the COI model are calculated using (2.4).

### 3.3.2 Data Pre-processing and Training

From the generated dataset of operational scenarios, a 70-30% train-test split is adopted, utilising the 5-fold Cross-Validation technique—which randomly splits the training set

into  $k$ -folds, whereby each  $k - 1$  set is used for training. Feature standardisation of the dataset is achieved by using the StandardScaler [57], which is a pre-processing technique that rescales every feature so that the mean is zero with a unit variance. Following training, the stored mean and standard deviation, through *Inverse-Transform*, are used to rescale the data for testing and evaluation. The scaling process is represented as:

$$z_{i,j} = s_i \frac{x_{i,j} - \mu_i}{\sigma_i}, \quad \forall i = [1, \dots, 2 \times N + M] \quad (3.11)$$

where  $z_{i,j}$  is the standard score,  $x_{i,j}$  is the input value,  $\mu_i$  is the mean,  $\sigma_i$  is the standard deviation and  $s_i$  is the scaling factor of the  $j^{\text{th}}$  sample.  $N$  represents the number of buses—multiplied by two metrics (RoCoF and nadir) per bus, while  $M$  are the operational and physical variables of the network, as input features. Therefore, given a trained ML model,  $\rho$ , and a feature input vector  $\mathbf{x}$  of size  $(M \times 1)$ , the locational frequency dynamics of the system are captured as follows:

$$[\hat{f}_{\text{inst},i}, \hat{f}_{\text{Nadir},i}] = \rho([l, g, d, h, o]), \forall i = [1, \dots, N] \quad (3.12)$$

where, in a network with  $N$  (monitored) buses,  $\hat{f}_{\text{inst},i}$  and  $\hat{f}_{\text{Nadir},i}$  are the predicted RoCoF and nadir at the  $i^{\text{th}}$  bus, respectively. The feature input vector of the model,  $\mathbf{x}$ , consists of  $l$ , the system loading;  $g$ , the CIG active power output;  $d$ , the active power dispatch of SGs;  $h$ , the inertia of SGs in MVA.s; and  $o$ , the generator MVA ratings. To identify a suitable combination of hyperparameters giving the best performance from the ML model's training, we apply the sklearn-GridSearchCV. This is a tool that automates the process of finding the best combination of hyperparameters for a given ML model by exhaustively searching over a user-defined parameter grid [57,58]. Sequel to model training, the stored mean and standard deviation, through *Inverse-Transform*, are used to rescale the data back to their original values for testing and evaluation.

### 3.3.3 Model Evaluation

The predictive performance of all the models, including the COI-based method, is evaluated by two main metrics: the Mean Absolute Error (MAE) (3.13) and the Root

Mean Squared Error (RMSE) (3.14), which are calculated by:

$$\text{MAE}_j = \frac{1}{N_e} \sum_{i=1}^{N_e} |y_{i,j} - \hat{y}_{i,j}|, \forall j = [1, \dots, N] \quad (3.13)$$

$$\text{RMSE}_j = \sqrt{\frac{\sum_{i=1}^{N_e} (y_{i,j} - \hat{y}_{i,j})^2}{N_e}}, \forall j = [1, \dots, N] \quad (3.14)$$

where  $N_e$  is the total number of OCSs,  $y_{i,j}$  and  $\hat{y}_{i,j}$  are the true and predicted regression targets (RoCoF and/or nadir in this case) at the  $j$ -th location of the network, respectively, while  $N$  is the total number of monitored locations. The MAE is used to assess how the model would perform on critical cases, that is, those close to the stability margin. Errors in critical cases impact key decisions, such as ancillary service procurement, such as inertia, multi-speed frequency response, load damping, etc. [15]. An overestimate of RoCoF and/or nadir may result in system operators over-procuring costly ancillary services. Conversely, an underestimate may result in a SO procuring insufficient ancillary services, leaving the system potentially vulnerable. For locational predictions, we focus on the maximum, mean and minimum errors of the inference model as follows:

$$\text{Max. RMSE} = \max_{j=1, \dots, N} \text{RMSE}_j \quad (3.15)$$

$$\text{Mean RMSE} = \frac{1}{N} \sum_{j=1}^N \text{RMSE}_j \quad (3.16)$$

$$\text{Min. RMSE} = \min_{j=1, \dots, N} \text{RMSE}_j \quad (3.17)$$

### 3.4 Case Studies and Results

This chapter presents the results of the proposed method, along with a quantitative comparison with the conventional COI-based approach.

### 3.4.1 Test Network and Case study

In this study, a modified version of the IEEE 39-bus network, as shown in Fig. 3.3, is used [3]. We modify the network by subdividing it into two areas, Area 1 and Area 2, whereby Area 2 is a low-inertia region due to the penetration of CIG. This configuration has been intentionally selected to create two distinct regions: high and low inertia, allowing a clear demonstration of the impact of high CIG integration on locational frequency dynamics. This is similar to the GB network, with England being the high-inertia region and Scotland the low-inertia region [15]. CIG is modelled using the Western Electricity Coordinating Council (WECC) Type IV Wind Turbine Generator (WTG) control model [62], which is integrated into the grid through a fully rated converter at Bus 16. We determine the number of online SG units using (6.15) and use the AC-OPF for generator dispatch as described in Chapter 3.3.1 through MATPOWER. The polynomial cost function coefficients of the SGs are as shown in Table 3.1, while the cost of CIGs is considered to be zero. RMS-TDS are conducted in DIgSILENT PowerFactory, and the instantaneous CIG penetration was between 50 MW to 1200 MW, representing around 40% and 95% of system and area generation, respectively. The applied contingency is  $N - 1$  of the highest generation in the low inertia Area 2 (G 06), contributing between 20% and 50% of the total area generation. This is done to effectively demonstrate the locational frequency stability aspects due to the high local penetration of CIGs [14]. The whole dataset consists of 1,728 samples from varying system loading, the number of online SGs units in Area 2, and the CIG output as described in Chapter 3.3.1. The stability limits are considered to be 59.3 Hz and  $-0.5$  Hz/s for nadir and RoCoF, respectively [12]. All RMS-TDS were carried out on an 11<sup>th</sup> Gen Intel (R) Core (TM) i7-11700 @ 2.50 GHz with 16 GB installed RAM, which took about 80 minutes.

### 3.4.2 Locational Frequency Response (RoCoF) Variations

To illustrate the inherent risks associated with COI-based methods —which cannot capture spatially distributed dynamics —RMS-TDS are conducted on all the solutions generated by the AC-OPF. Shown in Fig. 3.4 is RoCoF from sample locations across

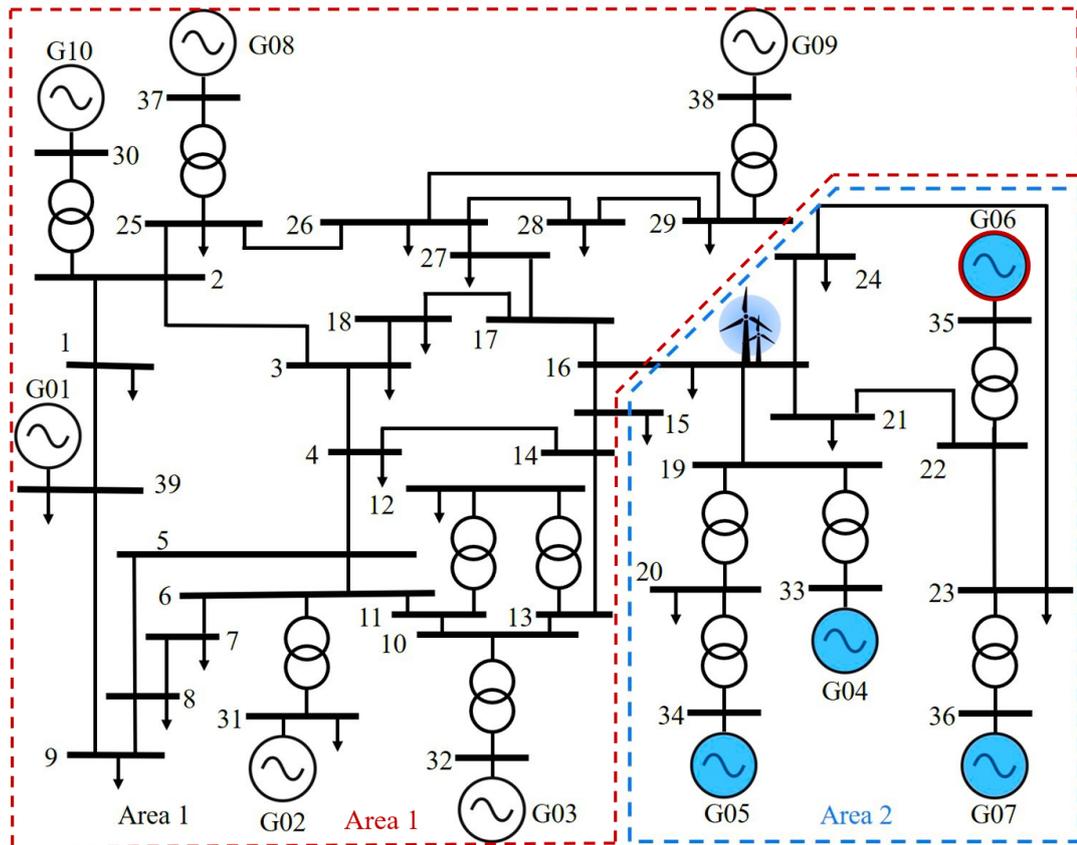


Figure 3.3: The modified IEEE 39-Bus network highlighting Area 2 where CIG is connected at Bus 16 and location of the disturbance SG-6 (G06)

Table 3.1: Cost function coefficient values for Synchronous Generator (SG) polynomial cost-function [3, 6]

	G 01	G 02	G 03	G 04	G 05	G 06	G 07	G 08	G 09	G 10	CIG
$C_2$	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0
$C_1$	2.5	6.9	4.7	2.8	3.6	2.8	3.7	4.8	3.7	3.9	0
$C_0$	0	0	0	0	0	0	0	0	0	0	0

the network, and the COI-based RoCoF predictions, plotted against CIG penetration between 50 MW to 1,200 MW. This represents a maximum CIG penetration of up to 40% of the entire system generation. While it is expected for the system to exhibit some locational dynamics variations, the increasing intensity of such variations with respect to the increasing CIG penetration is obvious. Specifically, for CIG penetration less than 250 MW, locational frequency variations are not significant enough to cross the -0.5 Hz/s threshold. In such cases, the COI-based approach does not bear any risks of unforeseen frequency violations. In contrast, for penetration levels above 250 MW, the variations become significant enough to allow some locations to cross the -0.5 Hz/s threshold. As a result, due to the widening margin between COI-based predictions and the spatially distributed extremes, the frequency violations are undetectable to the COI —up to 70 cases observed in this case. Consequently, this could lead to unforeseen activation of locational relays, potentially causing widespread blackouts. This highlights the risks; the COI-based method is becoming increasingly unreliable as more CIGs are integrated into power systems, necessitating new approaches to ensure secure operation.

### 3.4.3 Quantification of Locational Frequency Response (RoCoF) Variations

The locational frequency response variations are analysed and quantified to underscore the impact of significant CIG penetrations. We focus on the RoCoF metric as this is the most vulnerable metric to locational variations [12, 14]. Given a set of buses  $\xi^{\text{Area 1}}$ ,  $\xi^{\text{Area 2}}$  in Area 1 and Area 2 respectively (see Fig. 3.3), the maximum RoCoF variation,  $\mathbb{D}_{\text{RoCoF}}$ , is given by (3.18) as follows:

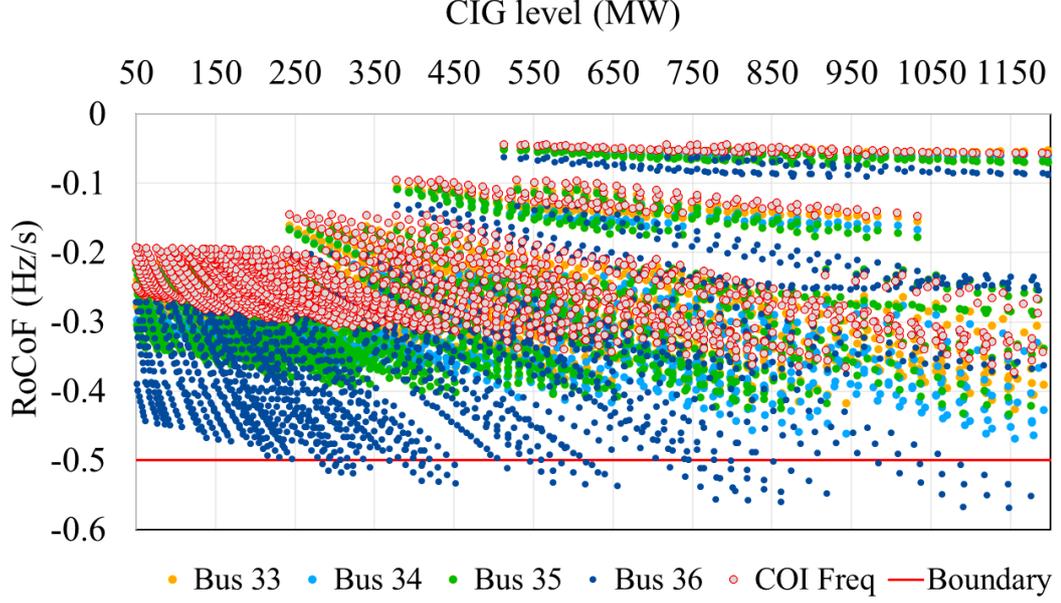


Figure 3.4: Locational RoCoFs at selected buses in the network compared with COI-based predictions

$$\mathbb{D}_{\text{RoCoF}} = \max(|\hat{f}_{\text{inst},i} - \hat{f}_{\text{inst},j}|), \forall i \in \xi^{\text{Area 1}}, \forall j \in \xi^{\text{Area 2}}, \xi^{\text{Area 1}} \cap \xi^{\text{Area 2}} = \{\} \quad (3.18)$$

The maximum difference/variations seen in the network based on  $\mathbb{D}_{\text{RoCoF}}$  are plotted in Fig. 3.5 against the loading of the system, with respect to the inertia (in MVA.s) in Area 2. Two scenarios, that is, the ‘*No CIG*’ scenario with 0% CIG penetration and the ‘*Area 2 CIG Location*’ where there is a varying penetration level of the CIG (up to 40% of system generation) located at Bus 16 in Area 2, are presented. Without the CIGs, there is a maximum RoCoF variation from 0.1450 Hz/s to 0.2505 Hz/s observed in the ‘*No CIG*’ scenario between 0.6 p.u. and 1.025 p.u. system loading, respectively. In contrast, the *Area 2 CIG Location* scenario increases these variations to range between 0.1550 Hz/s and 0.3487 Hz/s between 0.6 p.u. to 1.025 p.u. loading conditions, respectively. Specifically, locational RoCoF variations grow at every system loading as the inertia in Area 2 reduces, from high to medium levels —due to CIGs displacing SGs. When inertia is highest in the system, around the ‘*No CIG*’ scenario

where all SGs are connected,  $\mathbb{D}_{\text{RoCoF}}$ , is not as significant as when the inertia is lower at a medium level, as highlighted in green.

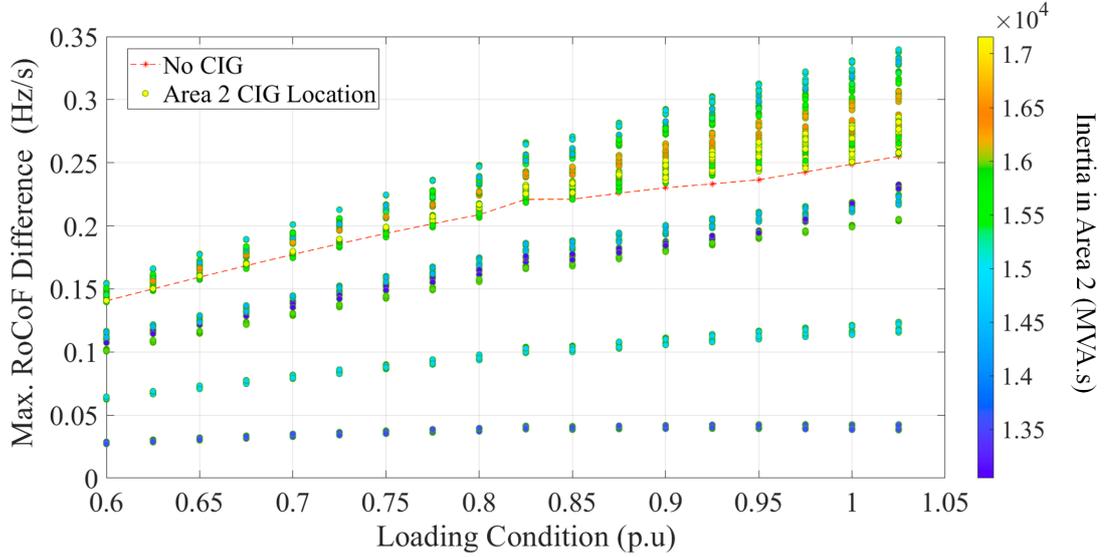


Figure 3.5: Maximum locational RoCoF variations following the outage of SG-6

It is worth noting in Fig. 3.5 that the varying disturbance size, fixed at SG-6, is responsible for the reduced locational RoCoF variations of some OCs in the *Area 2 CIG Location* scenario. This variation results from changes in the number of online SG units, including at this power plant, as stated by (6.15) in Chapter 3.3.1. Consequently, a smaller disturbance leads to lower system RoCoFs, causing a corresponding reduction in locational variations. In contrast, the *No CIG* scenario involves no displacement of SGs, retaining the highest number of online SGs, and therefore, the highest possible disturbance—and highest possible RoCoF variations. Similarly, changes in the disturbance size are also influenced by system loading. For instance, at a higher loading of 1.025 p.u., machines experience greater loading, leading to larger disturbances with more pronounced system impacts, as shown by the figure.

#### 3.4.4 Predictive Accuracy of the Machine Learning Models

The predictive accuracy of ML models and the COI-based method is evaluated in this chapter for both RoCoF and nadir, based on the evaluation metrics given in Chapter

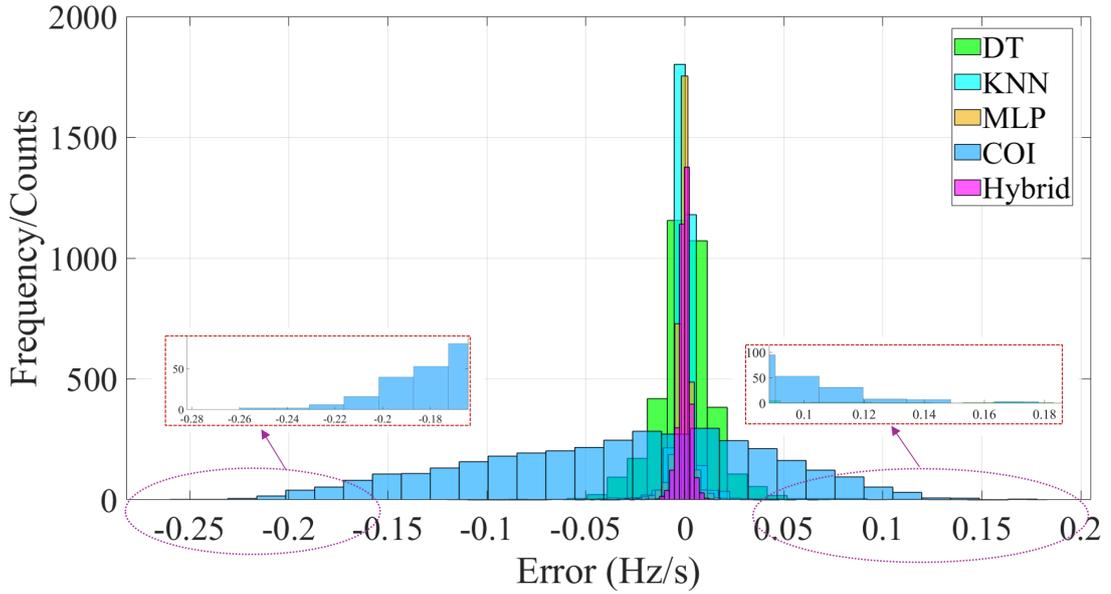


Figure 3.6: A distribution of locational RoCoF prediction errors by different models

3.3.3. RoCoF errors are given in Fig. 3.6, where the COI-based method incurs the widest range of under- and over-estimations (since it is simply a weighted sum of the metric across the network), ranging between  $-0.25$  Hz/s to  $0.18$  Hz/s. The Hybrid method has the least RoCoF error variation, seconded by the MLP and then the KNN model. Similarly, the nadir errors in Fig. 3.7 show that the COI-based method also has the widest error range between  $-1$  Hz and  $1$  Hz. The Hybrid model has the narrowest error range, almost zero. This results from the additional information from COI-based predictions, which preserves key physical aspects governing the dynamic system, potentially missed by the purely data-driven ML models. Positive errors mean the model gives a false alarm, while negative errors mean the model gives a false sense of security for OCs close to stability margins. Both are undesirable as false alarms lead to excessive allocation of costly control resources, while false security risks unforeseen system blackouts. Consequently, an ideal model has the narrowest shape of distribution errors, which, according to the charts, is achieved by the Hybrid model, followed by the MLP.

The predictive accuracy is quantified in Table 3.2. The Hybrid model has the mean RoCoF MAE and RMSE of  $0.0022$  Hz/s and  $0.0031$  Hz/s, respectively, outperforming the rest of the models, seconded by MLP and then KNN. The COI-based method bears

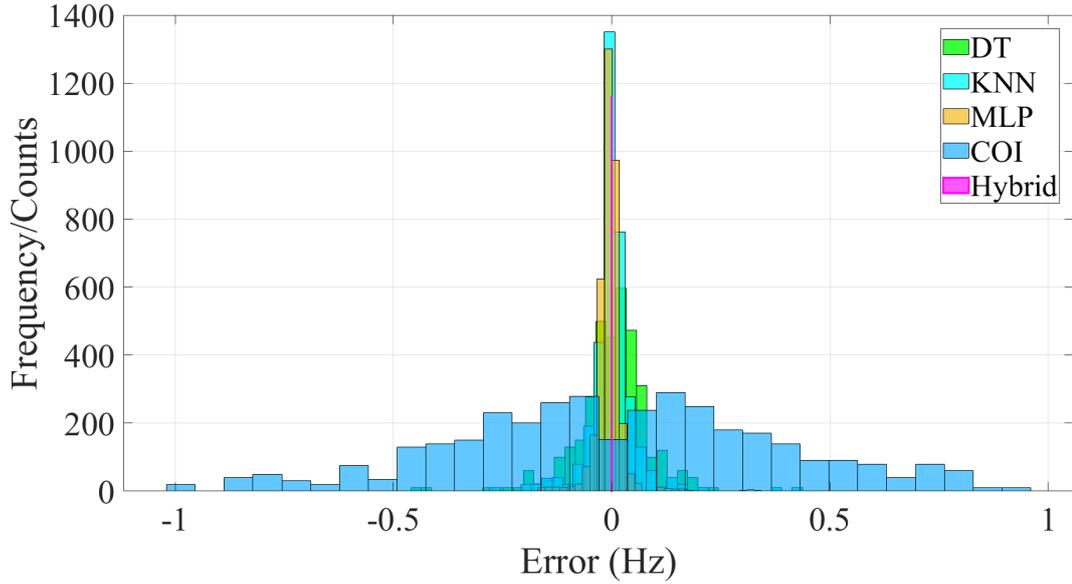


Figure 3.7: A distribution of locational nadirs prediction errors by different models

the largest mean RoCoF MAE and RMSE of 0.0631 Hz/s and 0.0587 Hz/s, respectively. The nadir errors also show that the Hybrid model consistently outperforms the rest of the ML models by bearing a mean MAE and RMSE of 0.0001 Hz and 0.0002 Hz, respectively. The COI's mean MAE remains high at 0.2980 Hz. This is because nadir is greatly influenced by larger-scale system characteristics such as load and generation balance, etc., and occurs much later after the initial condition dynamics have settled. In contrast, RoCoF is calculated right after the disturbance, before the activation of control resources necessary to maintain stability, resulting in high vulnerability to complex dynamics, which the COI struggles to model accurately. Nevertheless, while the ML models portray significant accuracy, specifically the Hybrid and MLP models, overall performance can also be improved through several ways such as implementing locational inference models, rather than an all-in-one model, as suggested in [3], and exploring a wider range of hyper-parameters, model architecture and/or algorithms, among others [61].

Table 3.2: Mean Prediction Errors for RoCoF and Frequency Nadir

Model	<i>RoCoF</i>		<i>Nadir</i>	
	MAE	RMSE	MAE	RMSE
<b>DT</b>	0.0107	0.0164	0.0584	<b>0.1387</b>
<b>KNN</b>	0.0031	0.0070	0.0321	0.0822
<b>MLP</b>	0.0030	0.0048	0.0191	0.0589
<b>COI</b>	0.0630	0.0587	0.2980	0.0343
<b>Hybrid</b>	<b>0.0022</b>	<b>0.0031</b>	<b>0.0001</b>	<b>0.0002</b>

### 3.4.5 Computational Aspects

Concerning computational requirements, inference time is a crucial aspect necessary for real-time or near real-time applications to enable sufficient time for preventive/corrective actions. Moreover, if there are numerous scenarios to be assessed, a fast-screening model can excel in situations where conventional methods make it computationally prohibitive. The proposed method employs ML to capture the system’s locational frequency dynamics, offering a computationally attractive alternative that enables near-instant predictions since solving the system’s DAE is no longer required to make predictions. We assess the prediction time of each method, including RMS-TDS used to capture network dynamics in detail, on a subset of 500 OCs. As summarised in Table 3.3, results show that an assessment of up to 300 seconds is required by RMS-TDS, while, for the whole set, 0.0808 seconds, 0.0361 seconds, 0.00203 seconds, and 0.001 seconds are required by Hybrid, MLP, KNN and DTs, respectively. The algorithms training—an offline and a once-off task—required 155.047 seconds, 74.853 seconds, 0.030 seconds, and 0.006 seconds for Hybrid, MLP, KNN, and DT, respectively. These results demonstrate how ML models can balance accuracy and computational efficiency, achieving over 3,000 times faster than the RMS-TDS in this case. This enables high-granularity frequency response assessments with minimal computational overhead.

## 3.5 Conclusion

We have demonstrated that the penetration of Converter Interfaced Generation (CIG) increases locational variations in the system frequency response due to changes in iner-

Table 3.3: Comparison of Computational Requirements for Different Models

Model	Assessment Time (s)	Training Time (s)
<b>RMS-TDS</b>	300	-
<b>DT</b>	<u><b>0.0010</b></u>	<u><b>0.006</b></u>
<b>KNN</b>	0.0020	0.030
<b>MLP</b>	0.0361	74.853
<b>Hybrid</b>	0.0808	155.047

tia distribution. Consequently, conventional Centre of Inertia (COI)-based approaches, which assume a global frequency response, fail to capture these locational variations, leading to unforeseen protection relay activations. To address this, we propose a Machine Learning (ML) approach to capture and represent the system’s locational frequency stability metrics, specifically frequency nadir and Rate of Change of Frequency (RoCoF). ML models can establish complex system relationships based on observational data more quickly and without the need for computationally intensive time domain simulations, making them advantageous for online settings. Several algorithms were tested on the modified IEEE 39-bus test network. The Hybrid approach, which integrates the COI model and a neural network, achieved the highest predictive accuracy, closely followed by the MLP model. Both models demonstrated a high degree of accuracy and computational efficiency, achieving speedups of over  $3,000\times$  faster than traditional methods. However, since both models perform well, we apply the MLP model in subsequent chapters due to its simplicity and practical considerations. Unlike the Hybrid model, which relies on COI-based calculations, it is less vulnerable to practical challenges such as sensor deployment limitations, effective inertia estimations, among others [29].

## Chapter 4

# Enhancing Understanding of Frequency Dynamics and Preventive Stability Management

Power system operators often face challenges in understanding the evolving frequency dynamics landscape driven by rising CIG penetration—even when detailed modelling methods, such as those presented in Chapter 3, are available. This, in turn, affects their ability to implement effective preventive and corrective actions to maintain system stability. We address this aspect in Chapter 4, by proposing an ML-based method that utilises the wealth of available system data to uncover underlying system relationships. Specifically, we apply explainable ML techniques that use a game-theoretic approach to quantify the contributions of system features towards influencing a given stability metric of interest—We focus on RoCoF and frequency nadir. This contribution is published in the International Journal of Electrical Power & Energy Systems (IJEPES) [**Pub. B**]. The contributions of this chapter are, therefore, summarised as follows:

- i. Enhancing operators’ understanding of what drives/affects the system frequency dynamics by uncovering potential system variables that might impact the frequency stability boundary at both local and global levels.
- ii. Introducing high-level constraints into optimisation models by formulating in-

sights from explainable ML models into actionable preventive control actions for improving system stability.

## 4.1 Explainable Machine Learning: A SHAP Value-Based Approach to Locational Frequency Stability

Modern power systems are growing in complexity due to the increasing penetration of CIGs. This technology is leading to increased locational frequency responses in power systems as discussed in Chapter 3.4.2. This is because CIGs introduce varying types of controls and system interactions different from the conventional SG technology. As a result, conventional analytical approaches used to extract insights into system behaviour struggle in terms of accuracy and computation overhead, leading to instability risks and a diminished understanding of these dynamics. To address this challenge of complexity leading to diminished system understanding, we propose enhancing the operator's understanding of the system's frequency dynamics through an ML-based approach. This approach integrates neural networks from Chapter 3.2.1 and SHAP, an additive feature attribution method that quantifies features' contribution to model predictions, thereby enabling operators to identify critical aspects of the system influencing the stability. A key contribution is that we go beyond the notion of ML as being just a powerful black-box predictor into explainability, thereby building the necessary trust towards real-world applications. The ML model is used to predict the dynamic frequency response of the system to disturbance, while SHAP uncovers the underlying system relationships as captured by the ML inference model.

The literature in Chapter 2.2 indicates that existing explainable ML methods focus on the accuracy of the black-box models used for inference, and do not go beyond to demonstrate how such insights can be validated as practical system interventions. They simply provide explanations from the explainable ML models without further validating them within the power system stability context, making their effectiveness uncertain. In contrast, we propose going beyond such explanations by: i) integrating ML and SHAP for enhanced situational awareness related to frequency stability of the system and

ii) introducing a SHAP-informed optimisation formulation. This integration enables the use of insights from the SHAP model as actionable optimisation constraints for preventive actions for enhancing system stability —or as a starting point for further system analyses. Consequently, we validate and quantify the practical effectiveness of the explainable ML models within power system operations.

## 4.2 ML Explainability for Understanding Locational Frequency Dynamics through SHapley Additive exPlanations (SHAP)

We use SHAP to enhance the understanding of frequency dynamics based on locational ML models trained to predict frequency stability metrics at each bus. This enables operators to uncover location-wise variables influencing the stability boundary, enabling more informed decisions. SHAP is a model-agnostic game-theoretic approach for explaining the output of parametric ML models (often black-box models). It seeks to provide a linear explanation based on additive feature attribution methods [63]. The explanation can be in the form of a *Local Explanation*, which is based on a single sample, or a global explanation, which is the global measure of feature effect computed by averaging *Local Explanations*. Concerning this work on frequency stability, SHAP values can be interpreted as a change in the *expected model* prediction (an average from the training dataset) of the stability metric of interest, in Hz for nadir or Hz/s for RoCoF. Therefore, the sum of all SHAP values equals the difference between the mean (*expected model*) prediction and the actual prediction of the inference model.

The process of deriving SHAP values is as follows: first, a SHAP explainer creates a ‘dummy model’ (a model without any feature inputs) whose prediction is the mean of predictions from the training dataset. This prediction is essentially known as the *Expected value*  $E[f(x_i)]$ . Given an input feature set  $F$ , the effect of each  $x_i$  input feature is computed by retraining the model on all feature subsets  $S$ , i.e.,  $S \subseteq F$ . An importance value for each feature is based on the difference in predictions between the model with and without the feature, i.e.,  $f_{S \cup \{x_i\}}$  and  $f_S$ , respectively. The effect of

each feature,  $\phi_{x_i}$ , i.e., the SHAP value, is then computed based on an average weighting of all the combinatorial differences as follows (4.1):

$$\phi_{x_i} = \sum_{S \subseteq F \setminus \{x_i\}} \frac{S!(|F| - |S| - 1)!}{|F|!} [f(S \cup \{x_i\}) - f(S)] \quad (4.1)$$

where  $S$  is the number of coalitions excluding the feature  $x_i$ ,  $F$  is the total number of coalitions, while the expressions  $f(S \cup \{x_i\})$  and  $f(S)$  are the coalitions predictions with and without the feature  $x_i$  respectively. Therefore,  $\phi_{x_i}$  is interpreted as the average marginal contribution from all coalitions that contain  $x_i$  feature, representing the effect of the feature on model predictions. The sum of all the effects of the  $M$  input features in set  $F$  is equivalent to the difference between the *Expected value*,  $E[f(x_i)]$ , and the prediction by the original ML model  $f(x_i)$ , i.e,  $f(x_i) = E[f(x_i)] + \sum_{i=1}^M \phi_{x_i}, \forall x_i \in F$ .

The process of fitting the SHAP explainer is computationally intensive and time-consuming due to the exponential increase in feature coalitions—theoretically a total of  $2^n$ . This is particularly problematic for the power system, which has numerous devices and associated monitored variables. To expedite the fitting process, some authors have proposed sampling-based approximations [64]. However, there are still some computational challenges and inefficiencies associated with these methods when conducting detailed SHAP analysis on large datasets [64]. Alternatively, explainers can adopt different algorithms, such as *KernelExplainer*, *DeepExplainer*, *TreeExplainer*, etc., which focus on the most impactful coalitions to compute SHAP values. For instance, the *DeepExplainer*, which is used in this chapter and is specific to deep neural networks, leverages neural networks (such as MLPs) to approximate  $\phi_{x_i}$ , enabling it to handle high-dimensional spaces [65]. Specifically, it uses the already trained neural network’s weights and activations to efficiently compute feature attributions. Overall, these approaches ease the computational burden imposed by the brute-force computation of SHAP values, which considers all coalitions. Nevertheless, this process is conducted offline and only once, so that thereafter, the computation of SHAP values for any previously unseen OC is done almost instantaneously.

### 4.3 Methodology for Enhancing Understanding of Frequency Dynamics

This chapter presents the methodology for understanding local frequency characteristics using ML. This allows system operators and planners to predict locational frequency metrics and identify important system variables shaping the frequency stability boundary at any location in close to real-time without requiring computationally intensive RMS-TDS. Such an approach could also be useful in planning timescales (days, weeks or months), where real-time is not as important, but when the number of scenarios becomes too large to handle. Moreover, system optimisation models can leverage simple rules extracted by insights based on SHAP values regarding important variables defining security boundaries to generate more secure and reliable solutions, particularly useful during operational and planning timescales.

#### 4.3.1 Proposed ML-based approach to understanding Local Frequency Characteristics

The proposed methodology is summarised in two stages in Fig. 4.1. The first stage is the *Offline Training* stage, which is for data generation, ML algorithms training, and the generation of the SHAP explainer. In this stage, the AC Optimal Power Flow (OPF) model is used to generate operational scenarios. The scenarios are then relayed to a dynamic simulation environment for frequency stability assessment under the largest generation outage contingency. Frequency stability metrics (i.e., nadir and RoCoF) are calculated for each location in the network to compile the frequency stability dataset. Locational ML algorithms are implemented and trained for each metric and location in a regression task. The ML models are thereafter explained with SHAP, which is a game-theoretic approach for feature attribution. In the second stage, i.e., the *Application* stage, the trained ML models are used to predict locational frequency stability metrics across the network of previously unseen scenarios, and their respective SHAP explainers are used to discover important system variables defining the stability boundaries of respective locations. Based on these insights, a formal guide is provided to the system

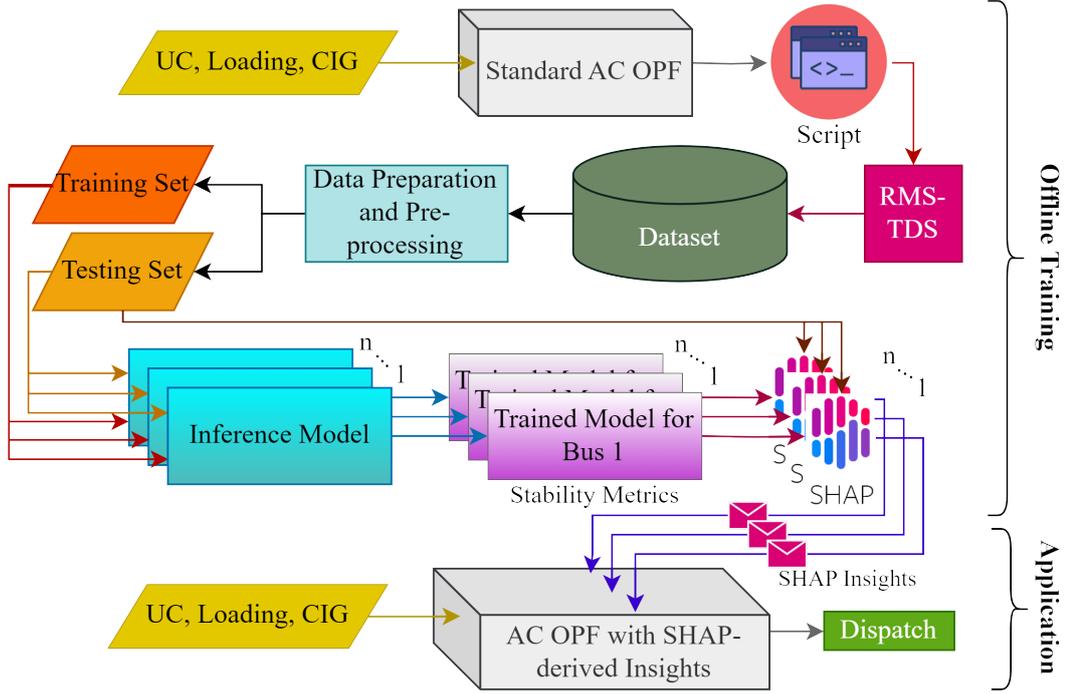


Figure 4.1: A SHAP values-based understanding of system dynamics with concise security rules

optimisation model for secure dispatch, referred to as SHAP-informed optimisation.

### 4.3.2 SHAP-Informed System Optimisation for Secure System Dispatch

This chapter proposes the use of SHAP as a means to enhance the understanding and management of locational frequency issues, which are becoming important in systems with high penetration of CIGs. By uncovering the effect of various power system variables on the predictions of locational ML models and their shaping of the frequency stability boundary, power system operators and planners, combined with domain knowledge, can minimise the risks of instability through SHAP-informed decision-making and optimisation strategies —leading to more secure and reliable system operation. To embed the SHAP-derived insights into the OPF, let  $\mathcal{L}$  be the set of loads,  $\mathcal{B}$  the set of buses, and  $\mathcal{G}$  the set of generators indexed by  $i$ . The active-power output of generator  $i$  is denoted by  $P_{G,i}$  and  $C_i(P_{G,i})$  is its generation cost. We can therefore write the OPF

problem as:

$$\min \sum_{i \in \mathcal{G}} C_i(P_{G,i}) \quad (4.2a)$$

$$\text{subject to: } g_0(x_s, u) = 0 \quad (4.2b)$$

$$h_0(x_s, u) \leq 0 \quad (4.2c)$$

$$\Phi_{x_i}(u) \leq 0 \quad (4.2d)$$

where:  $x_s$  and  $u$  are the state and control variables under steady-state conditions, respectively. (4.2b) enforces power balance, (4.2c) enforces limits in the network such as line-flows, generator output, etc., and (4.2d) is the additional inequality constraint based on the proposed SHAP-informed stability constraints—for operational and planning timescales. A more detailed formulation of the OPF problem can be found in [66].

This work leverages the actionable insights derived from SHAP values by incorporating a set of one-sided constraint inequalities,  $\Phi_{x_i}$ , within the optimisation model to provide frequency stability-related security constraints. This is achieved by identifying the most significant variable and determining the direction of its impact on the stability metric of interest. If the most significant feature has a negative direction, minimising it would result in a positive effect on the stability metric of interest (i.e., stability gain) and vice versa. Regardless of the direction, the goal is to bring a positive effect on the model prediction as shown below:

$$\bar{\phi}_{x_i} = \frac{1}{N} \sum_{j=1}^N \phi_{x_i,j}, \quad \forall x_i \in F \quad (4.3)$$

$$\phi'_{x_i} = \max(|\bar{\phi}_{x_i}|), \quad \forall x_i \in F \quad (4.4)$$

$$\Phi_{x_i} = \begin{cases} x_i > \frac{\phi'_{x_i} x_i^{\max}}{\phi_{x_i}^{\max} - \phi_{x_i}^{\min}}, & \text{if } \bar{\phi}_{x_i} > 0, \phi'_{x_i} = |\bar{\phi}_{x_i}| \\ x_i < \frac{\phi'_{x_i} x_i^{\max}}{\phi_{x_i}^{\max} - \phi_{x_i}^{\min}}, & \text{otherwise, } \phi'_{x_i} = |\bar{\phi}_{x_i}| \end{cases} \quad (4.5)$$

where the superscripts: *min* and *max*, represent the minimum and maximum value of

## Chapter 4. Enhancing Understanding of Frequency Dynamics and Preventive Stability Management

a variable (SHAP value or input feature), respectively. This set of constraints is specifically applicable to continuous variables, such as generator dispatch. We first compute the mean SHAP value for each  $i^{\text{th}}$  feature across  $N$  samples. Next, we identify the feature with the largest absolute mean SHAP value, as it is the most critical to the stability metric of interest. Based on the computed SHAP value, we then implement a feature-wise constraint: if the mean SHAP value is negative, the constraint aims to minimise the feature's value; if positive, it aims to maximise it. For generator ratings, typically determined by the number of units of a specific size connected, the optimisation process considers only the extremes as dictated by the SHAP values, i.e. minimum capacity for negative SHAP values and maximum capacity for positive SHAP values. This ensures that the constraints are applied in the appropriate direction, aligning with SHAP-derived insights to enhance system stability. This approach enables the optimisation model to dispatch generators more securely by pushing the frequency stability metric of interest farther away from the stability boundary. For instance, actionable insights from SHAP values can be used to manage and improve the stability of vulnerable locations by identifying important variables shaping the stability boundary of those locations.

The domain knowledge and data-informed decisions can be integrated within the frequency control resource allocation framework to enhance efficiency. RMS-TDS simulations on the solutions from the SHAP-informed optimisation model are applied to evaluate the constraint's impact. The constraints in this study implement the first and most significant actionable SHAP insight to i) avoid potential optimisation convergence issues and ii) to accurately evaluate the effectiveness of individual insights. Note that while analytical stability-constrained optimisation can also improve the system's stability, it often fails to account for the locational frequency dynamics due to increased complexity and/or significant computational resources required, as noted by [40]. This can result in either over-securing the system, which comes at a cost or exposing the system to risks of locational frequency violations. In contrast, our data-driven method balances accuracy and computational efficiency while identifying critical features that influence a stability metric (and even location) of interest.

### 4.3.3 Caveats of Using SHAP Explanations for Extracting Stability Constraints

Although explainability models such as SHAP can uncover the correlations that predictive ML models use to make accurate predictions, these relationships are not necessarily causal [65]. Consequently, the constraints derived from the proposed approach, implemented within an optimisation model, provide general guidance for overall improvement but do not guarantee a consistent impact across specific operating conditions. In some cases, they may even lead to deterioration under particular operating conditions. However, predictive models can sometimes capture causal effects if key features have a strong independent component, meaning that their predictive power is not redundant with other measured features or influenced by unmeasured confounders [65, 67]. Confounders are variables that influence both the independent and dependent variable(s), potentially leading to misleading conclusions [67]. Therefore, if implementing the SHAP-derived constraint does not alter other unconstrained system features in a way that cumulatively creates antagonistic effects, stability metrics tend to improve. Consequently, the extraction and utilisation of constraints following the process described is expected to yield an average improvement in stability metrics rather than a guaranteed improvement for every single case. This is further discussed and highlighted in subsequent sections.

### 4.3.4 Estimating Locational Frequency Dynamics with Machine Learning

Locational ML models are adopted as surrogate models to predict the frequency at every location or bus in the network, and SHAP explainers to identify important system variables defining stability boundaries. We adopt this locational ML model approach to derive locational- and metric-specific SHAP explanations. Following the performance results of different ML models in Chapter 3.4.4, we therefore adopt the MLP neural network (Chapter 3.2.1) in this chapter. We adapt the model to utilise  $M$  input features comprising physical and operational characteristic variables of the network, that is, system demand/loading, CIG output, SG active power set points, *Tripped SG* —an

## Chapter 4. Enhancing Understanding of Frequency Dynamics and Preventive Stability Management

index indicating the location of the disturbance, that is, the generator being disconnected —and SG ratings, to make a prediction of nadir and/or RoCoF at each of the  $N$  network buses.

Given a power system with  $N$  monitored buses, a total of  $N$  locational MLPs are required to capture the locational frequency dynamics metric associated with each location. In this chapter, the frequency nadir and RoCoF are used to represent the locational dynamics; consequently, up to  $2 \times N$  MLPs were trained. To enhance the understanding of system frequency characteristics, each MLP has an associated SHAP explainer model implemented for additive feature attribution through SHAP values. Moreover, an additional MLP is also implemented to capture the worst-case frequency stability metric (nadir and RoCoF) at the system level to identify associated interactions, thereby demonstrating how operators can identify actionable insights at both local and global scales.

### 4.4 Case Studies and Results

In this chapter, the results of the proposed method are presented. Section 4.4.1 provides details of the case study and the modelling of CIG. Section 4.4.2 evaluates the accuracy of the proposed method’s ability to track locational frequency stability metrics. SHAP is then implemented in Chapter 4.4.3 to enhance the understanding of the system’s frequency dynamics by uncovering the important variables that shape stability boundaries. Section 4.4.4 aims to enhance the locational and global system stability by applying actionable SHAP insights through targeted interventions, i.e., one-sided optimisation constraints, whose effectiveness is evaluated using RMS-TDS. Lastly, Section 4.4.5 discusses the application of the proposed methodology by power system operators to enhance system security. Details on the computational speed of the methodology are also provided in Chapter 4.4.6.

#### 4.4.1 Test Network and Case Study

- **IEEE 39-bus Test Network:** In this chapter, the modified IEEE 39-bus network (Fig. 3.3) is used, and the details are as described in Chapter 3.4.1.
- **Generator Dispatch via AC Optimal Power Flow:** In this chapter, the polynomial cost function of the SGs utilises the same coefficients as specified in Table 3.1 in Chapter 3.4.1, while the cost of CIG is considered to be zero.
- **Operational Scenarios:** In this chapter, we generate the dataset using the same approach as defined in Chapter 3.3.1. We vary the number of online SGs by considering every SG to be an equivalent of four equal-sized units; the system loading ranges between 0.6 p.u. to 1.025 p.u. in steps of 0.025 p.u., and the CIG penetration alternates between high and low penetration levels. The AC-OPF is employed for generation dispatch. The AC-OPF model is used for generation dispatch. This approach ensures that the generated scenarios are operationally feasible and consistent with the actual network's limits, which better reflects the practice by system operators.
- **Dynamic Simulations using RMS-TDS:** In this chapter, the considered disturbance is the disconnection of the largest SG dispatch, excluding G 01, which represents a large equivalent generation. This approach differs from that in Chapter 3.4.1, where the disturbance is fixed at a single generator. By varying the disturbance according to the SG dispatch, we implement a more realistic worst-case contingency each time. Additionally, if multiple SGs have the same active power set points, the same number of simulations is performed, considering each SG as a separate contingency.
- **Frequency Stability Datasets and ML Model Training:** In this chapter, we maintain the approach detailed in Chapter 3.3.2 for data handling. The datasets are standardised using the StandardScaler (which rescales every feature so that its mean is zero and its variance is one) [57] as a pre-processing step to ensure optimal training. Following the standardisation, we adopt a dataset split of 70-

Table 4.1: Frequency Stability Metric (RoCoF (Hz/s) and Nadir (Hz)) Estimation Performance by Different Models

<i>Model</i>	<i>Evaluation Metric</i>	<i>Max RoCoF</i>	<i>Max Nadir</i>	<i>Mean RoCoF</i>	<i>Mean Nadir</i>
<b>MLP</b>	MU	0.0374	0.0005	0.0311	0.0002
	MAE	0.0374	0.0005	0.0311	0.0003
	RMSE	0.0077	0.0001	0.0067	0.0000

30% for model training and testing, respectively. Hyperparameter tuning of the ML models is done using the 5-fold Cross-Validation technique —randomly splits the training set into  $k$ -folds, whereby each  $k-1$  set is used for training, while the remaining one is used for validation.

#### 4.4.2 Accuracy of the Locational ML Models for Predicting Frequency Stability Metrics

In this chapter, the accuracy of the locational ML models to predict frequency stability metrics is presented. It is noted that it is important to have high prediction accuracy for the ML models for the SHAP explainers to extract meaningful insights. A single ML model is required per location per stability metric (frequency nadir and RoCoF are used in this chapter). As a result, the 39 buses in the network required 78 MLPs. . Using the two evaluation metrics discussed in Chapter 3.3.3 to assess the performance of the ML models, that is, RMSE and MAE, in this chapter, we also introduce the Maximum Underestimate, MU, metric to quantify the model’s vulnerability to giving false negatives, which leads to unforeseen violations. Table 4.1. It can be seen that the locational models have the RoCoF maximum and mean MU of 0.0374 Hz/s and 0.0311 Hz/s, respectively, demonstrating the ability to effectively capture the complexities in frequency dynamics. The trend remains consistent for the nadir stability metric, though with less significant differences. This is attributed to the specific case but still demonstrates the general applicability of the MLP method as being more robust than the COI-based method, locational-wise, without any loss of generality.

### 4.4.3 Understanding the System-wide Minimum RoCoF with SHAP Values

In this chapter, we use results from the SHAP *DeepExplainer* [65] to enhance understanding of frequency dynamics. The system-wide minimum RoCoF is used as an example to demonstrate how SHAP can uncover important variables defining the system's stability boundary.

#### i) Local SHAP Explanation - Revealing Effects Around a Specific Operating Condition

Local SHAP explanations enable the characterisation of variables or features affecting a stability metric (RoCoF in this case), specifically for a particular OC. Local SHAP explanations enable the characterisation of variables or features affecting a stability metric (RoCoF in this case), specifically for a particular OC. By calculating SHAP values that are scenario-specific, contributions from the variables toward the stability metric are uncovered, enabling the identification of the most influential features shaping the system's stability boundary. Shown in Fig. 4.2 is a waterfall plot of the local SHAP values, where the horizontal axis represents SHAP values (the difference between the expected value,  $E[f(x)]$ , and the model prediction,  $f(x)$  in Hz/s) whilst the vertical axis contains a sorted list of variables starting with the most significant. The bars in red indicate a positive effect where an increase in the magnitude of the variable leads to an increase in the model's prediction (compared to the expectation), while the bars in blue represent the inverse, i.e., an increase in the prediction of the model is achieved by decreasing the magnitude of the variable.

For the sampled scenario shown in Fig. 4.2, the expected value,  $E[f(x)]$ , is  $-0.327$  Hz/s. In this plot, the most significant variable affecting the model's prediction is the 297 MW dispatch of G 07 (G 07 MW). It reduces the model's prediction by  $-0.24$  Hz/s. This is followed by the 800 MVA rating of G 04 (G4 MVA), which increases the prediction by  $+0.16$  Hz/s. The sum of all the SHAP values moves the prediction of the ML model from  $E[f(x)]$  to the left towards the final prediction,  $f(x)$ , of  $-0.353$  Hz/s. Thus, for the sampled scenario, minimising the dispatch of G 07 suggests an improvement in the system-wide minimum RoCoF. This may occur because the  $N - 1$

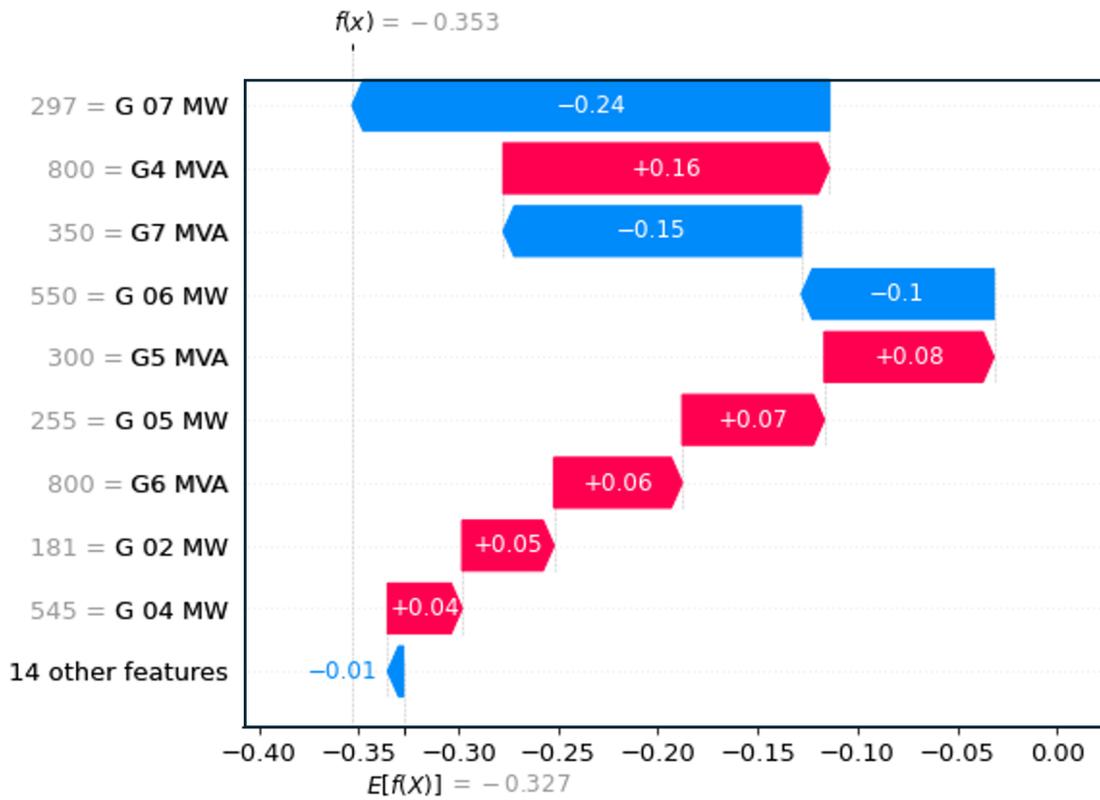


Figure 4.2: Local SHAP explanation of a randomly sampled scenario for the system-wide minimum RoCoF. SHAP values (horizontal axis) are plotted against system features (vertical axis)

SG candidate is likely shifted outside of this critical low-inertia region, or the actual disturbance is reduced, thereby leading to a more stable response. It can also be noted that the most significant variables for this OC belong to the group of SGs situated in the low inertia region, *Area 2*, i.e., G 04 to G 07, which aligns with the findings by [15].

It is important to note that driving the model's prediction towards a desired direction by altering a variable can lead to unintended results due to the combined effect of the remaining unconstrained variables, which can counteract the former. For instance, variables related to generator dispatch are more dependent on the system loading condition than on generator rating, although the rating of the generator determines the maximum dispatch. Therefore, varying the dispatch of a particular generator while maintaining the system loading means that several other generators alter their dispatch as well, whose combined effect may lead to negative effects.

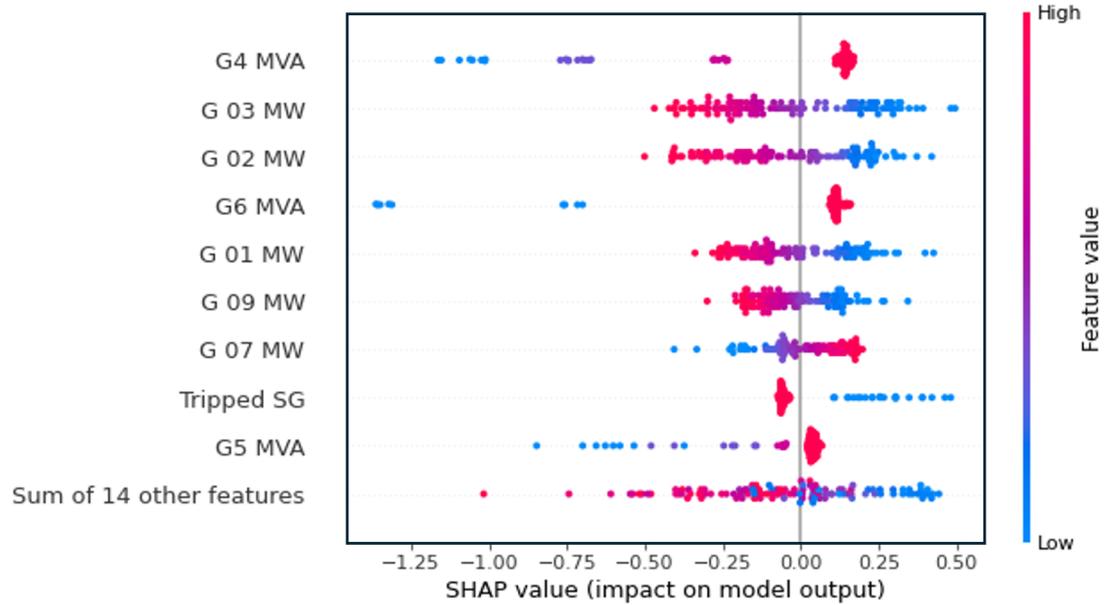


Figure 4.3: Global SHAP explanation for the system-wide minimum RoCoF. Colour represents the magnitude of the feature

### ii) Global SHAP Explanation - Revealing Overall System Trends

A global perspective of SHAP values can also be used to get a holistic understanding of how various variables affect particular aspects of a metric across multiple OCs. This is generated based on the scenarios dataset as an average of *Local Explanations*. Global SHAP explanation provides a summary of the most significant variables across a specific dataset, affecting the model’s prediction of the regression target. In Fig. 4.3, a beeswarm summary plot shows the system variables (vertical axis) against the global explanation SHAP values (horizontal axis) of the test dataset. The colour bar represents the normalised magnitude of features, from low (blue) to high (red), and the rating of G 04 (*G4 MVA*) is the most significant variable, followed by the dispatch of G 03 (*G 03 MW*) and G 02 (*G 02 MW*).

The plot gives an insight into the most significant variables in the system, which dominate in defining the system-wide minimum RoCoF. The explainer indicates that by increasing the rating of G 04, the RoCoF of the system can potentially be improved by up to +0.25 Hz/s. G 04 is located in the low inertia region, *Area 2*, and its in-

creased rating improves the inertial response available in this vulnerable region. In modern large and interconnected power systems, this kind of insight can inform frequency control resource allocation decisions, thereby enhancing efficient resource utilisation. Additionally, the plot also indicates that minimising the dispatches of the two generators, G 03 and G 02, can potentially improve the stability metric by +0.50 Hz/s and +0.44 Hz/s, respectively. This is because the two are among the largest SG units in the system, whereby minimising their dispatch improves their response to system imbalances following a disturbance due to the freed higher ramp rates as demonstrated in [21] —demonstrating that the initial operating point of an SG largely influences the generator response following a disturbance, i.e., heavily loaded SGs exhibit a damped response as compared to lightly loaded SGs.

Following the above, SHAP attributions can effectively act as a good starting point to provide guidance/insights to system operators and planners related to variables affecting complex power system dynamics without having to conduct brute force analyses through the computationally demanding RMS-TDS. In addition, other than local and global plots, SHAP explainers can also enable the inspection of a specific variable of interest to analyse its influence on stability metrics through *Feature Dependency Analysis*. This knowledge can be beneficial as well, for instance, where operators would like to investigate the impact of a CIG plant on the stability boundary. This is discussed in the next section.

### iii) Feature Dependence Analysis - A Focus on Particular Feature Effects

SHAP explainers can analyse the effect of a variable or feature on the stability metric for additional insights. For example, system operators may want to understand the impact of CIGs on frequency stability. Shown in Fig. 4.4 is a feature dependence plot for the CIG in the system, where the vertical axis represents SHAP values and the horizontal axis represents the CIG output in MW. The colour bar represents a closely related variable identified in the system, which is the dispatch of G 04 (*G 04 MW*).

This plot demonstrates how integrating domain knowledge with SHAP insights can enhance decision-making for system operators and planners. A moderate CIG penetration of up to 400 MW, combined with high generation output from G 04, suggests

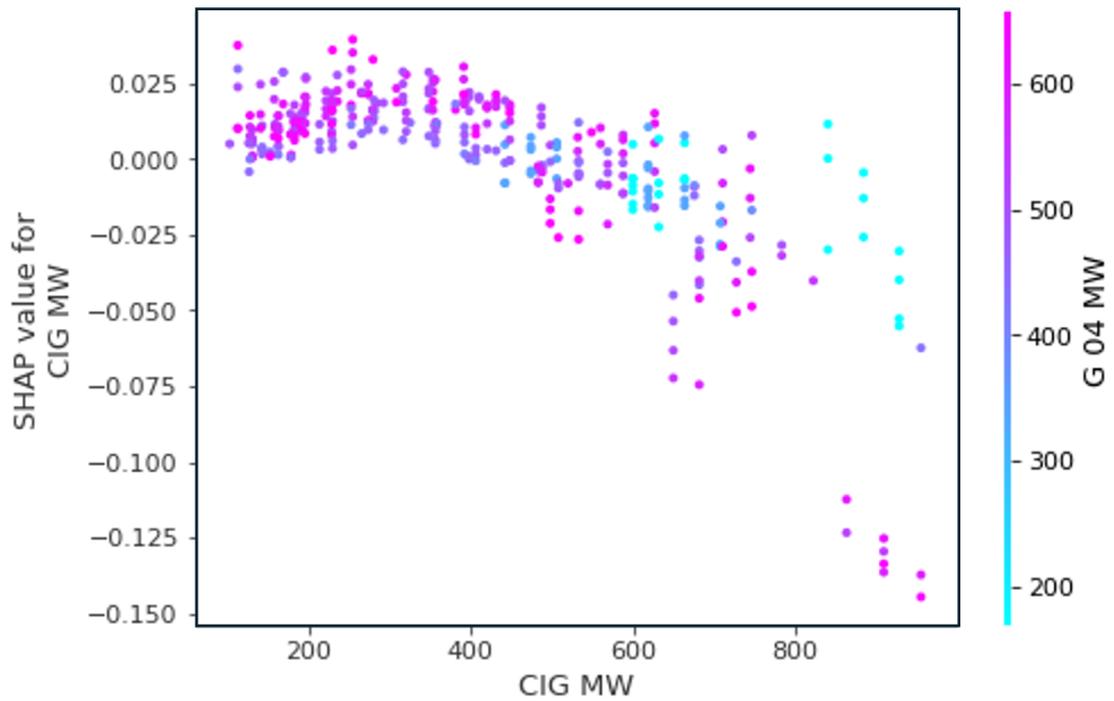


Figure 4.4: System-wide minimum RoCoF Dependence plot for CIG MW penetration. Colour represents G 04 dispatch in MW

a positive effect on the system-wide minimum RoCoF, with a potential improvement of up to  $+0.035$  Hz/s. Furthermore, the *CIG Dependence Analysis* plot indicates that at higher CIG penetration levels (beyond 400 MW), the RoCoF can potentially deteriorate by up to  $-0.150$  Hz/s (especially when combined with high G 04 loading). This insight aligns with existing literature on the impact of CIG and inertia reduction, which validates the SHAP explanations. However, our proposed method can provide specific details for specific units and thresholds as well as provide quantifiable indications of potential impacts, something that is not straightforward to do otherwise. By uncovering these interactions, system operators and planners can gain more confidence in SHAP-derived insights for improved system stability.

#### 4.4.4 SHAP-Informed Optimisation - enhancing frequency stability through actionable SHAP insights

The proposed method uses SHAP values to identify key variables that define the frequency response of the power system. This can be useful for implementing targeted interventions within unit commitment and/or dispatch models to enhance system security. In this chapter, we demonstrate how our method utilises SHAP explanations to provide a simple guide, in the form of one-sided SHAP-informed constraints, to the system optimisation model as discussed in Chapter 4.3.2 using (4.5). The RoCoF of the system is used as the stability metric to evaluate the performance of the proposed approach as follows:

**i) Case Study 1 - Localised Stability Improvement:** *Improve the local RoCoF at the most vulnerable location in the network, using the most significant feature from the corresponding global SHAP explanation. RMS-TDS results show that Bus 34, in the low inertia region (Area 2), is the most vulnerable, contributing over 45% of the lowest RoCoF metric values. Bus 38 follows this, contributing 32%, while the remaining 22% is shared among the remaining buses.*

Localised stability improvements can be achieved by applying insights from global SHAP explanations at the location of interest. Fig. 4.5 shows a global SHAP explanation for RoCoF at the vulnerable location of interest, Bus 34. The most significant feature is *Tripped SG*, which determines the location of the disturbance. Since this chapter looks at the system-wide  $N - 1$  SG outage, this feature is considered uncontrollable. Consequently, the subsequent feature, i.e., *G4 MVA*, which concerns the rating of G 04, is chosen. The SHAP-informed constraint on this feature ensures that the rating of G 04 is minimised. We, therefore, limit the rating to the minimum possible value of 200 MVA. Results obtained from RMS-TDS following incorporation of this constraint within the AC-OPF are shown in Fig. 4.6 and are compared to *Base Case* —involving the AC-OPF without any SHAP-related constraints. The *Base Case* indicates that the RoCoF at Bus 34 ranges from  $-0.4122$  Hz/s to  $-0.1491$  Hz/s and has a mean of  $-0.2676$  Hz/s. The *1st SHAP Feature* case incorporates the SHAP-informed constraint, and this improves the initial RoCoF range at the bus to between  $-0.2938$

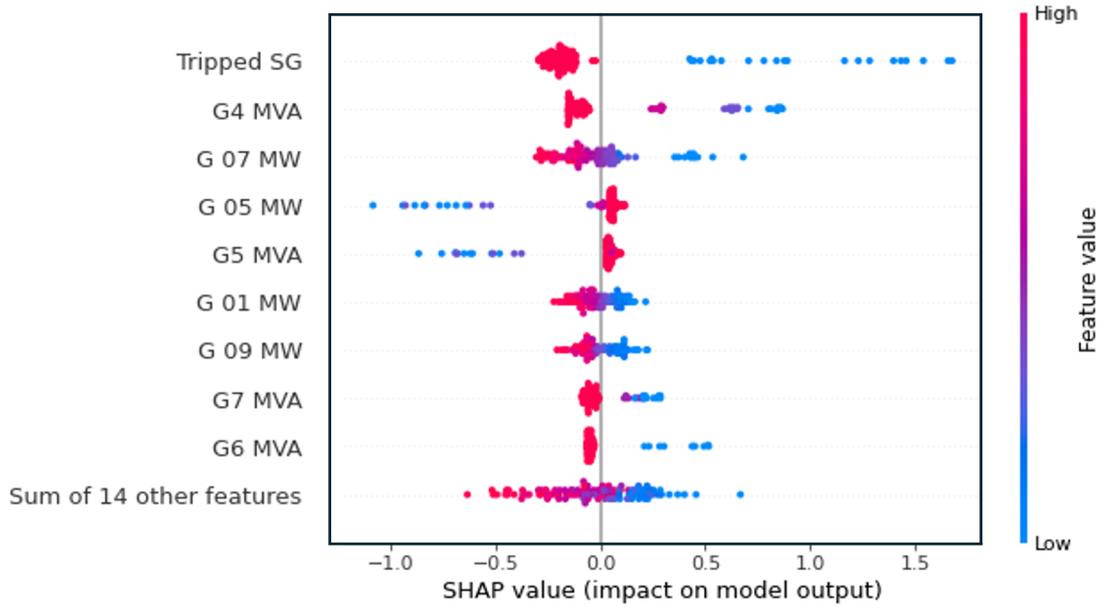


Figure 4.5: Global SHAP explanation for localised RoCoF at Bus 34. Colour represents the magnitude of the feature

Hz/s and  $-0.1403$  Hz/s, with a mean of  $-0.2112$  Hz/s. Thus, 21% and 29% improvement in the mean value and worst-case scenario, respectively. In addition, even when considering outliers, the lowest RoCoF is  $-0.3369$  Hz/s, which remains an improvement over the lowest RoCoF observed in the *Base Case* by 15%.

ii) **Case Study 2 - System-wide Stability Improvement:** *Using a separate ML model trained to predict the minimum RoCoF across the system, improve the system-wide RoCoF using the most significant feature from the system-wide global SHAP explanation. We use a dedicated ML model for this task to reduce the complexity of handling constraints resulting from deriving the system-wide minimum RoCoF from all the previously trained locational models, which could also introduce convergence issues.*

Improvement in the minimum RoCoF observed in any system bus can be achieved using insights from the global SHAP explanation of the system-wide stability metric as described in Chapter 4.4.3. Therefore, the MLP model applied for this task is trained to capture the system-wide worst-case RoCoF (global minimum), and the SHAP explainer is used to identify associated global interactions. This allows system operators and planners to improve the system stability, i.e., RoCoF, across the entire network. In

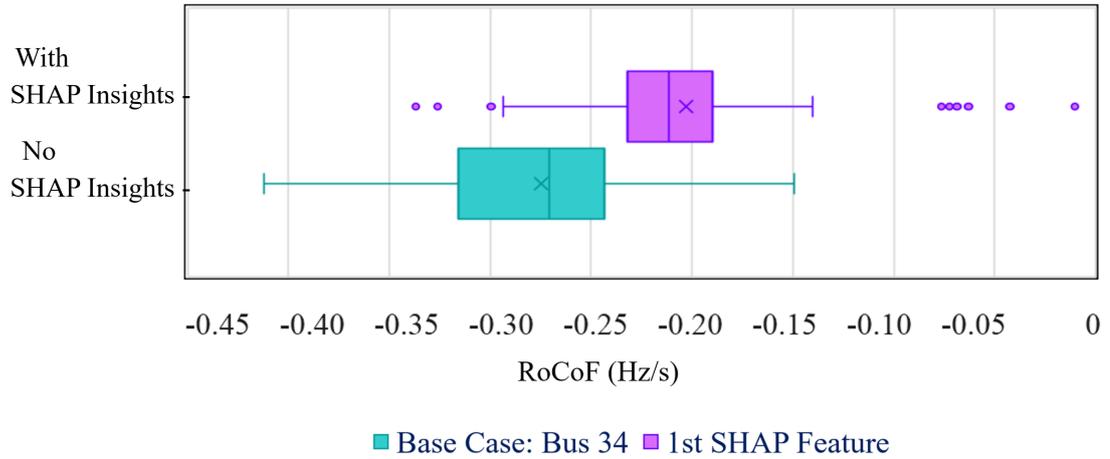
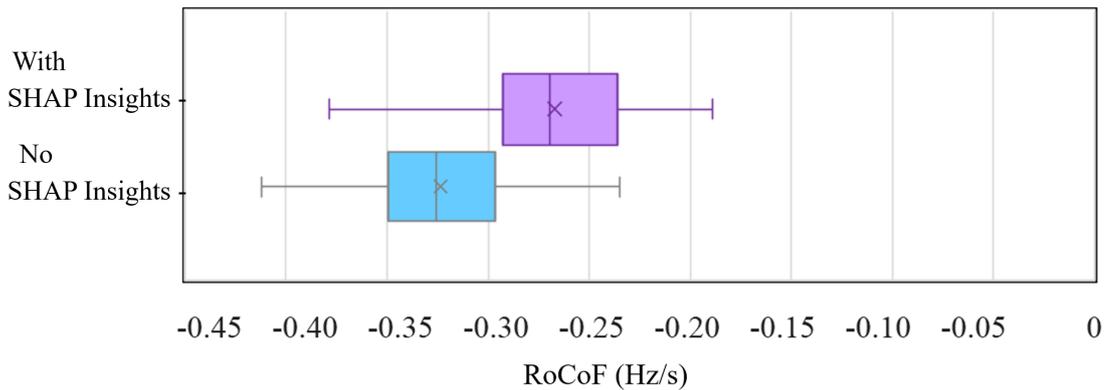


Figure 4.6: Case Study 1: RMS-TDS results of the RoCoF at Bus 34. The *Base Case* is where the standard AC-OPF includes no SHAP insights as constraints

Fig. 4.3, the global SHAP Explanation plot for the system-wide minimum RoCoF shows that the most significant feature is the rating of G 04, i.e.,  $G_4$  MVA. The SHAP-informed constraint is implemented using this feature, which aims to maximise the generator’s rating. We, therefore, constrain the limit to an available maximum rating of 800 MVA. RMS-TDS results following this implementation are shown in Fig. 4.7. The *Base Case* indicates that the system-wide minimum RoCoF ranges between  $-0.4123$  Hz/s to  $-0.2350$  Hz/s, and the mean is  $-0.3236$  Hz/s. However, this is improved in *1st SHAP Feature: System-wide* by the SHAP-informed constraint on G 04 rating, which increases this range to between  $-0.3784$  Hz/s and  $-0.1892$ , with a mean of  $-0.2700$  Hz/s. This represents a 15% improvement in the mean value and a 10% improvement in the worst-case scenario. Note that this insight is in the opposite direction of the SHAP insight for *Case 1* (Subsection 4.4.4) on the same feature, highlighting SHAP’s ability to accurately identify the direction of impact.

**iii) Case Study 3 - Localised and System-wide Stability Improvement:** *Improve both local and system-wide minimum RoCoF by identifying the first common feature between global SHAP explanations for the RoCoF at a specific location (Bus 34, identified as the most vulnerable) and the system-wide minimum RoCoF.*

Combined stability improvements can be achieved by applying common features from respective global SHAP explanation insights. This is demonstrated using lo-



■ Base Case: System-wide ■ 1st SHAP Feature: System-wide

Figure 4.7: Case Study 2: RMS-TDS results of the system-wide minimum RoCoF. The *Base Case* shows a scenario where the standard AC-OPF includes no SHAP insights as constraints

calised RoCoF (at Bus 34, which is identified as most vulnerable) and the system-wide minimum RoCoF, whose global SHAP explanations are shown by Fig. 4.5 and Fig. 4.3 respectively. The first common feature across the two plots is the rating of G 05, i.e., *G5 MVA*. This insight suggests that stability can be enhanced by maximising the MVA capacity or increasing the number of online units at power plant G 05. Therefore, the SHAP-informed constraint is designed to maximise the MVA capacity of G 05. Therefore, this is considered for the SHAP-informed constraint, which aims to maximise the capacity of the generator. We, therefore, limit the rating to a maximum available rating of 300 MVA. The RMS-TDS results are shown in Fig. 4.8. A simultaneous improvement across the two chosen stability metrics is observed. The RoCoF range, specific to Bus 34, is seen to improve to between  $-0.3223$  Hz/s and  $-0.1508$  Hz/s, with a mean value of  $-0.2425$  Hz/s. This represents a 10% improvement in the mean value and a 22% improvement in the worst-case scenario compared to the base case. Likewise, the system-wide minimum RoCoF experiences a similar but more pronounced improvement to between  $-0.3223$  Hz/s and  $-0.1909$  Hz/s and a mean of  $-0.2640$  Hz/s, representing 18% mean value and 22% worst-case scenario improvements. Thus, by maximising the capacity of G 05, a simultaneous improvement in the local and system-wide minimum RoCoFs is achieved.

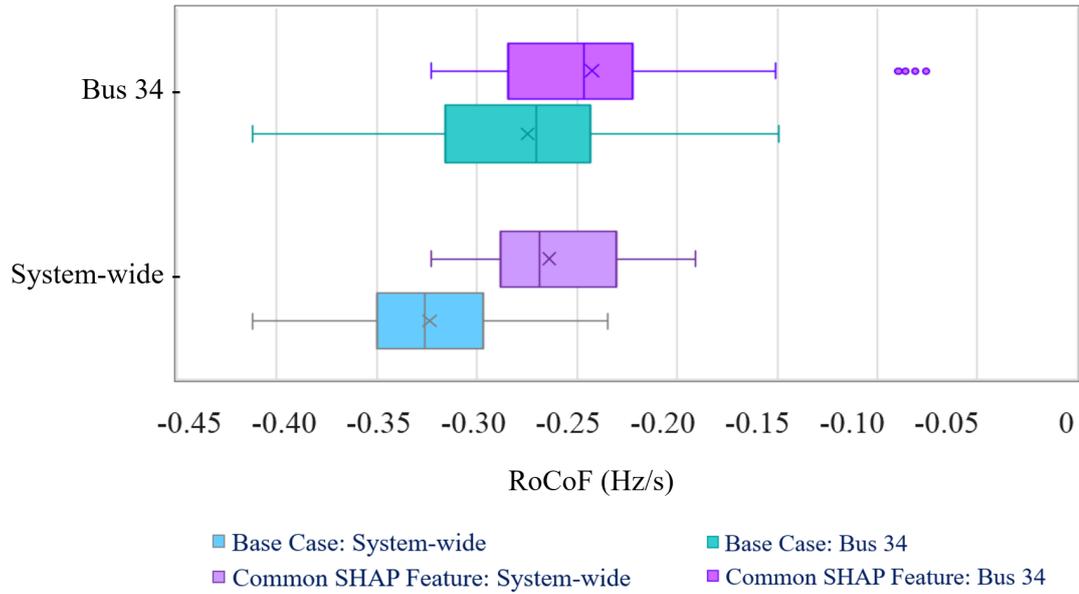


Figure 4.8: Case Study 3: RMS-TDS results of the system-wide minimum and Bus 34 RoCoF. The (*Base Case*) shows a scenario where the standard AC-OPF includes no SHAP insights as constraints

#### 4.4.5 A Discussion on SHAP-Derived Rules for Frequency Stability Enhancement

The proposed approach applies an explainable ML model to identify important system interactions in real-time. A sample of scenarios demonstrating the performance of the proposed approach is shown in Fig. 4.9 using results obtained in *Case Study 1*, from Subsection 4.4.4. The figure shows the RoCoF at Bus 34 following a disturbance for the *Base Case* and the SHAP-informed optimisation case, i.e., *SHAP: Bus 34*. Generally, while significant improvements are achieved in most cases, some scenarios show RoCoF deterioration following the inclusion of SHAP-informed constraints. This is primarily due to the remaining unconstrained features, whose aggregated effect may dominate and lead to unexpected outcomes.

Table 4.2 provides a detailed quantification of the RoCoF changes due to the SHAP-informed constraints for the three case studies of Section 4.4.4, both in terms of their magnitudes and the proportion of samples affected. In *Case Study 1*, which focuses on localised RoCoF at a specific location of interest (Bus 34 was considered as it was

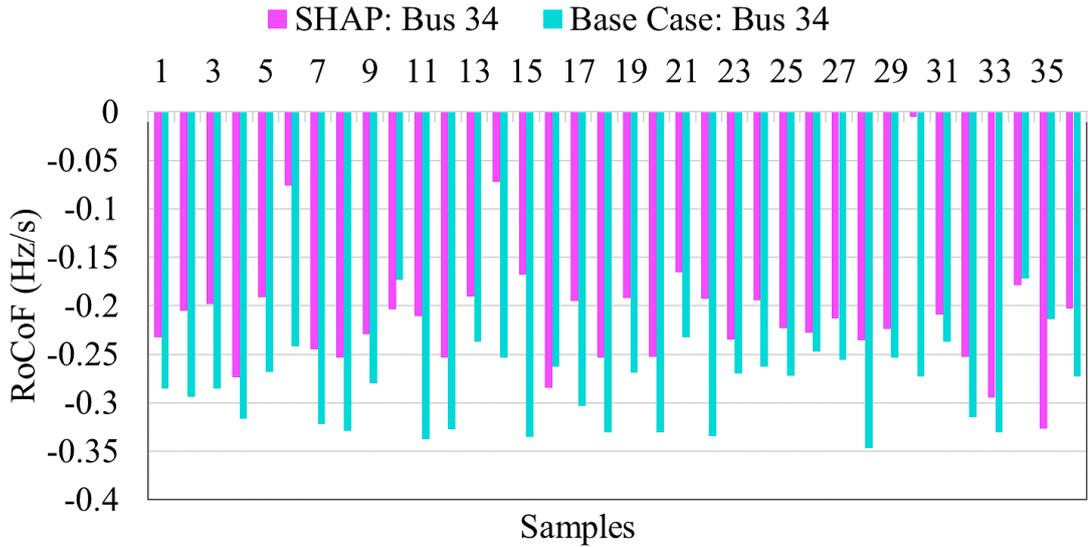


Figure 4.9: Sample scenarios showing the RoCoF at Bus 34 of *Base Case*, and following the incorporation of SHAP-informed constraints within the AC-OPF

identified as the most vulnerable), 90% of the sample experienced a RoCoF improvement following the SHAP-informed constraints, recording a maximum improvement of +0.232 Hz/s. A maximum RoCoF deterioration of  $-0.170$  Hz/s was observed in almost 10% of the remaining sample, indicating some negative impacts due to the introduced constraints. In contrast, 57% of the samples in *Case Study 2*, targeting the system-wide minimum RoCoF, experienced improvement, recording a maximum improvement of +0.208 Hz/s. No deterioration was observed in any of the samples, suggesting that the constraint effectively avoided negative outcomes at the system level (system-wide), albeit with a smaller sample improvement compared to *Case 1*. Furthermore, 42% of the sample experienced no change. Additionally, for *Case 3*, 85% of the samples showed improvement, recording a maximum improvement of +0.192 Hz/s. On the contrary, 10% of the sample exhibited some RoCoF deterioration, with a maximum recorded value of  $-0.067$  Hz/s, indicating a negative, but not significant in magnitude, effect of the introduced SHAP-informed constraints. Note that while the feature with the dominant SHAP value may not always be directly controllable at operational timescales, lower-ranked features may still offer viable insights for guiding practical interventions —as observed in *Case Study 1* in Chapter 4.4.4, where interventions to improve local

Table 4.2: Performance of the SHAP-informed Optimisation Model for three Case Studies

<i>Description</i>	<i>Case Study 1</i>	<i>Case Study 2</i>	<i>Case Study 3</i>
<b>Max. Deterioration</b>	-0.170 Hz/s	0.000	-0.067 Hz/s
<b>Max. Improvement</b>	+0.232 Hz/s	+0.208 Hz/s	+0.192 Hz/s
<b>OCs with Improvement</b>	90.52%	57.14%	85.09%
<b>OCs without Changes</b>	0.00%	42.86%	4.39%
<b>OCs with Deterioration</b>	9.48%	0.00%	10.53%

stability were informed by the second most significant feature, since the most significant one, i.e., *Tripped SG* (the location of the disturbance), was considered uncontrollable. This suggests that even when the most significant feature is not directly controllable, SHAP explanations may still support the identification of practically useful interventions.

The insights derived from SHAP values lead to a better understanding of the system’s frequency dynamics. The proposed method can facilitate the understanding of the system by providing a foundation for further analysis of system interactions without relying on the computationally intensive brute force RMS-TDS. Important variables shaping the stability boundary at critical or vulnerable locations or system-wide stability boundary can be uncovered, and targeted interventions, combined with domain knowledge, can be implemented. Moreover, the efficiency of decisions, such as the allocation of frequency control resources, can be enhanced through such actionable insights. Note that the proposed method can effectively identify crucial interventions for improving system stability, even if the disturbance location changes due to SHAP-informed re-dispatch. This is because it accounts for the worst-case disturbance in every scenario. However, while the insights offer general improvement, this approach cannot offer guarantees that opposite effects will not appear for specific operating conditions, as discussed in this chapter.

#### 4.4.6 Computational Aspects

Concerning computational requirements, the SHAP explainer (*DeepExplainer*), computed once and offline, following the MLP training phase, took approximately 70 sec-

## Chapter 4. Enhancing Understanding of Frequency Dynamics and Preventive Stability Management

onds for the test dataset containing around 1,200 samples. It generated a global explanation plot in only 0.1494 seconds. When explanations are performed online, local SHAP explanations are generated, requiring approximately the same amount of time. However, authors acknowledge that training  $N$  separate MLPs can be burdensome, particularly for very large-scale power systems with significant datasets, despite being a one-time exercise. To address this, there are two main mitigation strategies as follows:

- i. instead of training MLP models per bus, the models can be trained at a regional level, significantly reducing the overall number of required models. In such cases, the models can be trained to predict the respective region's worst-case responses to avoid unforeseen violations. Similarly, SHAP value-based explanations can remain useful in providing regional-level actionable insights for improving the stability within respective regions.
- ii. utilising a single MLP to capture the worst-case locational frequency response across the entire system, as demonstrated by Case 2 (Section 4.4.4). Ultimately, the derived SHAP value-based explanations remain valid in extracting actionable insights for improving the overall system's stability, without requiring  $N$  separate models.

### 4.5 Conclusion

In this chapter, we have demonstrated that Machine Learning (ML), combined with explainability methods like SHapley Additive exPlanations (SHAP), can effectively capture and provide insights into the local characteristics of system frequency response. This approach also offers ways to improve local frequency characteristics through targeted interventions in the form of constraints during security-constrained optimisation. As local frequency characteristics become more prominent due to the increasing CIG penetration, conventional COI-based approaches become ineffective in detecting local violations of frequency stability metrics, particularly the Rate of Change of Frequency (RoCoF). This can lead to unforeseen relay activations, potentially resulting in load shedding or local disconnection of distributed resources.

## Chapter 4. Enhancing Understanding of Frequency Dynamics and Preventive Stability Management

To this end, we propose a machine learning-based approach for estimating frequency nadir and RoCoF across various system locations. We also use SHAP values to enhance the interpretability of high-accuracy black-box ML models, such as neural networks, and understand the locational aspects of frequency dynamics. Using both local and global SHAP explanations, system operators can gain detailed insights into system parameters affecting local frequency characteristics for specific OCs and general insights across various conditions. Finally, we propose implementing targeted interventions derived from SHAP insights in the form of optimisation constraints to improve local frequency characteristics.

The method is applied and tested in the modified IEEE 39-bus network and validated with dynamic TDS, showcasing improvements in frequency response characteristics, both targeted at specific locations as well as across the entire system. The impact of SHAP insights implemented as actionable insights in an AC-OPF is also validated, showing overall improvement. However, limited cases (up to 10% in our studies) were observed where local frequency characteristics deteriorated due to the cumulative effect of the remaining unconstrained features, indicating the need for caution, as no formal guarantees can be offered for no violation by controlling only a very limited number of features.

## Chapter 5

# Security Constrained Optimisation for Preventive Frequency Stability

In Chapter 5, we address challenges encountered by conventional analytical security-constrained optimisation models. We propose a neural network-constrained OPF model to fulfil the conditions for locational frequency stability. Locational frequency dynamics, including the ‘hard-to-model’ interactions, can be captured by the neural network during offline training and embedded within the OPF. Consequently, secure generation dispatch, ensuring no frequency violations across all the monitored locations in the network, is achieved without imposing significant additional computational requirements on the optimisation process —which arises from solving DAEs/ numerically equivalent algebraic equations. We then implement multiple case studies and emphasise the criticality of the growing locational frequency responses driven by the increasing integration of renewables. Thereafter, comprehensive comparisons are conducted on a standard network, comparing the proposed approach and baseline methods from the literature. The paper presenting this methodology is currently under review with IEEE Transactions on Sustainable Energy as [Pub. D], building on [Pub. G]. Therefore, in response to the existing limitations associated with achieving computationally efficient and secure generator dispatch —preventing frequency violations across all network

locations —this chapter makes the following contributions:

- i. A neural network-constrained Optimal Power Flow (OPF) model that ensures locational frequency stability, by directly incorporating locational dynamics —including the ‘hard-to-model’ interactions —through the embedded neural network.
- ii. A method to identify and classify most active and most inactive regions of neural networks, based on empirical activation statistics —enabling identification of high-risk and low-risk regions for safe linearisation.
- iii. Abound-tightening technique that assigns fixed output bounds to highly active regions of the neural network, while pruning out highly inactive regions, thereby minimising linearisation errors, and facilitating convergence in optimisation.

## 5.1 Artificial Neural Network Constrained Optimal Power Flow for Locational Frequency Stability

The reliability of traditional analytical security-constrained models is diminishing due to increasingly localised frequency dynamics, driven by areas with high local CIG penetration. These models typically rely on the COI-based model, i.e., the average frequency response of the system weighted by inertia contribution from SGs, hence inherently incapable of capturing the extremes of spatially distributed dynamics. As a result, where such spatial variations in frequency dynamics are not considered, the approach can lead to a false sense of security by missing the possible locational violations (extremes). In contrast, to avoid such scenarios from occurring, operators resort to over-securing the system (allocating more frequency control resources than would be necessary) to avoid such scenarios from occurring, potentially introducing unnecessarily high operational costs. To address this challenge, there are two major obstacles: i) the growing complexity of modern power systems difficult to fully capture analytically, and ii) including more detailed analytical approaches by solving the DAEs of the network requires significant computational resources, hindering adoption in real-time (or close to real-time) applications or large-scale analysis of scenarios. This explains why simplified COI-based

methods remain widely used despite their limitations.

[14, 15] proposed a numerical approach to describe the frequency response of an  $N$ -region network using the COI and an additional term accounting for the  $N$  regional frequency differences —two regions demonstrated in this paper. Although this approach improves the performance of the COI-based model, the definition of such regions remains a challenge [41]. Furthermore, using conservative regression may lead to tighter frequency-security regions, especially if the security boundary is highly non-linear, potentially leading to convergence issues. [40] captures the detailed frequency behaviour of each SG in the system using the Trapezoidal rule. This is a numerical integration method that converts differential equations into their numerical equivalent algebraic form. These algebraic equations are then embedded into the optimisation model to represent the locational frequency characteristics. However, this approach substantially increases computational requirements, particularly due to the large set of constraints and the inherent complexity of solving such systems, which eventually limits its scalability. Moreover, with each additional SG, the number of DAEs to be solved increases proportionally, further increasing the required computational requirements. [44] proposes a security-constrained unit commitment using neural networks for RoCoF estimation. A linearised neural network, incorporated as optimisation constraints, is used for transient stability in [33]. However, while the use of ML models as security constraints has demonstrated potential, the challenge of securing against local frequency stability while accounting for the detailed frequency dynamics in converter-dominated systems has yet to be fully addressed.

## 5.2 Stability Constrained Optimal Power Flow Problem

The AC Optimal Power Flow (OPF) is a well-studied nonlinear optimisation problem in power systems. The objective is to obtain an operating point that minimises the cost of meeting electricity demand, subject to generation and network constraints. Let  $\mathcal{L}$  represent a set of loads,  $\mathcal{B}$  a set of buses, and  $\mathcal{G}$  be the set of generators indexed by  $i$ , and let  $C_i(P_{G,i})$  and  $P_{G,i}$  be the cost of generation and active power output of the  $i^{\text{th}}$  generator, respectively. The OPF problem can be written as follows:

$$\min \sum_{i \in \mathcal{G}} C_i(P_{G,i}) \quad (5.1a)$$

subject to

$$g_0(x_0, u_0) = 0 \quad (5.1b)$$

$$h_0(x_0, u_0) \leq 0 \quad (5.1c)$$

$$h_w(x_w, u_w) \leq 0 \quad (5.1d)$$

where the subscripts 0 and  $w$  represent the steady-state and post-contingency operating points of a power system, respectively, while  $x$  and  $u$  represent the state and control variables of the system. The steady-state equality (5.1b) and inequality (5.1c) constraints consist of power balance and bounds, which are expressed as follows. The power balance equations can be written as,  $\forall i, k \in \mathcal{B}, (i, k) \in \mathcal{L}$ :

$$P_i - \sum_{k \in N(i)} P_{ik} - P_{Li} = 0 \quad (5.2a)$$

$$Q_i - \sum_{k \in N(i)} Q_{ik} - Q_{Li} = 0 \quad (5.2b)$$

where  $P_i$  and  $Q_i$  are the active and reactive power generation at bus  $i$ ,  $P_{ik}$  and  $Q_{ik}$  are the active and reactive power flow from bus  $i$  to  $k$ .  $P_{Li}$  and  $Q_{Li}$  are the active and reactive power demand at bus  $i$ .  $N(i)$  represents the set of buses directly connected to bus  $i$ . The active and reactive power flows are nonlinear functions of bus voltages, and a complete formulation of the power flow constraints can be found in [66]. The bounds on control and state variables, as well as the apparent power limits on transmission lines, can be written as follows:

$$P_{\min,i} \leq P_{G,i} \leq P_{\max,i}, \quad \forall i \in \mathcal{G} \quad (5.3a)$$

$$Q_{\min,i} \leq Q_{G,i} \leq Q_{\max,i}, \quad \forall i \in \mathcal{G} \quad (5.3b)$$

$$S_{\min,ik} \leq P_{ik}^2 + Q_{ik}^2 \leq S_{\max,ik}, \quad \forall (i, k) \in \mathcal{L} \quad (5.3c)$$

$$V_{\min} \leq V_i \leq V_{\max}, \quad \forall i \in \mathcal{B} \quad (5.3d)$$

The inequality (5.1d) represents the post-contingency operating point of the system. In this work, to explicitly capture the locational aspects of the frequency response, a method based on Artificial Neural Networks (ANNs) is proposed. The standard OPF is reformulated to include Neural Network-based stability constraints, providing accurate stability indications to the solver during the optimisation process. This approach permits the handling of complex relationships that cannot be easily represented in the standard optimisation formulation. Consequently, this enables system operators to directly account for the frequency response at every location in the system (instead of relying on COI), including the ‘hard-to-model’ frequency relationships. Using MATPOWER’s callback function framework [66], which enables customisation of the OPF process by allowing user-defined code to be executed at specific stages of the OPF algorithm, we implement the proposed approach.

The resulting OPF model is a non-linear optimisation model. It is solved in MATLAB by extending MATPOWER’s AC-OPF formulation, with MATLAB’s Interior Point Method (MIPS) employed as a solver. Nonlinear optimisation models can encounter numerical issues, especially with very large network sizes. These issues may include non-convergence, infeasibility, or convergence to local optima. To address such challenges, the nonlinear AC OPF can be replaced by a suitable convexified or linearised model [68]. However, such approximations may compromise solution quality [69].

### 5.3 Artificial Neural Networks (ANNs) for Locational Frequency Stability

We use ANNs, described in Chapter 3.2.1, as the ML model due to their proven effectiveness and extensive application in power systems, including the capability to be linearised and integrated as constraints in optimisation problems [19, 33, 43, 44, 70–72]. ANNs have been demonstrated to be scalable, adaptable and able to capture complex relationships (including critical region performance refinement), which is crucial to applications in modern large-scale systems. A deep neural network algorithm using the MLP algorithm is used to capture the locational frequency stability metrics. As shown in Fig. 3.1, the MLP is a feed-forward ANN that contains a minimum of three layers: the input layer, the hidden layer with  $\sigma_v$  neurons each, and the output layer. An MLP with hidden layers,  $V$ , feature vector  $\mathbf{x}$ , weights matrix  $\mathbf{W}_v$  ( $\sigma_v \times \sigma_{v+1}$ ) and bias vector  $\mathbf{b}_v$  ( $\sigma_v \times 1$ ), can be represented by (5.4a)-(5.4c).

$$\mathbf{z}_1 = \mathbf{W}_1^T \mathbf{x} + \mathbf{b}_1 \quad (5.4a)$$

$$\hat{\mathbf{z}}_1 = \Theta_1(\mathbf{z}_1) \quad (5.4b)$$

$$\mathbf{z}_{v+1} = \mathbf{W}_{v+1}^T \hat{\mathbf{z}}_v + \mathbf{b}_{v+1} \quad (5.4c)$$

The model's activation function is represented by  $\Theta$ . We use the Rectified Linear Unit (ReLU) activation function given by  $\hat{\mathbf{z}}_v = \max(0, \mathbf{z}_v)$  due to its superior properties compared to other functions, such as linear and sigmoid activation functions, among others. To train the MLP algorithm, the first step is to propagate the features up to the output layer, a step known as forward propagation. Thereafter, based on the output, the error,  $e$ , is calculated whereby the weights in the hidden layer(s) are adjusted to minimise the same (back-propagation) [57]. This is achieved by optimising the parameters,  $\delta^\theta$  (weights and biases), through the error derivative at a specified learning rate,  $\eta$ , as follows:

$$\delta_i^\theta \leftarrow \delta_i^\theta - \eta \nabla e \quad (5.5)$$

Finally, these steps are repeated several times over epochs to establish the best model parameters. The predicted frequency stability metric,  $\hat{y}$ , is derived from the last layer as:

$$\hat{y} = \mathbf{W}_{V+1}^T \hat{\mathbf{z}}_V + \mathbf{b}_{V+1} \quad (5.6)$$

## 5.4 Methodology for Deep Learning Constrained Optimal Power Flow

As highlighted in Chapter 3, deriving the locational frequency response across different network buses and formulating associated locational constraints is challenging. This task is particularly difficult in systems with CIG penetration, where the dynamic response of CIGs can affect the locational frequency characteristics. Maintaining the detailed locational frequency response characteristics through RMS-TDS is increasingly important in such systems. However, this approach incurs significant computational cost, especially if done as part of an optimisation routine. To address this challenge, we propose a method that captures detailed locational frequency characteristics by training neural networks. The trained neural network is then reformulated as constraints to implement stability constraints into an optimisation problem efficiently.

### 5.4.1 Overview of the Methodology

The methodology is summarised in the three stages shown in Fig. 5.1. In the Offline stage, the database for training and testing the ML model for locational frequency stability metric prediction is generated based on RMS-TDS results from a phasor domain simulation software. After training and testing, the frequency dynamics ANN model is linearised and formulated as an optimisation constraint(s). In the Online stage, which is for application, generators are dispatched economically while observing the embedded locational frequency stability constraints. Thus, the proposed approach excludes the

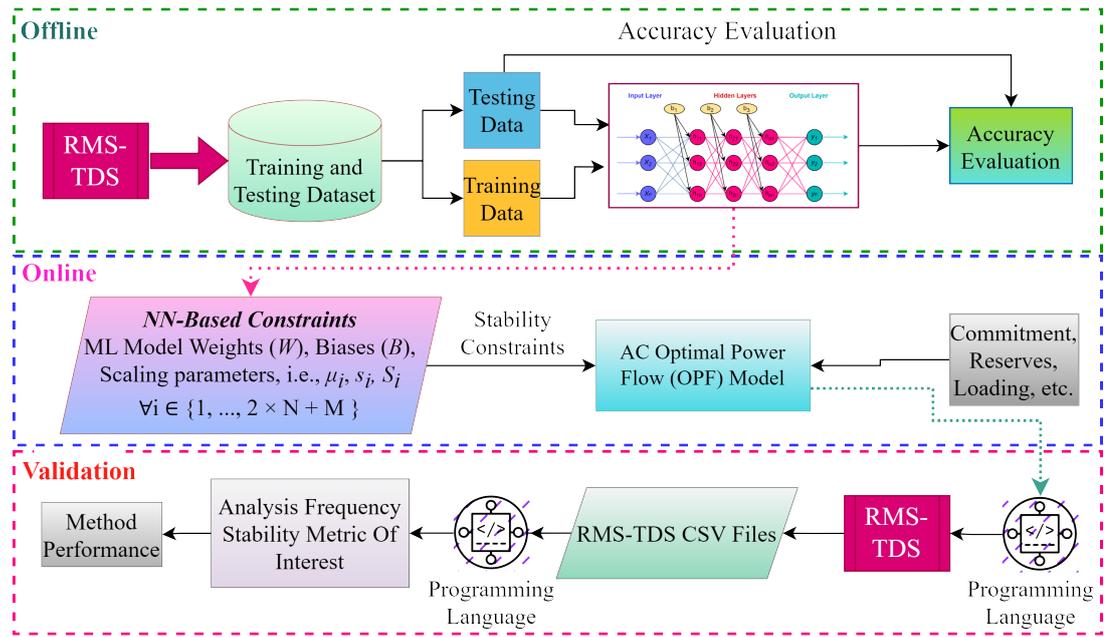


Figure 5.1: A three-stage methodology for the development of a Neural Network Constrained Optimal Power Flow (OPF)

need to iteratively assess the dynamic response of the OC using RMS-TDS, which are computationally intensive and can be very slow for large-scale power systems [29,59,72]. Moreover, there is no need to formulate and embed explicit system equations within the optimisation. Lastly, the Validation stage is an offline stage for performance evaluation using the ground-truth RMS-TDS to ascertain that there are no violations across the whole network.

### 5.4.2 Dataset Generation and ML Model Training

The AC-OPF model is applied in the generation of system operational scenarios by varying three variables, i.e., the system loading, the number of SGs connected and the CIG output. Details regarding the penetration of CIGs and subsequent displacement of SGs are provided in Chapter 3.3.1. The penetration of CIG is modelled using the Type 4 WECC control model [62, 73]. The CIG is integrated into the grid through a fully rated converter without provision for any form of synthetic inertia. The output of the CIG remains fixed at the power dispatch based on (6.16). All the SGs maintained

the default active and reactive power output limits of 0.2 to 0.85 p.u. and -0.3 to 0.7 p.u. respectively, as per rating. The standard quadratic polynomial power generation cost function in the optimisation problem is represented as:

$$\text{Cost } (\mathcal{L}/hr) = \sum_{i \in \mathcal{G}} c_{2,i} P_{G,i}^2 + c_{1,i} P_{G,i} + c_{0,i}, \quad \forall i \in \mathcal{G} \quad (5.7)$$

where  $c_{0,i}, c_{1,i}, c_{2,i}$  are constant, linear and quadratic cost coefficients of the cost curve of SG  $i$ , respectively [3, 66]. The cost function coefficients for each SG are independently sampled from a normal distribution iteratively within minimum and maximum cost boundaries, i.e.,  $c_{0,i} = 0, 1 < c_{1,i} < 10, 0.1 < c_{2,i} < 1, \forall i \in \mathcal{G}$ . This iterative process reflects scenarios where some SGs become more expensive, leading to different dispatches necessary for the training dataset [74, 75]. Details of the SG cost function coefficients used for testing the proposed method are found in Table 3.1 in Chapter 3.4.1, while the generation cost of CIG is considered to be zero. The locational frequency stability dataset is generated by simulating the largest generation loss, that is, an  $N - 1$  SG outage, based on the AC-OPF solutions. RMS-TDS are conducted, and the detailed locational frequency stability metrics, i.e., nadir and RoCoF, are calculated and recorded. Standardisation of the dataset is as described in Chapter 3, Section 3.3.2.

The fully connected MLP architecture used in the chapter has the ReLU activation function across three hidden layers with 400 neurons each. The maximum number of iterations is 3000, while the alpha and the learning rate of the model are 0.01 and 0.001, respectively. These parameters are optimally determined by the sklearn-GridSearchCV [57], which is a hyperparameter tuning technique automating the process of finding optimal hyperparameters for an ML model over a user-defined range. The regression target of the ANN model is the locational frequency stability metrics of interest (e.g. nadir and RoCoF) that will be monitored following a contingency. The inputs to the ANN are the steady-state OC (pre-disturbance) system variables, i.e., SGs and CIG active power set points and system loading, and some physical characteristics of the system, i.e., SG ratings and disturbance index (thus the generator, which is disconnected). The ML model is trained by simulating the largest loss of generation

for different OCs using RMS-TDS and calculating the worst-case response from the detailed locational frequency stability metrics for each of the  $N$  buses of the network, i.e. a single output for nadir or RoCoF. Consequently, the ANN learns the complex underlying relationships related to the locational frequency response metrics.

### 5.4.3 Incorporating Stability Constraints Within Optimisation Models

This chapter describes how the standard OPF formulation is modified to incorporate the proposed ML-based locational frequency stability constraints. The standard general optimisation vector,  $\bar{\mathbf{x}}$ , of the power system is expressed as:

$$\bar{\mathbf{x}}^T = [\Theta_m \quad V_m \quad P_g \quad Q_g] \quad (5.8)$$

where  $\Theta_m$  are the voltage angles,  $V_m$  are the voltage magnitudes,  $P_g$  and  $Q_g$  are the active and reactive power generation, respectively. To incorporate the locational frequency stability constraints, the optimisation vector  $\bar{\mathbf{x}}$  is merged with the vector  $\bar{\mathbf{n}}$ , which includes static, user-defined additional variables associated with frequency prediction, i.e., ML model biases and static system characteristic features. This unified set is required for the final stability decision vector,  $\mathbf{x}_{\text{freq}}$ , which is optimised. Dynamic constraints, as opposed to standard linear constraints, are used in this study to represent the frequency dynamics in the system optimisation model. This approach enables the customisation of the optimisation process and the addition of new variables. It enables handling more complex tasks compared to the standard OPF formulation, as well as attaining more efficient solutions. Therefore, the optimisation of system variables for locational frequency stability takes the form:  $f_{\min} \leq \mathbf{A}\bar{\mathbf{x}} + \bar{\mathbf{n}} \leq f_{\max}$ , where  $f_{\min}$  and  $f_{\max}$  are the stability bounds (e.g., minimum nadir, or RoCoF), and  $\mathbf{A}$  is a sparse system state matrix specifying the active variables required as the model's inputs (inactive variables are represented by 0). The choice of the upper stability metric bound has a significant impact on the optimisation solution space. Consequently, this must be significant enough not to limit optimisation's convergence. Thus, the more relaxed the upper bound is, the easier it is for the optimisation problem to converge.

This provides a trade-off between easily finding feasible solutions and a tight frequency security boundary. The frequency stability metrics thresholds for RoCoF and nadir threshold are  $-0.5$  Hz/s and  $59.6$  Hz, respectively [12, 76]. Since linearisation is required to implement stability constraints based on the trained ANN, a direct approach is to train a linear model. However, this approach fails to capture the full extent of the highly complex relationships in the data, which can be captured by the ReLU activation function. Alternatively, binary variables can be used in specialised and/or commercial solvers to linearise the ReLU function. However, the binary variables require element-wise computation at each neuron since activation functions are implemented for each neuron, leading to high computational requirements and/or longer simulation times. To address this, we adopt a convex relaxation of the ReLU activation function using the Triangle relaxation method [77, 78], discussed in the next Chapter.

#### 5.4.4 Triangle Relaxation of the ReLU Activation Function for Generating Frequency Stable Optimisation Solutions

This chapter presents the proposed approach, referred to as  $\Delta - OPF$ , which ensures locational frequency stability by incorporating ANN-based stability constraints into the optimisation process. Introducing non-linear stability constraints into a power system optimisation model significantly increases the complexity of the dispatch problem, making it difficult to solve efficiently. To address this, we propose neural network linearisation. This enables us to strike a balance between the computational simplicity of a linear ANN and the complexity of the ReLU activation function. The Triangle relaxation method [77, 78] aims to approximate the ReLU activation function in a bounded input space by offering a purely linear representation. Fig. 5.2 illustrates the method's approximation of the ReLU through linear inequalities, thereby minimising complexity in optimisation. It over-approximates the ReLU function except at the input boundaries, and this overestimation is the highest for values close to zero. This property guarantees the generation of stable scenarios and makes the method suitable for verification tasks [77]. Given a trained neural network with ReLU activation functions, the Triangle relaxation constrains the output of every neuron within a minimum pre-activation lower

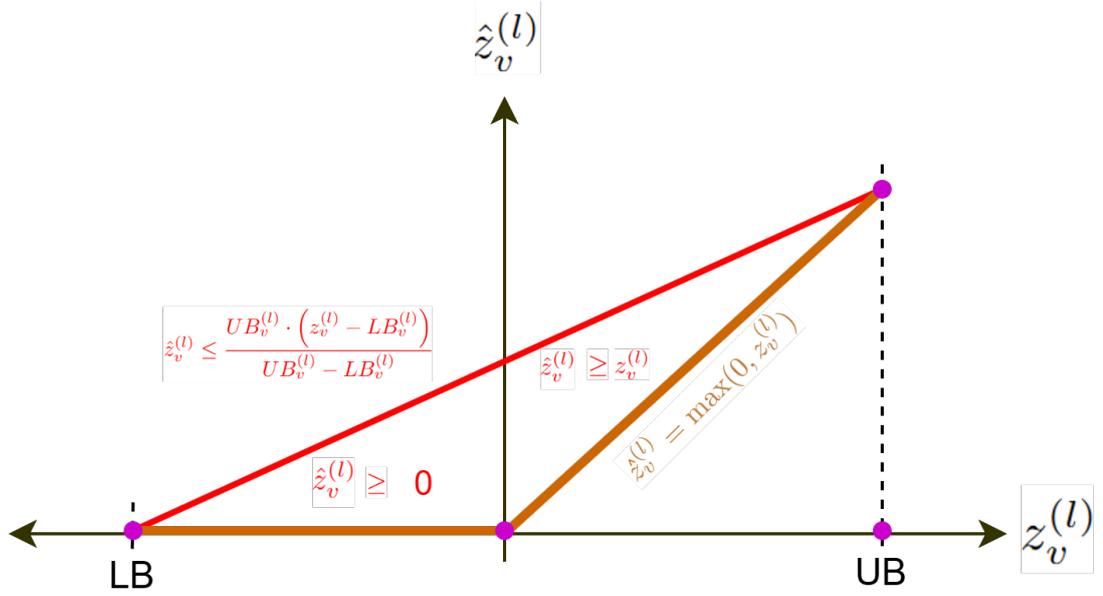


Figure 5.2: Triangle relaxation (red) of the ReLU activation function (brown) with bounded inputs  $z_v \in [LB, UB]$  for layer  $v$  neuron  $l$

bound,  $LB$ , and a maximum pre-activation upper bound,  $UB$ . The training dataset is used to compute these neuron-wise upper and lower bounds, assuming it is fully representative of the data distribution. We denote the pre-activation output,  $z$ , of neuron  $l$  in layer  $v$  of a neural network with  $V$  hidden layers as  $z_v^{(l)}$ . The scaled upper bound constraint for Triangle relaxation is expressed as,  $\forall l = [1, \dots, L], \forall v = [1, \dots, V]$ :

$$\hat{z}_v^{(l)} \leq \frac{UB_v^{(l)} \cdot (z_v^{(l)} - LB_v^{(l)})}{UB_v^{(l)} - LB_v^{(l)}} \quad (5.9)$$

This can be rewritten as an affine expression as follows:

$$G_{v,1}^{(l)} = \frac{UB_v^{(l)}}{UB_v^{(l)} - LB_v^{(l)}}, \quad G_{v,2}^{(l)} = \frac{UB_v^{(l)} \cdot LB_v^{(l)}}{UB_v^{(l)} - LB_v^{(l)}} \quad (5.10)$$

Substituting into the original inequality (5.9), the output of every neuron can be constrained as follows:

$$\hat{z}_v^{(l)} \leq G_{v,1}^{(l)} \cdot z_v^{(l)} - G_{v,2}^{(l)}, \forall v = [1, \dots, V], \forall l = [1, \dots, L] \quad (5.11)$$

The complete relaxed neural network from the input to the final prediction can therefore be represented as follows:

$$z_1^{(l)} = \mathbf{W}_1^T \mathbf{x} + \mathbf{b}_1^{(l)}, \quad \hat{z}_1^{(l)} = G_{1,1}^{(l)} \cdot z_1^{(l)} - G_{1,2}^{(l)} \quad (5.12a)$$

$$z_2^{(l)} = \mathbf{W}_2^T \hat{z}_1^{(l)} + \mathbf{b}_2^{(l)}, \quad \hat{z}_2^{(l)} = G_{2,1}^{(l)} \cdot z_1^{(l)} - G_{2,2}^{(l)} \quad (5.12b)$$

⋮

$$z_v^{(l)} = \mathbf{W}_v^T \hat{z}_{v-1}^{(l)} + \mathbf{b}_v^{(l)}, \quad \hat{z}_v^{(l)} = G_{v,1}^{(l)} \cdot z_v^{(l)} - G_{v,2}^{(l)} \quad (5.12c)$$

$$z_V^{(l)} = \mathbf{W}_V^T \hat{z}_v^{(l)} + \mathbf{b}_V^{(l)}, \quad \hat{z}_V^{(l)} = G_{V,1}^{(l)} \cdot z_V^{(l)} - G_{V,2}^{(l)} \quad (5.12d)$$

$$\hat{y} = \mathbf{W}_{V+1}^T z_V^{(l)} + \mathbf{b}_{V+1}, \quad \text{or} \quad z_{V+1} = \mathbf{W}_{V+1}^T z_V^{(l)} + \mathbf{b}_{V+1} \quad (5.12e)$$

As described in Section 5.4.3, to integrate the stability constraints within our optimisation model, we need them expressed linearly in the form  $f_{\min} \leq \mathbf{A}\bar{\mathbf{x}} + \bar{\mathbf{n}} \leq f_{\max}$ . To achieve this, we collapse the weights, biases and  $G_v^{(l)}$  parameters in (5.12) as follows:

$$\bar{\mathbf{W}}_1 = G_{1,1} \circ \mathbf{W}_1^T, \quad \bar{\mathbf{b}}_1 = G_{1,1} \circ \mathbf{b}_1 - G_{1,2} \quad (5.13a)$$

$$\bar{\mathbf{W}}_2 = G_{2,1} \circ \mathbf{W}_2^T, \quad \bar{\mathbf{b}}_2 = G_{2,1} \circ \mathbf{b}_2 - G_{2,2} \quad (5.13b)$$

⋮

$$\bar{\mathbf{W}}_V = G_{V,1} \circ \mathbf{W}_V^T, \quad \bar{\mathbf{b}}_V = G_{V,1} \circ \mathbf{b}_V - G_{V,2} \quad (5.13c)$$

$$\bar{\mathbf{W}}_{V+1} = \mathbf{W}_{V+1}^T, \quad \bar{\mathbf{b}}_{V+1} = \mathbf{b}_{V+1} \quad (5.13d)$$

$$y = \bar{\mathbf{W}}_{V+1} \left( \bar{\mathbf{W}}_V \left( \dots \left( \bar{\mathbf{W}}_2 \left( \bar{\mathbf{W}}_1 \mathbf{x} + \bar{\mathbf{b}}_1 \right) + \bar{\mathbf{b}}_2 \right) + \dots \right) + \bar{\mathbf{b}}_V \right) + \bar{\mathbf{b}}_{V+1} \quad (5.13e)$$

where  $G_v^{(l)}$  parameters are applied using element-wise multiplication ( $\circ$ ), while a dot product is used by the weights to transform the input from the previous layer. This simplifies the level of implementation complexity and the number of new parameters introduced into the optimisation model. The final stability decision vector that is optimised then becomes:

$$\mathbf{x}_{\text{freq}}^T = \left[ \Theta_m \quad V_m \quad P_g \quad Q_g \quad \mathcal{F}_i \quad \theta_j \right] \quad (5.14)$$

where  $\mathcal{F}_i$  represent additional variables (ANN input features) not in the standard optimisation vector, which are: the system loading condition, SG rating, CIG output and disturbance index, while  $\theta_j$  is the set of model biases,  $\mathbf{b}_v$ . All active variables in the sparse system matrix,  $\mathbf{A}$ , (with corresponding values  $\neq 0$ ) are therefore updated with weighting parameters derived from the linearised ANN (5.13).

#### 5.4.5 Selective Activation Pruning OPF (SAP- $\Delta$ -OPF)

The Triangle relaxation introduces looser bounds, particularly in neurons with high variance, which provides some guarantee to generate frequency-stable solutions. However, this approach can significantly reduce the feasible region space, potentially leading to long optimisation times with low convergence rates. To address this and approximate the ReLU activation behaviour within OPF constraints more efficiently, we introduce Selective Activation Pruning OPF (SAP- $\Delta$ -OPF), a method for bounding neurons based on their empirical activation ratios, rather than imposing hard constraints on every neuron—which introduces significant conservatism, particularly around the region  $z_v^{(l)} \approx 0$ . We analyse neuron-wise behaviour following training to identify the most active and most inactive neurons as an optimal set for linearisation, leaving the ambiguous ones unconstrained. Given a trained neural network, we assign  $\text{UB}_v^{(l)} = \text{LB}_v^{(l)} = 0$  to neurons mostly inactive, and  $\text{UB}_v^{(l)} = 1, \text{LB}_v^{(l)} = 0$  to neurons mostly active. We classify neurons using the activation ratio,  $\Sigma_v^{(l)}$ , for each neuron  $l$  in all  $V$  layers as follows:

$$\Sigma_v^{(l)} = \frac{1}{N_e} \sum_{n=1}^{N_e} \pi \left[ z_{v,n}^{(l)} > 0 \right], \forall v = [1, \dots, V], \forall l = [1, \dots, L] \quad (5.15)$$

where  $N_e$  is the total number of data samples,  $z_{v,n}^{(l)}$  is the preactivation value of neuron  $l$  in layer  $v$  for training sample  $n$ , and  $\pi[\cdot]$  is the indicator function, returning 1 if the condition is true, and 0 otherwise. Therefore,  $\Sigma_v^{(l)}$  lies in the range  $[0, 1]$ , and we use 0.95 as a threshold to classify neurons. Specifically, a neuron with  $\Sigma_v^{(l)} \geq 0.95$  is considered most active, while a neuron with  $\Sigma_v^{(l)} \leq 0.05$  is considered most inactive.

### 5.4.6 Model Evaluation and Baseline Methods

The predictive accuracy of the trained ANN model is assessed using the evaluation metrics RMSE and MAE, described in Chapter 3.3.3. Finally, the performance of the proposed security-constrained method,  $\Delta - \text{OPF}$ , is compared against the normal OPF without any frequency stability-related constraints (referred to as OPF) mainly in terms of computation time and generating frequency-stable solutions. Additionally, two baseline methods from the literature are also included for comparison as follows:

**i) Mixed Analytical-numerical method (Numerical method):**

To obtain a linear formulation of the constraints to guarantee that the frequency nadir is above the stability threshold across the network, a numerical estimation is applied on the integrals of post-fault frequency deviation [14, 15]. These integrals are expressed as a linear combination of key system variables, including system inertia, available reserves, system damping, and the size of the disturbance, as shown below:

$$\int_0^{t_{\text{Nadir},i}} \Delta f_i(t) dt = \sum_{j=1}^N (c_{j,i} H_{t_j} + d_{j,i} D_j + e_{j,i} R_j) + g_i \Delta P + m_i \quad (5.16)$$

where  $c_{j,i}, d_{j,i}, e_{j,i}, g_i, m_i$  represent the regression weights,  $R_j$  represents the available reserves,  $D_j$  represents the available damping, and subscript  $j$  represents the location out of the  $N$  areas of the network. The regression weights of the frequency nadir are calculated through a linear regression on training samples generated by RMS-TDS, focusing on those on the stability boundary. Further details of this approach can be found in [14, 15].

**ii) Iterative RMS TDS-based method (TDS-Opt):**

This method uses a two-stage approach to optimise power system operations and maintain stability [21]. In the first stage, AC-OPF is performed for an economic dispatch of the generators. In the second stage, RMS-TDS is conducted on the AC-OPF solution to assess the system's frequency stability. If the system is stable, the process terminates. If instability is detected, following the loss of the largest generating SG as the disturbance, the process enters an iterative loop, progressively tightening reserve requirements for each SG until stability is ensured. The reserve requirements for each

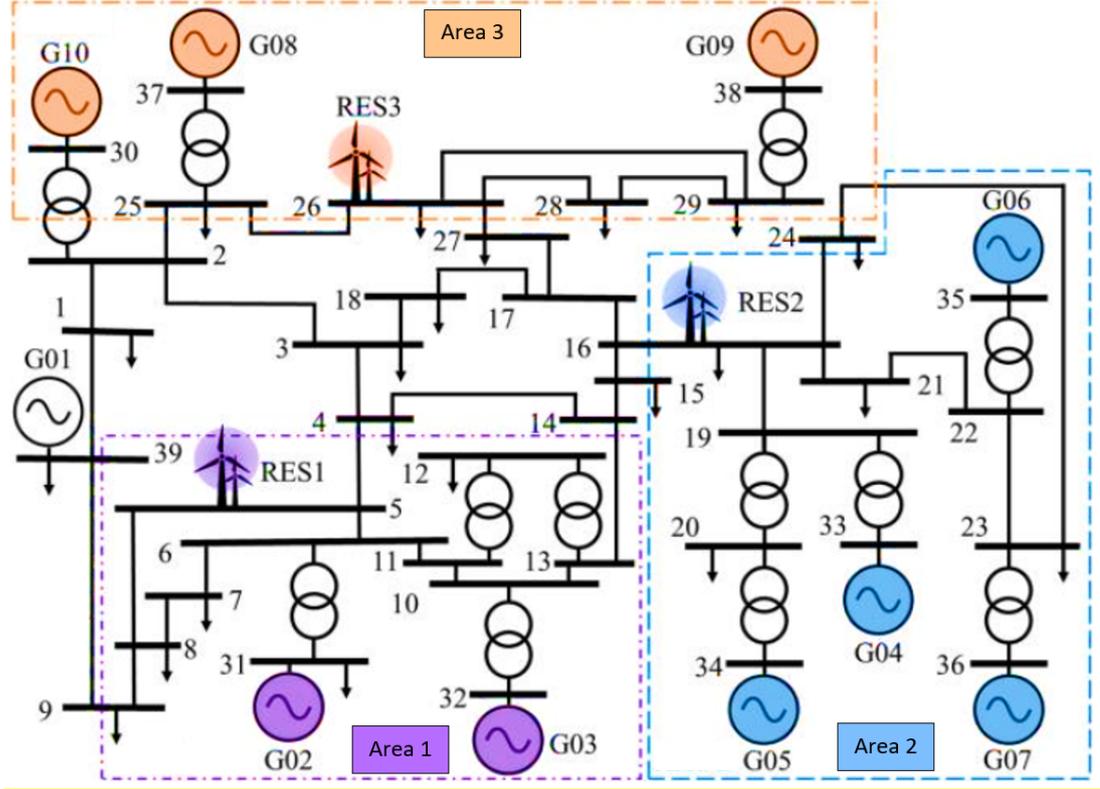


Figure 5.3: The modified IEEE 39-bus network highlighting the three network areas with CIG penetrations at buses 5, 16 and 26 [3]

SG in the network are calculated as shown below,  $\forall i \in I, i \neq l$ :

$$R_i \leq 2\rho_i \frac{M_H(f_0 - f_{\min} - f_{db})}{\Delta P_l} \quad (5.17)$$

where  $R_i$  is the reserve provision of SG  $i$ ,  $\rho_i$  is the  $i^{\text{th}}$  SG governor response,  $f_{db}$  is the dead-band frequency,  $I$  is the set of SGs, and  $\Delta P_l$  is the power outage due to SG  $l$ . Further details of this approach can be found in [75].

## 5.5 Case Studies and Results

In this chapter, the results of the proposed ANN-constrained OPF for locational frequency stability are presented using the modified IEEE 39 bus network shown in Fig. 5.3. Two case studies of varying complexity are considered as follows: i) *Case I*: a

case where the  $N - 1$  disturbance (loss of largest generation within the area) is in one location in the area with CIG (Area 2), and ii) *Case II*: a more complex case with CIG penetration in each of the three areas and the  $N - 1$  disturbance implemented in each area, i.e., the largest loss of generation for each area is considered as a disturbance. A single ML model is used in each case study, and the objective of the proposed ANN-constrained OPF is to generate solutions which, following a disturbance, do not have any frequency violation at any network location. Accuracy evaluation of the ANN is done using the metrics discussed in Chapter 3.3.3. The AC-OPF model is applied for generator dispatch using MIPS in MATPOWER [66], while the solver *find minimum of constrained nonlinear multivariable function* (fmincon) is employed in security-constrained dispatch due to its flexibility and handling complex constrained optimisation. Based on the scenarios generation criteria described in Chapter 5.4.2, the solutions from the optimisation are exported to DIGSILENT [17] using the Python programming language. RMS-TDS are then conducted on these scenarios to generate the locational frequency stability training and testing datasets for the ML models. All simulations were carried out on an 11th Gen Intel (R) Core (TM) i7-11700 @ 2.50 GHz with 16 GB installed RAM.

### 5.5.1 Observed Locational Frequency Deviations in the System Due to High CIG Penetration

Similar to locational frequency dynamics variations in Chapter 3.4.3, this chapter uses a different dataset described in Chapter 5.4.2, to demonstrate how COI-based methods are increasingly becoming unreliable in the face of high CIG generation. Thereafter, we test the proposed approach to demonstrate its ability to accurately capture these important locational frequency characteristics. As stated by (3.1) in Chapter 3.1.1, the accuracy of the conventional COI-based approach to represent the system frequency response inversely correlates with the deviation factor,  $A_i\omega_i(t)$ . An increase in the deviation factor implies that the frequency response is becoming more locational. Consequently, higher risks of locational frequency violations may exist. Using *Case II* as an example, Fig. 5.4 shows the locational RoCoF variations between any two locations

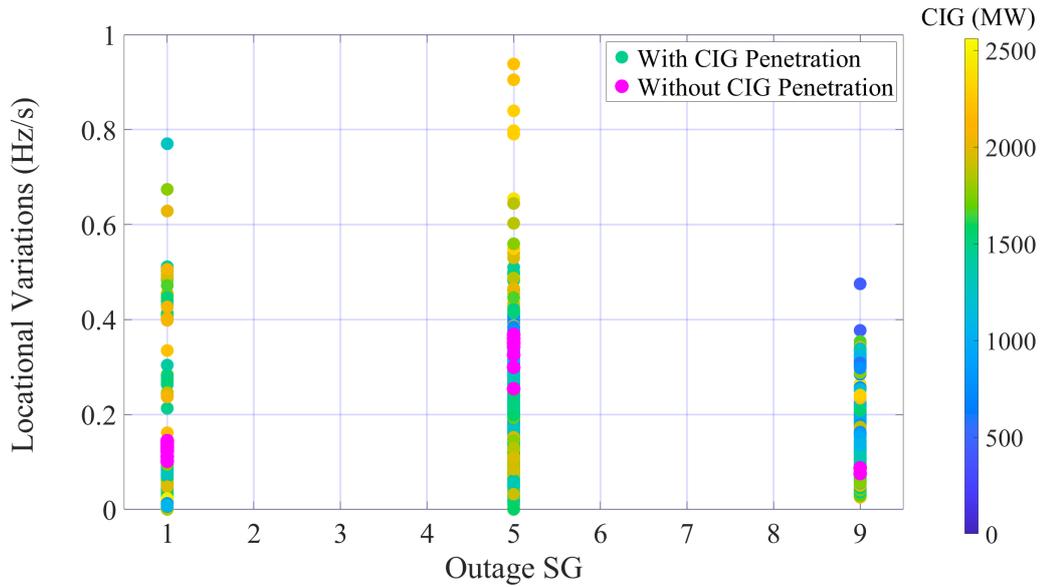


Figure 5.4: Impact of CIG penetration on the maximum locational RoCoF variations (in Hz/s) between any two locations in the modified IEEE 39 Bus network

across the network under two scenarios: i) *With CIG Penetration* and ii) *Without CIG Penetration*. The disturbances evaluated are the  $N - 1$  largest area-specific generating SG outages. The maximum variation for the scenario without CIGs is 0.15 Hz/s, 0.37 Hz/s, and 0.08 Hz/s for the disturbance outage of SG 01 (Area 1), SG 05 (Area 2) and SG 09 (Area 3), respectively. On the other hand, given different high CIG penetration levels across the network, the maximum variations significantly increase to 0.77 Hz/s (400% increase), 0.94 Hz/s (62% increase) and 0.48 Hz/s (500% increase) for the disturbance outage of SG 01 (Area 1), SG 05 (Area 2) and SG 09 (Area 3), respectively. Notably, this happens during instances of high CIG output, over 1,500 MW. Consequently, in systems with high CIG penetration, the significant variation in locational frequency dynamics can undermine the reliability of COI-based methods —potentially leading to unforeseen relay activations, as observed in Chapter 3.1.1. These results highlight that the widening gap between low-CIG and high-CIG systems necessitates more accurate and efficient modelling approaches to ensure secure system operation.

### 5.5.2 Case I: Fixed Disturbance Location With Single Area CIG Penetration

Operational scenarios are generated with the CIG integrated on Bus 16 in Area 2 across all scenarios. All SGs within the area, i.e., SG 4 - SG 7, are displaced systematically as stated by (6.15). Likewise, the scaling of the CIG follows the expression (6.16). Consequently, the penetration level was up to 40% of the overall system generation. The  $N - 1$  disturbance is the disconnection of SG 05, located in the same area as the CIG generation, contributing up to 25% of the area's generation. Overall, 2,114 operating scenarios are generated, and RMS-TDS is conducted to record the locational frequency nadir and RoCoF of the system following the disturbance. More details on dataset generation are given in Chapter 5.4.2.

#### i) Accuracy Evaluation of the Artificial Neural Network (ANN) Model

The ANN network is trained to predict the locational frequency stability metric, i.e., nadir and/or RoCoF, using the RMS-TDS dataset. The performance results of the trained model, whose parameters are determined by the GridSearchCV —a tool that automates the process of finding the best combination of hyperparameters for a given ML model by exhaustively searching over a user-defined parameter grid [57, 58], are summarised in Table 5.1. It is observed that the ANN can achieve a high degree of locational frequency stability metrics prediction accuracy by having a maximum RoCoF and nadir RMSE of 0.0021 Hz/s and 0.0095 Hz, respectively. It can also be observed that the mean RMSEs achieved by the model are 0.0020 Hz/s and 0.0093 Hz, respectively. This indicates that the model can consistently make accurate predictions, something that is an essential step before implementing it as a constraint in the system optimisation model. The ML model has adequately learned the locational frequency dynamics of the case study to generate solutions that adhere to frequency stability requirements at every bus in the network. The model parameters, i.e., weights, biases, and the training dataset standardisation parameters, are therefore saved and formulated into an OPF stability constraint using the method discussed in Chapter 5.4.3. The results of this approach are presented in the next subsection.

#### ii) Neural Network-Constrained Optimal Power Flow ( $\Delta - \text{OPF}$ ) Perfor-

Table 5.1: Accuracy of the Neural Network Model for locational frequency stability metrics prediction for Case I

<i>Metric</i>	<i>Max RMSE</i>	<i>Mean RMSE</i>	<i>Max MAE</i>
<b>RoCoF (Hz/s)</b>	0.0021	0.0020	0.0195
<b>Nadir (Hz)</b>	0.0095	0.0093	0.0489

Table 5.2: Model Optimisation Results for Locational Frequency Stability (Nadir) for Case I

<i>Performance</i>	<i>OPF</i>	<i>TDS-Opt</i>	<i>Numerical</i>	<i><math>\Delta</math>-OPF</i>
<b>Mean Iterations</b>	15.76	<b><u>8.20</u></b>	16.81	18.77
<b>Unstable OCs (%)</b>	82.98	<b><u>0.00</u></b>	<b>19.35</b>	<b><u>0.00</u></b>
<b>Mean OC Nadir (Hz)</b>	59.56	59.62	<b><u>59.84</u></b>	59.76
<b>Minimum Nadir (Hz)</b>	59.48	59.61	<b><u>59.75</u></b>	59.60
<b>Mean Opt. Secs. (MOS)</b>	<b><u>0.05</u></b>	57.39	<b><u>0.05</u></b>	0.06

### mance

The linearised model is embedded in the optimisation model as a locational frequency stability constraint, referred to as  $\Delta - \text{OPF}$ , following the procedure described in Sections 5.4.3 and 5.4.4. The modified optimisation model (as a security-constrained OPF model) is run to generate solutions which adhere to frequency stability requirements at every bus in the network. To validate this, RMS-TDS are conducted on the  $\Delta - \text{OPF}$  solutions, whereby the frequency stability metric of interest is recorded for every OC.

Using the frequency nadir as a stability metric and performance variable, the UFLS relay setting is considered to be 59.6 Hz. This threshold marks the activation point of the UFLS relays, indicating an unstable or insecure scenario. The performance of the  $\Delta - \text{OPF}$  is compared in Table 5.2 against the standard OPF without stability constraints (OPF), optimisation incorporating RMS-TDS in the loop (TDS-Opt) (5.17) and the numerical estimation model (Numerical) (5.16), as explained in Chapter 5.4.6. The numerical estimation method is applied across the generator buses in the network without the need to explicitly define network areas—leading up to ten stability constraints. It is observed that the mean and minimum nadir of the OPF model are 59.5613 Hz and 59.4834 Hz, respectively, with an overall nadir violation of 83%, i.e.

cases where the frequency causes any locational violation exceeding the pre-set threshold of 59.6 Hz in this chapter. The RMS-TDS has both the mean and minimum nadir above the 59.6 Hz threshold, bearing a 0% nadir violation. The Numerical estimation has the mean and minimum nadir of 59.8416 Hz and 56.4829 Hz, respectively, with an overall OCs violation of around 19%. The proposed  $\Delta$  – OPF achieves a mean and minimum frequency nadir of 59.76 Hz and 59.60 Hz, respectively, with no violations of frequency nadir limit observed in the simulated cases —indicating that the implemented frequency stability constraints can limit the frequency nadir within acceptable limits.

In Fig. 5.5, the extent of violations of the discussed models based on the results in Table 5.2 is presented. The figure summarises the nadir violations across a range of different CIG penetration levels. It can be seen in the figure that only two models, i.e., OPF and Numerical, bear violations with maximum values of 0.12 Hz. A positive correlation can be seen between the CIG penetration level and the nadir violations due to the reduced inertia and increased complexity introduced in the system. The Numerical method incurs violations only for OCs with the highest CIG penetration, i.e., at least 600 MW. On the contrary, the TDS-Opt and  $\Delta$  – OPF have no violation, i.e., 0 Hz.

### 5.5.3 Case II: Examining the Effect of Disturbances Across all Areas of the System

In the second case study, the modified IEEE 39 bus network is formed of three areas, as shown in Fig. 5.3. The generation of operating scenarios is based on the AC OPF as done in Chapter 5.5.2, considering the description in Chapter 5.4.2. The commitment of SGs is based on (6.15), and the scaling of the system loading is as described in Chapter 5.4.2. The locations of the CIGs within the areas remain fixed at the shown buses, i.e., Bus 5, Bus 16 and Bus 26 for Area 1, Area 2 and Area 3, respectively. The penetration of the CIGs in each area is determined by incrementally displacing the four equal-sized units of SGs as stated by (6.16). After an SG is completely displaced, the displacement of the next SG follows the same pattern, and the former SG reverts to

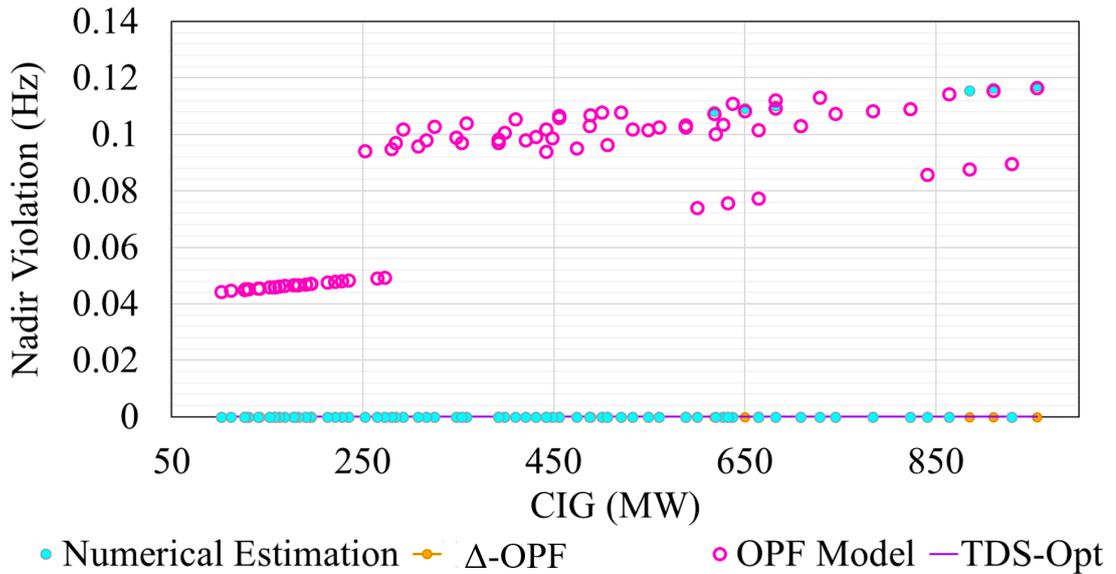


Figure 5.5: Locational Frequency Nadir Violations by four different models, including the proposed  $\Delta$  – OPF, across different CIG levels (up to 40% of system generation)

full capacity. Consequently, the total CIG penetration level with respect to the area generation for Area 1, Area 2 and Area 3 was between 4.5%-70%, 5%-60% and 6%-80%, respectively. Overall, CIGs generated between 6%-60% of the total system generation. For each OC, three RMS-TDS cases are performed for the loss of the largest generator in each of the three network areas. Consequently, the dataset for this case consisted of 4,785 OCs. The frequency nadir boundary is maintained as in *Case I* at 59.6 Hz.

**i) Accuracy Evaluation of the Artificial Neural Network (ANN) Model**

The ANN is trained to predict the locational (across 39 buses) frequency stability metric, i.e., nadir or RoCoF, using the RMS-TDS dataset. The accuracy of the ANN is then presented based on the buses, which are grouped into three areas, according to their designated areas shown in Fig. 5.3. Table 5.3 presents the RoCoF and nadir evaluations across the three areas. It can be observed that, similar to the performance in *Case I* (Table 5.1), the ANN maintains high inference accuracy. The worst mean RoCoF mean RMSE is in Area 2 at 0.0057 Hz/s, while the worst mean nadir RMSE is in Area 3 at 0.0208 Hz/s, thereby validating the model’s high accuracy performance requirements needed for the ANN-stability-constrained optimisation application.

**ii) Neural Network-Constrained Optimal Power Flow ( $\Delta$  – OPF) Perfor-**

Table 5.3: Accuracy of the Neural Network Model for Locational Frequency Stability Metrics Prediction for Case II

<i>Metric</i>	<i>Area</i>	<i>Max RMSE</i>	<i>Mean RMSE</i>	<i>Max MAE</i>
<b>RoCoF (Hz/s)</b>	Area 1	0.0224	0.0047	0.0168
	Area 2	0.0108	0.0057	0.0069
	Area 3	<b>0.0036</b>	<b>0.0033</b>	<b>0.0025</b>
<b>Nadir (Hz)</b>	Area 1	<b>0.0222</b>	<b>0.0205</b>	0.0159
	Area 2	0.0224	<b>0.0205</b>	<b>0.0156</b>
	Area 3	0.0225	0.0208	<b>0.0156</b>

### mance

The trained ANN is linearised and implemented as a constraint within the optimisation model for locational frequency stability, subject to any disturbance from any of the three areas. The performance of the neural network-constrained OPF using the Triangle relaxation of the ReLU ( $\Delta$  – OPF) is also compared with the Numerical method (Numerical) and the standard OPF without constraints (OPF). In this network with three distinct areas (Fig. 5.3), the Numerical estimation method is applied to each area, resulting in three corresponding stability constraints. The validation of each method is performed with RMS-TDS. Table 5.4 shows a summary of the performance of the three different methods following a disturbance in each area (implemented independently). The TDS-Opt, which includes RMS-TDS within its optimisation process, is not included in this table as it exhibits no frequency violation, with the associated significant computational requirements that were discussed in the previous section being the only disadvantage. The OPF model incurs 100%, 90% and 98% of frequency nadir violations for disturbances in Area 1, Area 2 and Area 3, respectively. This is, to some extent, expected since this approach is not security-constrained. The Numerical model demonstrates good performance for disturbances in Area 1 with total violations of around 1%, but its performance deteriorates for the disturbances in other areas, i.e., 76% and 97% of nadir violations in Areas 2 and 3, respectively. This is due to the varied behaviour of responses in these areas, which is difficult to capture by the regression model used in this approach. This can, however, potentially be improved by i) optimisation with different sets (or more) of data points within a varied distance

Table 5.4: Model Optimisation Results for Locational Frequency Stability (Nadir) for Case II

<i>Region</i>	<i>Performance</i>	<i>OPF</i>	<i>Numerical</i>	<i><math>\Delta</math>-OPF</i>
<b>Area 1</b>	Unstable OCs (%)	100	1.39	<b><u>0.00</u></b>
	Mean Nadir (Hz)	59.22	<b><u>59.77</u></b>	59.66
	Min Nadir (Hz)	59.02	59.59	<b><u>59.61</u></b>
<b>Area 2</b>	Unstable OCs (%)	91.67	76.39	<b><u>0.00</u></b>
	Mean Nadir (Hz)	59.08	59.56	<b><u>59.85</u></b>
	Min Nadir (Hz)	58.83	59.40	<b><u>59.60</u></b>
Area 3	Unstable OCs (%)	98.55	97.10	<b><u>0.00</u></b>
	Mean Nadir (Hz)	59.25	59.38	<b><u>59.90</u></b>
	Min Nadir (Hz)	59.04	59.18	<b><u>59.60</u></b>

from the stability boundary, ii) using bus-level constraints, rather than area-level, and iii) overestimating the predictions —though this may lead to reduced solution space. In contrast, the  $\Delta$ -OPF model demonstrates consistent adaptability across disturbances, by keeping the frequency within acceptable limits in all areas. However, it has an optimisation convergence of 63% due to over-approximating the ReLU activation function as described in Section 5.4.5. In the next Chapter, we present the neural network’s pre-activation statistics, then employ SAP- $\Delta$ -OPF to improve this convergence, while tolerating some risks.

#### 5.5.4 Neural Network Pre-activation Statistics and the Necessity of SAP- $\Delta$ -OPF

As described in Chapter 5.4.5, the Triangle relaxation can encounter convergence issues due to the high variance neurons. In Chapter 5.5.3,  $\Delta$ -OPF exhibited a 63% convergence rate because it is highly conservative; as a result, it reduces the solver’s feasible region. We assess this tendency by analysing the neural network’s pre-activation distributions. As shown in Fig. 5.6, the layer-wise mean pre-activation intervals are broad, especially for deeper layers 1 and 2 with ranges -60 to 10 and -30 to 10, respectively. These distributions are skewed towards 0 across all layers, where the gap between the ReLU and the relaxation is the widest (as depicted in Fig. 5.2). As a result, inference within this region is highly conservative, affecting the convergence rate. More gran-

Table 5.5: Model Optimisation Results for  $\Delta$ -OPF and SAP- $\Delta$ -OPF for Case II

Region	Performance	$\Delta$ -OPF	SAP- $\Delta$ -OPF
<b>Area 1</b>	Unstable OCs (%)	<u>0.00</u>	<u>0.00</u>
	Mean Nadir (Hz)	59.66	<u>59.93</u>
	Min Nadir (Hz)	59.61	<u>59.86</u>
<b>Area 2</b>	Unstable OCs (%)	<u>0.00</u>	1.67
	Mean Nadir (Hz)	<u>59.85</u>	59.69
	Min Nadir (Hz)	<u>59.60</u>	59.52
<b>Area 3</b>	Unstable OCs (%)	<u>0.00</u>	1.67
	Mean Nadir (Hz)	<u>59.90</u>	59.77
	Min Nadir (Hz)	<u>59.60</u>	59.53
<b>Computation</b>	Mean Iterations	<u>222.29</u>	386.22
	Mean Opt. Secs. (s)	5.44	<u>4.97</u>
	Convergence Rate (%)	63.89	<u>83.33</u>

ularity is provided in Fig. 5.7, where the distribution of neuron-wise minimum and maximum pre-activation bounds is plotted. In this figure, high variance neurons can be observed, for example, neurons 289, 107 and 249 (full pre-activation trace shown by Fig. 5.8) in layers 0, 1 and 2, respectively. These neurons are of special (critical) significance because they determine the overall tightness of the Triangle relaxation —largely determining the method’s conservativeness.

To improve the linearisation efficiency of the  $\Delta$ -OPF, and enhance optimisation convergence —while allowing some risks —we implement selective activation pruning (SAP- $\Delta$ -OPF) discussed in Chapter 5.4.5. We achieve this by identifying the most active and inactive neurons of the neural network using (5.15). A neuron is classified as most active if it has  $\Sigma_v^{(l)} \geq 0.95$ , i.e., the neuron was active at least 95% of the training time. In contrast, the most inactive neuron has  $\Sigma_v^{(l)} \leq 0.05$ , i.e., the neuron was only active less than 5% of the training time. After identifying these neurons, SAP- $\Delta$ -OPF is implemented on them. Results are shown in Table 5.5 using *Case II* data (Section 5.5.3), where  $\Delta$ -OPF managed to maintain the frequency within acceptable limits, but its convergence was only 63%, as mentioned before. In contrast, SAP- $\Delta$ -OPF has no violations in Area 1, but some 1.67% violations in Areas 2 and 3 —violating the acceptable threshold of 59.6 Hz by a maximum of 0.08 Hz and 0.07

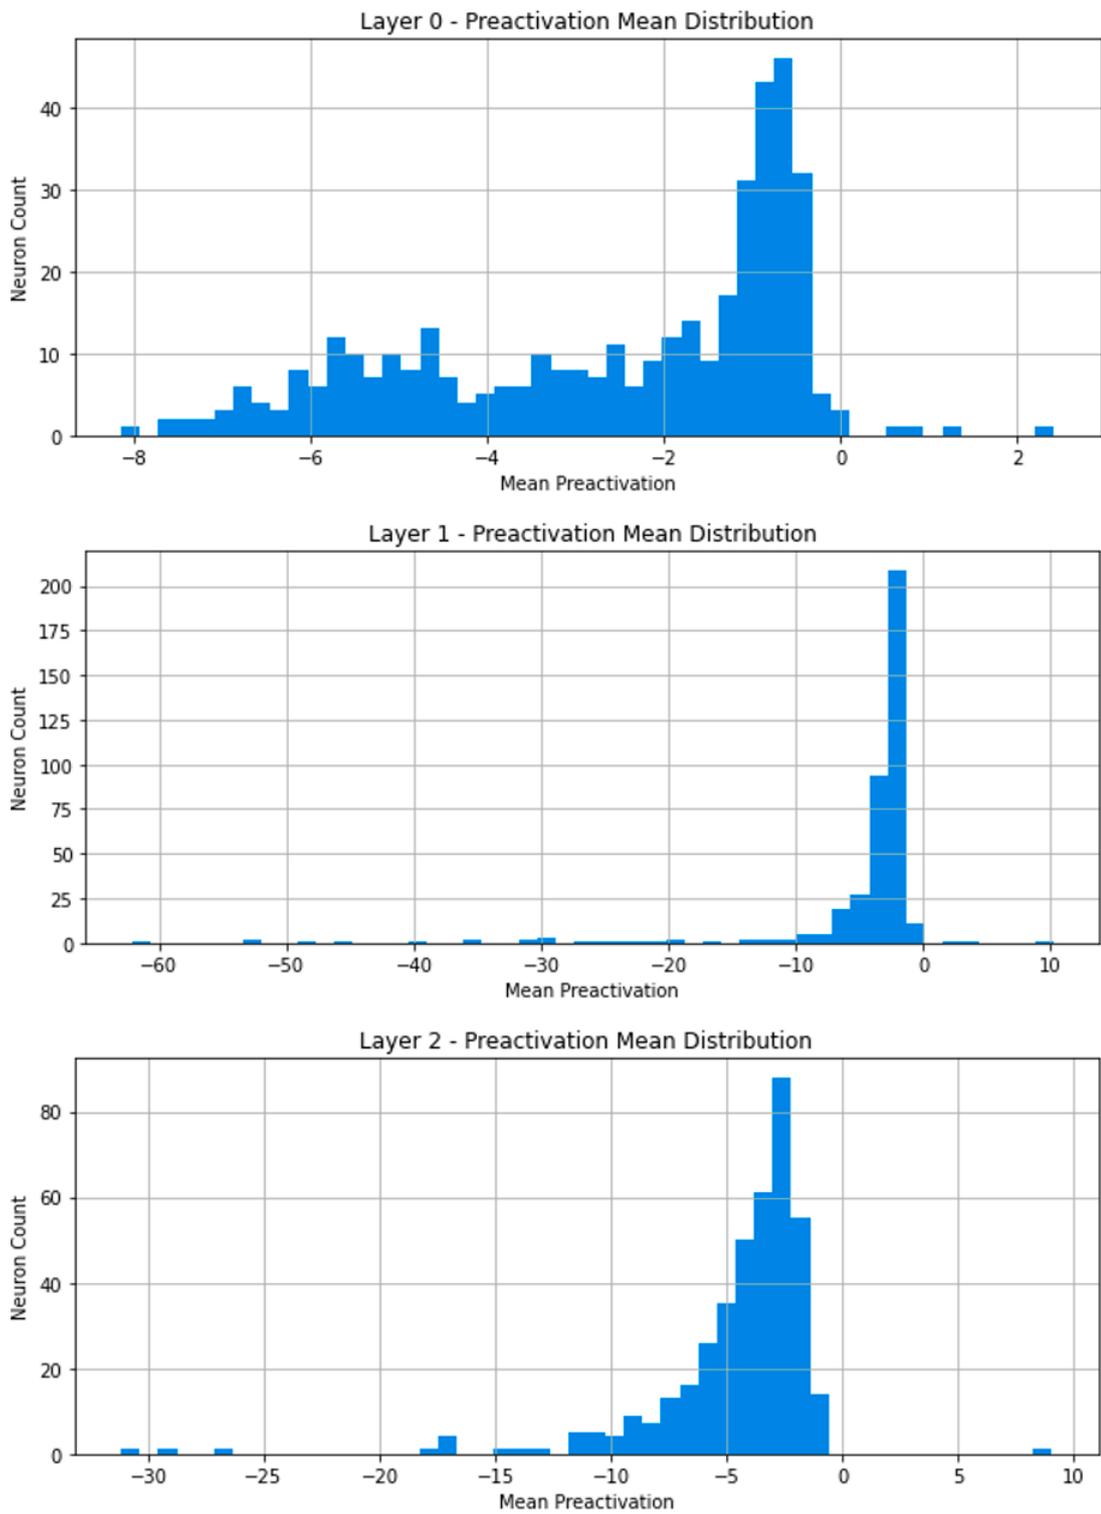


Figure 5.6: Layer-wise preactivation mean distribution of a neural network with three hidden layers

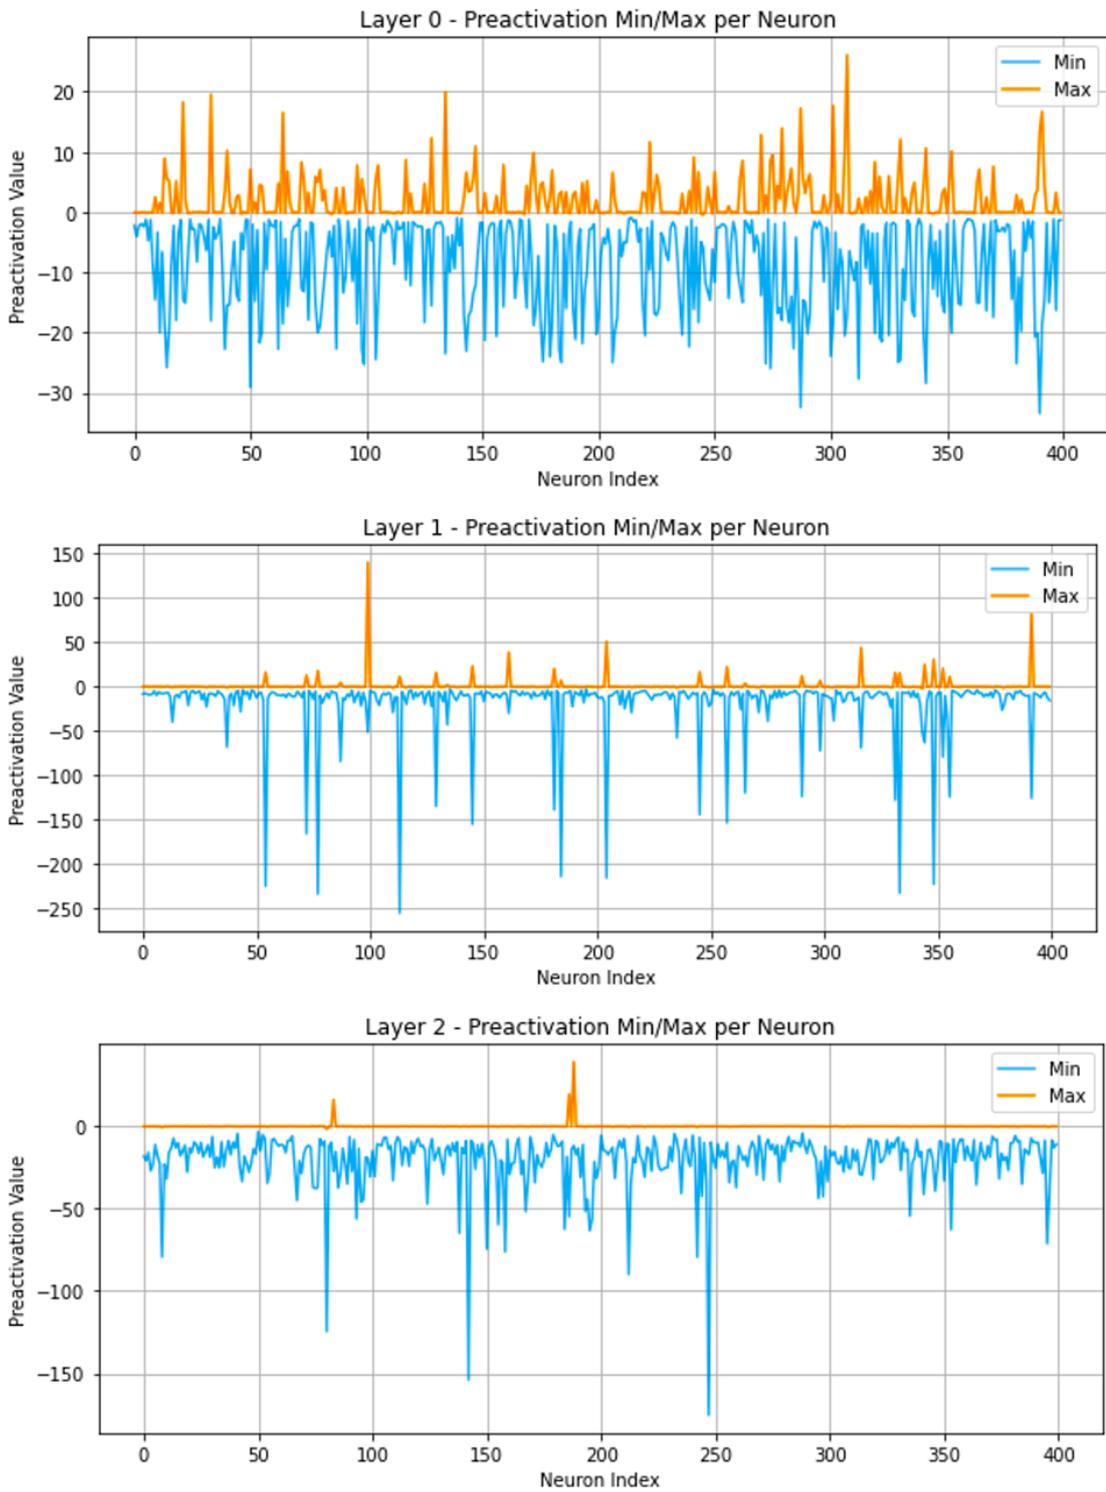


Figure 5.7: Neuron-wise minimum-maximum preactivation distribution of a neural network with three hidden layers and 400 neurons per layer

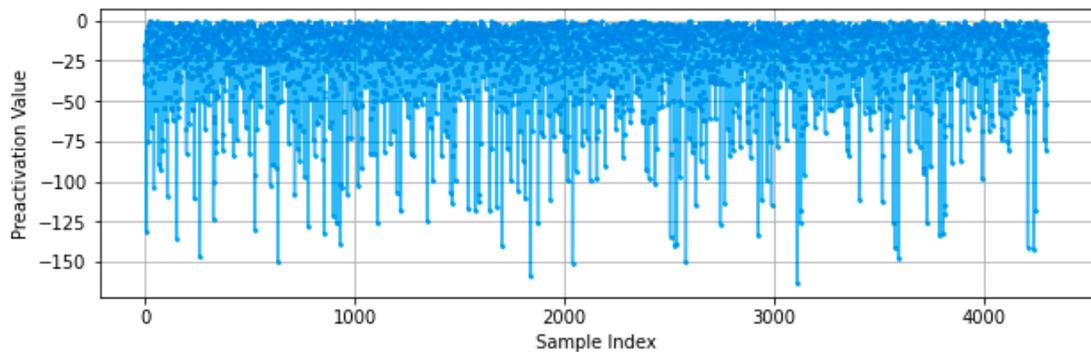


Figure 5.8: Full neuron preactivation trace for neuron 249 in Layer 2 (mostly inactive)

Hz in Areas 2 and 3, respectively. However, SAP- $\Delta$ -OPF significantly improves the optimisation convergence of  $\Delta$ -OPF by 32%, increasing it from 63% to 83% —indicating that selective neuron linearisation can be beneficial in scenarios where a degree of risk is acceptable. Furthermore, SAP- $\Delta$ -OPF reduces  $\Delta$ -OPF’s mean optimisation time (Mean Opt. Secs.) by 10%, that is, from 5.97 seconds to only 5.44 seconds. Such improvements would yield even greater benefits in larger-scale systems, where computational challenges and nonlinearity can be more pronounced. Note that while relaxing the boundary may improve  $\Delta$ -OPF’s convergence rate, it changes the nature of the problem and weakens theoretical guarantees. In contrast, SAP- $\Delta$ -OPF is a principled improvement —not a workaround —that preserves the original problem’s structure and improves convergence strictly within the existing feasible region.

### 5.5.5 Computational Considerations

As previously described in Chapter 5.5.2, both the  $\Delta$  – OPF and the TDS-Opt exhibit good performance. However, one of the key advantages of the  $\Delta$  – OPF is the improvement in computational time while maintaining the detailed locational frequency response characteristics. This is also one of the key motivations behind the proposed method. For example, concerning the computation performance in *Case I*, the OPF without stability constraints take the least amount of time with 0.05 seconds and 15.8 iterations per solution. This is followed by the Numerical estimation method (0.05 seconds, 16.81 iterations) and the  $\Delta$  – OPF (0.06 seconds, 18.77 iterations). TDS-Opt

takes the longest time per solution, up to around 57 seconds. This is due to the incorporation of the computationally intensive TDS within the optimisation loop, where several iterations,  $n$ , (8 on average), are required to find an optimal solution —thus 900 times slower than the  $\Delta$  – OPF. Despite the accuracy of this approach (and TDS in general), its application can be quite limited in real-time or close-to-real-time applications, especially for large-scale power systems. Furthermore, it was demonstrated in Chapter 5.5.4, using results from Case II (Chapter 5.5.3) that, where a degree of risk is acceptable, SAP- $\Delta$ -OPF (also proposed in this study) offers the potential to improve  $\Delta$ -OPF’s computational efficiency by 10%, further highlighting the potential of the proposed method in ensuring efficient (and secure) generator dispatch.

## 5.6 Conclusion

Performing TDS that accurately accounts for spatial frequency response considerations is a time-intensive process, making its integration into dispatch optimisation impractical. In this work, a machine learning-driven dispatch optimisation formulation is presented that models the increasingly important locational aspects of frequency response in systems with high penetration of Converter Interfaced Generation (CIG) while also offering the ability to account for the detailed dynamic behaviour of such devices. To this end, a Neural Network is trained to accurately predict the locational frequency characteristics, which is linearised and added as a security constraint to a typical Optimal Power Flow (OPF). The proposed method has the advantage of offering preferable qualities of computational speed while exhibiting high accuracy and no stability violations, comparable to approaches relying on much more computationally demanding iterative TDS or analytical approaches.

The proposed method was tested using two case studies on a modified version of the IEEE 39 bus network and compared against an OPF without frequency security constraints, a numerical approach introducing constraints into OPF and an iterative optimisation approach based on TDS. In the first case, results showed no frequency variations when using the proposed approach, compared to 83% violations in the non-constraint case and 19% violations in the case of the numerical approach. The second,

more complex case demonstrated consistent behaviour for more complex scenarios for disturbances across different areas of the test network, highlighting good performance against locational variations in the frequency response. While no violations can occur while using the TDS-based approach, the ANN-based proposed approach is up more than 900 times faster, enabling application in operational timescales.

Overall, the proposed method offers a combination of preferable characteristics of maintaining the detail of locational frequency response characteristics that time domain-based approaches have, combined with the speed of analytical-based approaches.

## Chapter 6

# Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

In Chapter 6, we address the challenges faced by conventional AUFLS methods in maintaining system stability, while minimising supply service interruption (load-shedding). Specifically, we focus on RL-based AUFLS approaches (AUFLS RL), which have shown significant advantages over analytical AUFLS methods. These advantages include their ability to learn even the ‘hard-to-model’ system dynamics through interactions with the system (simulation environment), while also offering good computational efficiency, allowing close to real-time control, as discussed in Chapter 2.4. We propose a Physics-Informed Deep RL technique for AUFLS by introducing a PS within the RL training. This shield utilises the system’s swing equation to guide the RL agent’s actions and eliminate non-essential actions. However, as the swing equation alone cannot adequately capture the increasingly complex locational/regional frequency dynamics, we introduce an additional step for coherence detection through the use of neural networks. This approach addresses a common limitation of most AUFLS RL methods in high-dimensional settings by providing a systematic dimensionality reduction technique, thereby facilitating practical adoption in large-scale power systems. It also increases the precision of the RL agent’s actions by enabling the agent to consider the OC-specific coherence

of buses within the network when optimising its decisions. The paper presenting this methodology is currently under review with IEEE Transactions on Power Systems as [Pub. F], building on the work [Pub. E], which was published in the proceedings of the 14th Mediterranean Conference on Power Generation, Transmission, Distribution and Energy Conversion (MEDPOWER 2024). The contribution of this work is threefold:

- i. We incorporate the governing physics of power system frequency dynamics into the RL training process through a PS to eliminate non-essential actions and subsequent fruitless exploration. This enhances both the optimality of solutions and the computational efficiency of current RL models, typically trained through excessive sampling.
- ii. We propose a neural network-based method for detecting coherence among participating frequency control loads, enabling RL agents to execute actions while considering scenario-specific aggregation. Unlike conventional geographic or static grouping, real-time identification of coherent areas enables efficient scalability in high-dimensional systems, increasing precision of the agent's actions.
- iii. We introduce a two-stage frequency protection architecture by integrating the conventional UFLS scheme with an RL-based control. The RL agent aims to minimise load shedding while preserving system stability, and the UFLS scheme remains fully armed as a last resort. This hybrid approach offers a robust control scheme that demonstrates the practicality of integrating RL-based strategies with legacy grid schemes.

## 6.1 Emergency Frequency Control in Power Systems

This chapter outlines the limitations of traditional UFLS schemes and discusses the challenges associated with existing AUFLS RL techniques proposed in the literature to address these limitations. The Under-Frequency Load Shedding (UFLS) scheme is a typical emergency frequency control strategy involving automated load shedding to

maintain power balance and frequency stability. The UFLS usually operates based on violation of pre-set frequency thresholds, triggering pre-defined shedding of a certain load percentage at various steps [45]. However, the assumptions and operational principles underlying this technique are increasingly being challenged in many ways. For instance, relying on pre-set thresholds assumes relatively predictable load and generation characteristics, and it neglects to account for the possibility of flexible loads [46,47].

### 6.1.1 Challenges of Existing Alternatives to the UFLS Scheme

Adaptive AUFLS has been proposed to address the limitations of the conventional UFLS, by using real-time system data to provide an adaptive load-shedding scheme [45, 48, 49]. However, most analytical approaches rely on simplifications to derive frequency responses and determine the appropriate load shedding. Consequently, their accuracy is increasingly challenged by the complexity and uncertainty introduced by CIGs in modern power systems [51]. Moreover, real-time computation of the load-shedding scheme, while considering the specific nature of the disturbance, can be computationally demanding. As a result, reinforcement learning-based AUFLS (AUFLS RL) has gained traction. RL models can learn optimal policies from observed or simulated data during training, allowing them to capture complex dynamics without explicit modelling of the system.

However, despite the significant advantages of existing AUFLS RL methods over analytical approaches as described in detail in Chapter 2.4, they are purely data-driven methods, with scaling challenges in high-dimensional settings. Consequently, they usually limit analysis to a few selected control variables. For example, these methods often select a few loads as load-shedding candidates in their control action space, or they arbitrarily cluster these loads to reduce the system's dimensionality. This practice reduces the effectiveness of these AUFLS RL techniques by neglecting the locational aspect of the participating loads. In addition, these methods ignore the physics governing the system's dynamic behaviour, leading to heavy reliance on large datasets to develop an optimal policy. Although it can be claimed that training datasets may capture the underlying physics, the explicit lack of it within the training loop contributes

to sub-optimal performance, long training time and difficulty in extrapolating beyond the training domain.

### 6.1.2 Enhancing the Effectiveness of Existing AUFLS RL Methods

In response to the aforementioned challenges, we propose a Physics-Shielded Deep RL technique for AUFLS. The method introduces a “Physics Shield” (PS) within the RL training, which incorporates the system’s swing equation to guide the agent’s learning process. This integration constrains exploration by eliminating non-essential load shedding actions that unguided and purely data-driven RL agents might otherwise waste considerable time evaluating. However, while the PS captures the key dynamics through the swing equation, it is insufficient to account for the impact of the spatial distribution of loads—a critical factor for effective load shedding. To address this, we therefore introduce real-time system coherence detection using the rate of change of frequency (RoCoF). Unlike traditional model-based coherence detection techniques, which often struggle with clustering non-SG-generating buses, our approach applies a data-driven technique using neural networks (NNs). This method imposes relatively little to no computational overhead, making it suitable for online applications [41]. By adaptively clustering the system based on real-time dynamic behaviour, the proposed method identifies *where*, *when*, and *how much* load is to be shed more efficiently. Finally, given that data-driven approaches can suffer accuracy bias resulting from limited training datasets and/or out-of-distribution edge-case scenarios, we integrate the RL agent within the existing UFLS protection architecture as a last resort, providing an extra layer of security.

## 6.2 Deep Reinforcement Learning Implementation

In this chapter, we describe deep RL and how we formulate our problem of AUFLS to keep the power system frequency within acceptable limits, through the RL agent’s location- and time-specific load shedding actions.

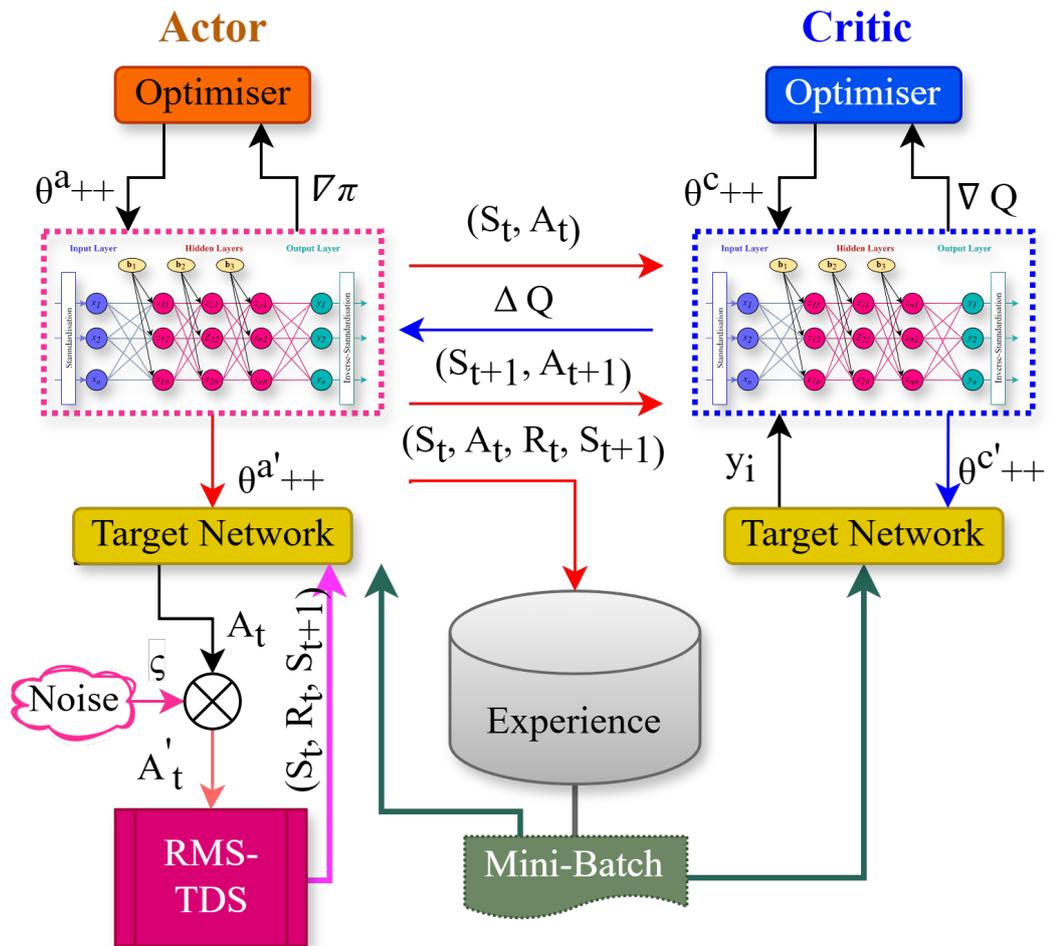


Figure 6.1: Architecture of the DDPG algorithm illustrating the actor and critic networks and the simulation environment (TDS)

### 6.2.1 Problem Formulation

RL uses Markov’s Decision Process (MDP) to implement sequential strategic decisions [79], where agents aim to maximise the reward, they get from their interactions with a given environment. By balancing between exploitative and exploratory actions, RL agents learn and establish the complex relationships of the environment. This is typically achieved using a  $Q$ -value [80], that is, an action value that tells an agent how good it is to take a certain action in a given state. The Deep Deterministic Policy Gradient (DDPG) RL algorithm [81], shown in Fig. 6.1, employs an *actor-critic* architecture. It uses two NNs: the *Actor* network, which is responsible for interacting with the environment and learning the policy mapping,  $\pi_\theta$ , between the current state and a given action; and the *Critic* network, which guides policy improvement by estimating the  $Q$ -value of the action chosen by the actor. Furthermore, the algorithm employs target actor and target networks, which are slow-moving copies of the actor and critic networks, with the primary purpose of stabilising the algorithm’s training [50, 80]. We employ the DDPG algorithm because it is well-suited for continuous action multidimensional spaces [50], making it appropriate for our implementation.

To initialise training and properly guide interactions with the environment, two key parameters are specified, i.e.,  $\epsilon_{\text{threshold}}$ , an exploration threshold, and  $\eta$ , sampled from a normal distribution  $[0, 1]$ . During training, if  $\eta > \epsilon_{\text{threshold}}$ , the agent’s action is exploitative and vice versa. If  $\eta \leq \epsilon_{\text{threshold}}$ , the action is exploratory, allowing the agent to select any expendable action in the search space. In addition, a decaying Gaussian noise,  $\zeta$ , is added to the agent’s actions to prevent the agent from getting trapped during training. During such interactions, the target actor and the critic networks learn and update (copy) their parameters as follows:

$$\left( \theta^{a'} \leftarrow \tau\theta^a + (1 - \tau)\theta^{a'}, \quad \theta^{c'} \leftarrow \tau\theta^c + (1 - \tau)\theta^{c'} \right) \quad (6.1)$$

where  $\theta^a, \theta^c$  are parameters of the actor and critic networks, respectively, while  $\theta^{a'}, \theta^{c'}$  represent the parameters of the target actor and target critic networks, respectively. Thus, the actor and critic are the main networks actively being trained to improve the agent’s performance. The parameters of the target networks are updated at a rate

related to  $\tau$ , the role of which is to prevent large updates, thus ensuring stable training. The aim is to regulate and ensure that the target networks are updated gradually to match the main networks. As such, a smaller  $\tau$  leads to slower updates, whereas a larger  $\tau$  leads to faster updates [50, 55]. The critic network evaluates and guides the actor in selecting optimal actions, allowing it to learn the optimal state-action mapping through the loss function:

$$J(\theta^c) = \frac{1}{m} \sum_{j=1}^m (y_j - Q(s^j, a^j, \theta^c))^2 \quad (6.2)$$

where  $m$  is the number of samples of the mini-batch (a subset of the stored transitions or experience),  $Q(s^j, a^j, \theta^c)$  is the predicted  $Q$  value from the critic for the *state-action* pair  $(s^j, a^j)$  using the current critic network parameters,  $\theta^c$ .  $y_j$  is the target or the real  $Q$  value for the  $j^{\text{th}}$  sample, and is calculated from the agent's reward function as follows:

$$y_j = \begin{cases} r & \text{for } i_{\text{step}} = 1 \\ r + \gamma Q(s^j, a^j, \theta^c) & \text{for } i_{\text{step}} > 1 \end{cases} \quad (6.3)$$

where the value of the action,  $y_j$ , equals the immediate allocated reward,  $r$ , if the episode ends during the first step,  $i_{\text{step}}$ . For all subsequent steps up to the last step, the cumulative reward reflects intertemporal relationships by being weighted by the discount factor  $\gamma$ , assigning a lesser value to future rewards. Thus,  $r$  is specially designed and is critical to the optimal learning and development of the agent's policy,  $\pi_\theta$ . In our implementation, the scenarios are modelled independently, as each is a separate TDS. The design choice for this function is described in Section 6.3.3. The loss function of the actor network is stated as follows:

$$J(\theta^a) = -\frac{1}{m} \sum_{j=1}^m Q(s^j, a^j, \theta^c) \quad (6.4)$$

where  $\theta^a$  are the actor's current parameters being optimised to minimise the overall loss,  $Q(s^j, a^j, \theta^c)$  is the  $Q$  value predicted by the critic (with parameter set  $\theta^c$ ) for the action  $a^j$  chosen by the actor. This function aims to minimise the negative  $Q$ -value of the actor's chosen actions across the minibatch,  $m$ . Its minimisation is equivalent to maximising expected return, so that the critic estimates higher values for the actions chosen by the actor. Therefore, in the context of this work, the objective of the RL agent is to maintain the frequency nadir after large disturbances (e.g.  $N - 2$ ) with

minimal load-shedding, by controlling *when*, *where*, and *what amount* of distributed loads is shed across the network. The state and action spaces of the environment are designed as follows:

**i) State ( $s$ )**

This represents the current status of the environment, consisting of system variables relevant to the given objective. The state space is crucial for the agent’s learning as it defines the available knowledge of the system. For our problem, we define  $s = [\rho, \mathbf{s}_g, \mathcal{P}_g, \mathcal{L}]$ , where  $\rho$  is the system’s coherence detection vector, containing bus indices labelled according to the coherence cluster, out of coherence clusters  $K$ . This grouping allows the RL agent to identify coherent buses and treat them as a single controllable entity, thereby improving both the precision and efficiency of the control strategy through dimensionality reduction. This is essential to capture detailed locational frequency dynamics of the system’s current state (Section 6.3.2 describes this in detail).  $\mathbf{s}_g$  is a vector of generator ratings,  $\mathcal{P}_g$  is a vector of active power set points of generators, and  $\mathcal{L}$  is the total system load scaling.

**ii) Action ( $a$ )**

At each state, the RL agent generates a continuous action vector of size  $K + 1$ , where the first  $K$  elements correspond to the AUFLS applicable to  $K$  pre-defined clusters of the system, while the additional element is the load shedding activation time (detailed in Section 6.3.2). Thus,  $a$  is executed in a single step as the optimal load shedding scheme,  $\xi_i, \forall i \in K$ , with a unified activation time,  $t_{\text{Act}}$ , i.e.,  $a = [\xi_i, t_{\text{Act}}]$ . The total load in the system available for frequency control is divided into two equal-sized groups: *flexible loads*,  $\mathcal{D}_{\text{flexible-loads}}$  —controlled by the RL agent through the RL-based scheme  $\xi_i, \forall i \in K$  —and *defensive loads*,  $\mathcal{D}_{\text{defensive-loads}}$ , whose protective relays remain fully armed —managed by the conventional UFLS scheme as a last resort.

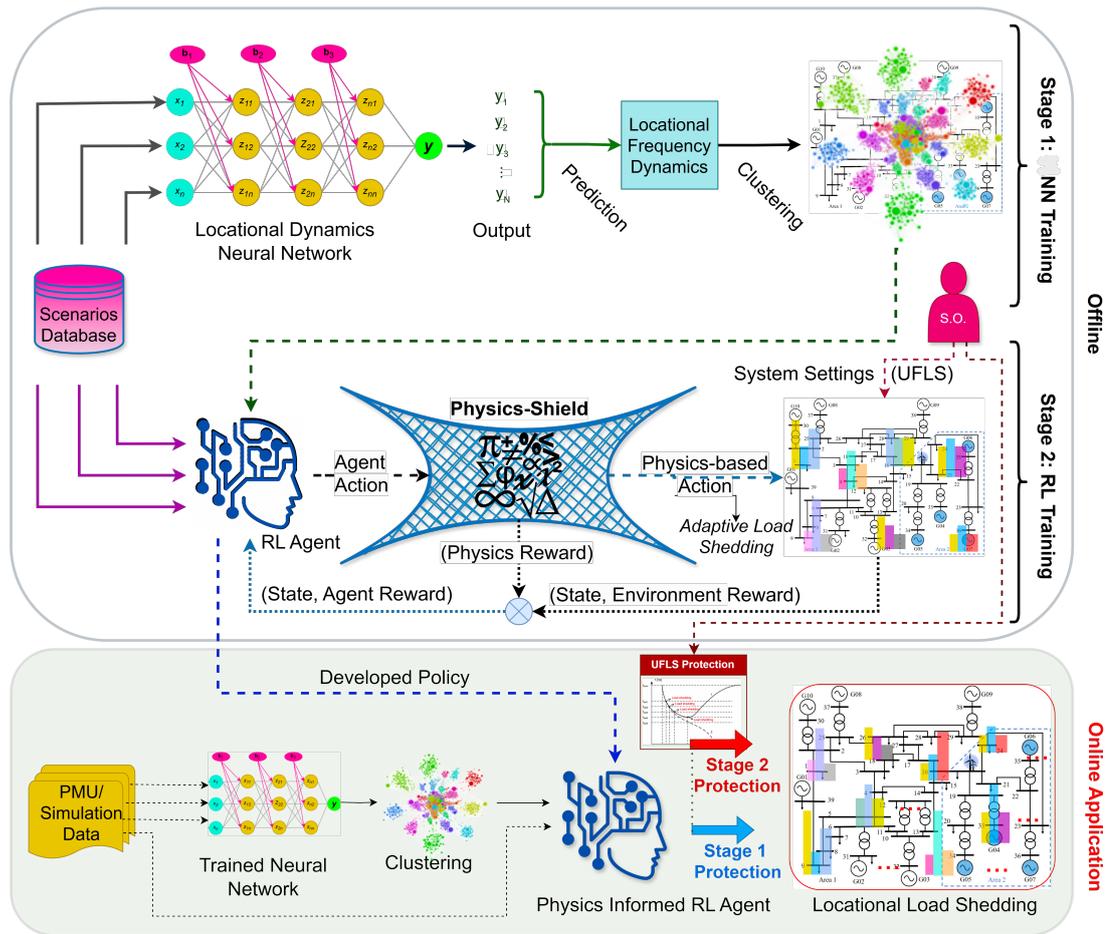


Figure 6.2: Overview of the NN-Physics-shielded RL agent implementation for adaptive emergency frequency control

### 6.3 Methodology for the Physics-Informed RL-based AU-FLS for Frequency Stability

As modern power systems grow in size, complexity, and uncertainty, avoiding local frequency violations throughout the system becomes more challenging. While RL-based AUFLS has shown promise, existing methods are mainly data-driven, which inherently limits their scalability and efficiency due to the excessive sampling required during training to explore the high-dimensional state-action space of complex power systems. We address this limitation through a physics-shielded approach that incorporates the system’s physical laws and relationships directly into the RL training process. By embedding domain knowledge through the PS, we constrain the search space and filter out non-essential actions that existing RL models might otherwise spend considerable time exploring, enabling faster and optimal convergence. This is summarised in Fig. 6.2, which illustrates the Offline (with Stages 1 and 2) and Online Application phases. To facilitate efficient training of the RL agent, providing it with information on the locational frequency characteristics of the system is essential. To this end, we introduce a method based on NNs and clustering to reduce the dimensionality of the action space. Instead of treating each PQ load bus hosting a flexible load as a separate variable, the RL agent can collectively control a group of loads in locations with similar frequency response behaviour. In the Offline phase, both the NN and the RL agent train by utilising physical and steady-state operational variables of the system from the training dataset generated by TDS, and the RL agent interacts with the same simulation environment. In Stage 1, an NN is trained to predict detailed locational frequency dynamics (i.e. RoCoF). Then, an unsupervised algorithm (K-Means) uses these predictions to cluster coherent buses of the network into  $K$  groups. In Stage 2, where the Physics-Shielded RL agent is trained, we introduce a PS that uses the system’s swing equation to systematically constrain the action space and eliminate actions with apparent violations. However, since this alone is insufficient to account for the locational aspects of efficient load shedding, we allow the agent to act using information from Stage 1, where coherent buses are clustered, and interact with the environment (phasor

domain simulation environment in our case) to learn detailed dynamics of locational frequency. This enables the agent to treat each coherent group as a single entity, effectively reducing the problem’s dimensionality in a structured manner. The PS filters the agent’s actions, ensuring that only essential actions are executed in the simulation environment (a computationally expensive action), and physics-inconsistent actions are penalised (*Physics Reward*). Furthermore, the agent receives an *Environment Reward* based on the system’s frequency response (from detailed phasor domain simulation) to its actions. After training, the Online Application phase is meant to be used with real measurements during close-to-real-time system operation (e.g. from PMUs) and employs the AUFLS RL agent trained in the Offline phase. The RL agent can be used alongside the traditional UFLS scheme based on pre-set thresholds, which can still be used as a last resort. The AUFLS RL scheme is envisaged to work with predefined flexible loads,  $\mathcal{D}_{\text{flexible-loads}}$ , able to respond to a signal from the RL agent to disconnect, providing the locational granularity needed, as further described by **Algorithm 1**.

---

**Algorithm 1** NN-Physics-Shielded RL for AUFLS Emergency Frequency Control

---

**Inputs:** System operating condition: bus clustering labels ( $\rho$ ), generator ratings ( $\mathbf{s}_g$ ), generator active powers ( $\mathcal{P}_g$ ), total system load scaling ( $\mathcal{L}$ )

**Output:** Load shedding amount (for specific loads in different locations) ( $\xi_i$ ), and activation time (after the disturbance) of the AUFLS RL ( $t_{\text{Act}}$ )

**Stage 1: NN Training:**

- 1: train neural network (NN) to predict bus-level RoCoF from pre-disturbance features
- 2: evaluate NN accuracy (RMSE and  $R^2$ )
- 3: **if** NN accuracy test passed, continue; **else** retune and retrain
- 4: initialise K-Means with  $K \leftarrow 2$
- 5: **while** true:
  - apply K-Means on predicted RoCoF; compute maximum intra-cluster inertia
  - if** intra-cluster inertia  $>$  5% of RoCoF threshold:  $K \leftarrow K + 1$
  - else**: store  $K$  as the cluster number of the power system, **break**

**Stage 2: RL Training:**

- 1: **for** state space,  $s$ , in training dataset **do**
  - 2: trained NN(s) predicts bus level RoCoF across PQ load buses of the system
  - 3: K-Means clusters the predicted RoCoFs into  $K$  coherent groups
  - 4: **update**  $s \leftarrow$  bus indices labelled according to the coherence clusters
  - 5: **for**  $j_{\text{episode}} \in [1, \dots, \mathcal{E}_0]$  episodes **do**
  - 6: **if** sampled exploration factor  $\eta > \epsilon_{\text{threshold}}$
  - 7: exploit agent's learned actions
  - 8: **select** action  $a$  based on the updated state  $s$
  - 9: **clip** action  $a$  + decaying Gaussian noise within  $[0,1]$
  - 10: **pass** action  $a$  through Physics Shield  $\Omega_{\text{PI}}$
  - 11: **get** Physics Shield-based reward,  $r_{\Omega_{\text{PI}}}$ , update  $a$
  - 12: **initialise** RMS-TDS and apply the disturbance
  - 13: **update** total agent reward  $r_{\text{RL}} \leftarrow r_{\Omega_{\text{PI}}} + r_f$
  - 14: store system transition in memory buffer,  $\mathcal{G}$
  - 15: **if** frequency nadir within safe bounds, i.e.,  $s_{\text{Metric}} < \mathcal{H}_{\text{Hz}}^{\text{limit}}$
  - 16: **break** (exit current episode loop)
  - 17: sample mini-batch  $\mathcal{B}$  from  $\mathcal{G}$ , update  $\theta^a, \theta^{a'}, \theta^c, \theta^{c'}$
  - 18: **if**  $s, a$  remain unchanged over  $\Psi$  interactions
  - 19: store scheme  $[\xi_i, t_{\text{Act}}], \forall i \in \mathcal{D}_{\text{flexible-loads}}$
  - 20: **break** (exit full training loop)
- 

### 6.3.1 Capturing Locational Frequency Dynamics with Artificial Neural Network

Traditional analytical approaches often struggle to capture locational frequency characteristics, especially in systems with a lot of CIGs, due to high complexity [14]. We circumvent this challenge using a data-driven neural network approach [82]. The neural

network (Multi-layer Perceptron (MLP)) [57] is trained to capture bus-level RoCoFs in a system, focusing on all PQ load buses since we aim to subsequently cluster those buses based on similarities in frequency response characteristics. This, in turn, will allow more granular and efficient frequency control by shedding loads in the most effective locations. The MLP uses an input vector consisting of the total system load scaling factor, CIG active power output, active power dispatch of SGs, and the MVA ratings of SGs. For effective hyperparameter tuning, we automate the process using sklearn-GridSearchCV [57], which performs an exhaustive search over a user-defined parameter range, while applying  $k$ -fold cross-validation (with 5-fold and 0.2 validation fraction), i.e., evaluating the model's performance across multiple  $k$  data splits. This enables us to systematically identify the number of hidden layers, the number of neurons per layer, the activation function, the model's alpha—which regulates the size of the model's weights—the learning rate, etc., without biases arising from manual tuning. This ensures that MLP is both expressive and generalisable. We adopt a 70%-30% dataset split for training and testing, respectively, and use the root mean squared error (RMSE) and coefficient of determination  $R^2$  for performance evaluations. Once sufficiently trained, the MLP enables real-time RoCoF-based coherence detection with minimal computational overhead, as detailed in the next subsection.

### 6.3.2 RoCoF-Based Adaptive System Coherence Detection

Coherence detection allows power system operators to identify system buses that exhibit similar dynamic characteristics, typically to reduce dynamic models [41]. In the context of this work, coherence detection aims to identify locations with similar frequency response (RoCoF) characteristics. The aim is to improve scalability by improving sampling efficiency while training the RL agent, which would otherwise be challenging in modern high-dimensional power systems. We adopt a data-driven approach to coherence detection, which offers several benefits over traditional model-based methods. Traditional coherence detection methods can be computationally intensive for real-time applications. Furthermore, these methods typically linearise the system matrix around an operating point, capturing only small disturbances around that point. In contrast,

our ML-based approach (described in Chapter 6.3.1) improves efficiency and is suitable for online application, since once trained, it only utilises (once) matrix multiplication of learned parameters to predict the system’s dynamic response after a disturbance. However, while training the model (including the generation of training datasets) can incur significant time, it is a one-time process done offline, where time is not critical. Once trained, the model can almost instantly perform predictions on streaming PMU data, enabling scalable, near real-time dynamic adaptation to changing operating conditions, with respect to RoCoF, and subsequent clustering of flexible PQ load buses for dimensionality reduction. Therefore, using the RoCoF predictions described in Chapter 6.3.1, we apply the K-Means [57], an unsupervised ML technique, to group coherent buses. K-Means assigns these buses to  $K$  clusters by minimising the inertia within the clusters, thereby grouping buses with similar RoCoF. This clustering enables each group of buses to be treated as a single controllable entity, improving the precision and efficiency of the control strategy. Each cluster is described by the mean of its samples, or centroid. Mathematically, the objective of K-means is to minimise the cluster inertia as follows:

$$\sum_{i=1}^n \min(\|\hat{y}_i - \mu_j\|^2), j = [1, \dots, K] \quad (6.5)$$

where  $\hat{y}_i$  is the  $i^{th}$  data point within the vector of the frequency dynamics predictions by the ML model (NN), across  $n$  PQ buses.  $\mu_j$  is the centroid of the  $j^{th}$  cluster out of the  $K$  number of clusters. For a given power system,  $K$  is determined by ensuring that the maximum inertia within each cluster is less than 5% of the stability threshold (by iteratively adjusting the value of  $K$ ). For example, in a system with a RoCoF stability boundary of  $-0.5$  Hz/s, the maximum permissible clustering MSE would be  $0.025$  Hz/s. Using this rule, we consider the clustering of buses with similar frequency dynamics to be acceptable.

### 6.3.3 Physics Shield and Rewards Integration of the RL Agent

In this chapter, the PS is proposed to regulate the actions taken by the RL agent, ensuring that they abide by the physical laws or relationships governing the system's frequency dynamics. As shown in Fig. 6.2, the PS filters all actions taken by the agent before being executed in the system. The agent is rewarded if the action taken lies within the PS boundaries to encourage consistency, and vice versa. The agent is penalised if the action lies beyond the PS to discourage similar actions in the future. Using the swing equation to represent the foundational system's frequency dynamics, we systematically guide and eliminate non-rewarding actions.

#### i) Physics Shield Structure

As a reminder, the agent has access to a fully detailed phasor time domain simulator to evaluate the actions in detail and learn the optimal policy. Therefore, the effectiveness of the agent's actions is evaluated based on the agent's ability to identify a combination yielding the least amount of load shedding while ensuring that the frequency remains within allowed limits. The PS is designed to apply the swing equation of the power system as a starting point to model emergency frequency control through load shedding. Following a disturbance, the system responds to the resulting power imbalance,  $\Delta P(t)$ , as follows:

$$\Delta f(t) = \int_{t_0}^{t_{\text{Nadir}}} \frac{df(t)}{dt} dt = - \int_{t_0}^{t_{\text{Nadir}}} \frac{\Delta P(t) f_n}{2 \sum_{i=1}^N s_i H_i} dt \quad (6.6a)$$

$$\Delta f(t_{\text{Nadir}}) = - \left( \frac{\Delta \mathcal{E}(t_{\text{Nadir}})}{2 \sum_{i=1}^N s_i H_i} f_n \right), \quad \text{where } \Delta \mathcal{E}(t_{\text{Nadir}}) = \int_{t_0}^{t_{\text{Nadir}}} \Delta p(t) dt \quad (6.6b)$$

where  $t_0$  and  $t_{\text{Nadir}}$  are the disturbance occurrence time ( $t = 0$ ) and the time when the frequency nadir occurs, respectively.  $\Delta \mathcal{E}(t_{\text{Nadir}})$  is the accumulated energy imbalance in the system following the disturbance and any subsequent load shedding actions. This can also be represented by mapping it to the summation of the disturbance,  $\Delta P_l$ , and the amount of load shed,  $\psi$ , weighted by a factor  $\varphi$ , which captures the 'hard-to-model' nonlinear relationships in the system, as follows:  $\varphi((-\Delta P_l + \psi)t_{\text{Nadir}}) \rightarrow \Delta \mathcal{E}(t_{\text{Nadir}})$ .

Thus,  $\varphi$  is effectively a simple approximation, a system-specific and operating condition-dependent variable that linearises a more complex relationship in reality, mapping the disturbance and any subsequent load-shedding actions before  $t_{\text{Nadir}}$  onto an effective system imbalance influencing generators' dynamic behaviour. The PS uses this to constrain the action space bounds, as described below.

To train the agent more efficiently, we constrain the action space within some permissible load-shedding range, between  $\psi_{\min}$  and  $\psi_{\max}$ . These theoretical limits eliminate apparent nonessential actions to speed up training. The maximum limit ensures that the agent does not explore actions which could lead to over-frequency and/or positive RoCoF. In contrast, the minimum limit ensures that the agent does not explore actions which could fail to keep the frequency within acceptable limits. While there is no direct knowledge of the exact limits, the RL agent must discover the optimal actions; apparent non-essential actions can delay and/or prevent the agent from discovering these optimal actions. Therefore, eliminating them from the action space improves convergence. We define this range by associating load shedding with a frequency response range as follows:  $\Delta f^{\max}$ , for maximum frequency decline and  $\Delta f^{\min}$ , for minimum frequency decline. Accordingly, (6.6b) is reformulated to define the PS's upper and lower bounds:

$$\left( \frac{\Delta f^{\max}}{f_n}, \frac{\Delta f^{\min}}{f_n} \right) = \left( \frac{\varphi((-\Delta P_l + \psi_{\min})t_{\text{Nadir}})}{2 \sum_{i=1}^N s_i H_i}, \frac{\varphi((-\Delta P_l + \psi_{\max})t_{\text{Nadir}})}{2 \sum_{i=1}^N s_i H_i} \right) \quad (6.7)$$

Thus, the RL agent is constrained to explore and optimise within the permissible region delineated by (6.7), as follows:

$$\psi_{\max}^{\Omega_{\text{PI}}} = |(\psi_{\max} - \Delta P_l)t_{\text{Nadir}}| = \frac{2 \sum_{i=1}^N s_i H_i}{f_n} \times \frac{\Delta f_{\text{COI}}^{\min}}{\varphi_{\max}} \quad (6.8a)$$

$$\psi_{\min}^{\Omega_{\text{PI}}} = |(\psi_{\min} - \Delta P_l)t_{\text{Nadir}}| = \frac{2 \sum_{i=1}^N s_i H_i}{f_n} \times \frac{\Delta f_{\text{COI}}^{\max}}{\varphi_{\min}} \quad (6.8b)$$

where  $\psi_{\max}^{\Omega_{\text{PI}}}$  and  $\psi_{\min}^{\Omega_{\text{PI}}}$  are the resultant maximum and minimum imbalances determined

## Chapter 6. Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

by the PS due to the agent’s load shedding actions,  $\psi$ .  $\varphi_{\min}$  and  $\varphi_{\max}$  are the minimum and maximum parameters observed in the system, specified as 0.1 and 0.4, respectively. We define these bounds based on empirical observation of our dataset, analysing the relationship between disturbances and the resulting frequency responses. We choose the most conservative as our representative bounds, giving the agent full flexibility to develop and identify optimal actions within these conservative limits, thus making  $\varphi$  an implicitly learnable parameter. As such, these values are system-dependent and can be adapted to suit any specific system without the loss of generality. (6.8a) defines the maximum feasible load shedding to prevent over-shedding (which may lead to over-frequency and/or positive RoCoF), while (6.8b) defines the minimum feasible load shedding to prevent system instability following the disturbance. However, for convenience, the lower bound is unenforced (set to zero), allowing the agent to minimise load shedding to the fullest extent possible. This is because the precise analytical derivation of locational frequency dynamics (and the exact load shedding amounts, location and time) is complex and hard, requiring finding a global minimum to a nonconvex function [14]. Therefore, we leave these hard-to-model relationships for the RL agent to discover, while only providing the swing equation, incorporated in the PS, as defined by (6.8a) and (6.8b), as its starting point. Every action of the agent,  $\psi_{\text{RL}}$ , that violates or goes beyond the permissible region of the PS is updated by uniformly scaling it as follows:

$$\psi_{\text{RL}} \leftarrow \begin{cases} \psi_{\text{RL}}, & \\ \text{if } \psi_{\min}^{\Omega_{\text{PI}}} \leq \psi_{\text{RL}} \leq \psi_{\max}^{\Omega_{\text{PI}}} & \\ \psi_{\text{RL}} \left( \min \left( 1, \frac{\psi_{\max}^{\Omega_{\text{PI}}}}{\psi_{\text{RL}}} \right) \times \max \left( 1, \frac{\psi_{\min}^{\Omega_{\text{PI}}}}{\psi_{\text{RL}}} \right) \right), & \\ \text{otherwise} & \end{cases} \quad (6.9)$$

Note that the effective imbalance in the system, after a load shedding action, depends on the activation time,  $t_{\text{Act}}$ , as follows:  $\varphi((\psi_{\text{RL}} - \Delta P_l)t_{\text{Nadir}} - \psi_{\text{RL}}t_{\text{Act}}) \rightarrow \Delta \mathcal{E}(t_{\text{Nadir}})$ . Note that while it is possible to improve the PS by adding ‘more detailed physics’ such as: modelling the influence of voltage-dependent loads, accounting for the equivalent inertia due to renewables and factoring in the variability of primary en-

ergy sources, we maintain the standard swing equation as a benchmark for the agent’s control actions to demonstrate how RL models can, even with bare bones, capture the hard-to-model relationships to establish an effective control strategy. Unlike traditional model-based controllers that rely on detailed system equations and parameter tuning, the RL agent learns from interactions with the power system simulator, observing state transitions and rewards, to uncover hidden dynamics and optimise performance while capturing detailed dynamics. Eventually, this highlights the potential of RL to generalise from limited physics, while managing the underlying physics behaviour of the power system. It also enables us to achieve faster and more efficient training, making the approach scalable to larger systems.

**ii) Physics Shield Reward,  $r_{\Omega_{\text{PI}}}$**

Rewards from the PS encourage the RL agent to consistently explore and exploit only essential actions, thereby speeding its convergence and improving the optimality of solutions. The PS’s reward,  $\Omega_{\text{PI}}$ , therefore depends on the extent of deviation from the PS’s bounds, with a weighting factor,  $\phi_{\Omega_{\text{PI}}}$  (selected through empirical evaluation of agent behaviour), as follows:

$$r_{\Omega_{\text{PI}}} = \begin{cases} \phi_{\Omega_{\text{PI}}}(\psi_{\text{max}}^{\Omega_{\text{PI}}} - \psi_{\text{RL}}), & \text{if } \psi_{\text{RL}} \geq \psi_{\text{max},\Omega_{\text{PI}}} \\ \phi_{\Omega_{\text{PI}}}(\psi_{\text{RL}} - \psi_{\text{min}}^{\Omega_{\text{PI}}}), & \text{if } \psi_{\text{RL}} \leq \psi_{\text{min},\Omega_{\text{PI}}} \\ 0, & \text{otherwise} \end{cases} \quad (6.10)$$

**iii) Agent Rewards Integration**

The overall reward of the RL agent,  $r_{\text{RL}}$ , consists of two parts: rewards from the PS,  $r_{\Omega_{\text{PI}}}$ , and rewards from the environment,  $r_f$  —which depends on the system’s frequency response. We adopt a physics-based reward function that steers the agent toward consistently performing actions that minimise deviation of the stability metric from its nominal value. This is achieved by including the observed system response,  $f_\lambda$ , within the function, rather than using binary stability indicators, as follows:

$$r_f = -\phi_{ff}(f_n - f_\lambda) \quad (6.11)$$

where  $f_\lambda$  is the minimum (or worst) frequency stability metric observed in the system.

## Chapter 6. Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

For the frequency nadir, for example, in a network of  $N$  buses,  $f_\lambda = \min(f_i), \forall i = [1, \dots, N]$ .  $f_n$  is the nominal or steady state stability condition (0 Hz/s for RoCoF and 50 Hz or 60 Hz for frequency nadir—depending on the system),  $\phi_{ff}$  is the stability metric weighting factor determining the importance of the deviation from the nominal value. The overall reward for the agent is modulated by a parameter  $\Xi$ , which is a large positive or negative number, based on the desirability of the agent's actions. We employ manual tuning to select the relative weights associated with the reward components, based on empirical agent behaviour (convergence within the specified number of interactions) and performance [56, 83]. Specifically,  $\Xi$  takes a positive value when the agent takes an action that keeps the frequency within acceptable limits, and a negative value when the action fails to prevent violation of those limits. This adjustment acts as a reward-shaping mechanism, amplifying rewards or penalties beyond the relatively modest incremental feedback provided by the environment and the PS, thereby accelerating convergence.

$$r_{\text{RL}} = \begin{cases} \Xi + r_{\Omega_{\text{PI}}} + r_f, & \text{if unstable} \\ -\Xi, & \text{otherwise} \end{cases} \quad (6.12)$$

### 6.3.4 Modelling of UFLS Relays and Implementation of Conventional UFLS Scheme

The conventional UFLS associates load-shedding with pre-set frequency thresholds for emergency frequency control. We model the scheme by discretising the load at each bus to ensure an even contribution from all loads in the system at every stage of the UFLS. This is done to avoid a potential locational bias or influence the scheme (positively or negatively) based on the arbitrary selection and allocation of loads at each UFLS stage. System protection relays are implemented using the standard Definite Time (DT) characteristic with an execution time delay of 200 milliseconds and a frequency sampling time of 0.06 seconds [17]. The amount of load shedding in the system due to the activation of UFLS relays is given as  $Z_{\text{UFLS}}$ , expressed as follows:

$$\mathcal{Z}_{\text{UFLS}} = \sum_{i=1}^{\mathcal{D}_{\text{defensive-loads}}} l_{[0/1],i}^{\text{Relay}} l_i^{\mathcal{D}_{\text{defensive-loads}}}, \quad \forall i \in \mathcal{D}_{\text{defensive-loads}} \subset \mathcal{L} \quad (6.13)$$

where  $l_{[0/1],i}^{\text{Relay}}$  indicates the status of the  $i^{\text{th}}$  relay (connected to the  $i^{\text{th}}$  load), returning 0 if closed and 1 if open,  $\mathcal{D}_{\text{defensive-loads}}$  is the set containing all loads participating in the UFLS defence scheme, while  $\mathcal{L}$  is a global set containing all loads. Since all loads participate in the UFLS, arbitrary pre-selection of certain loads to make up the total percentage required by the UFLS scheme (which sheds all loads uniformly) is not required. The Physics-Shielded Deep RL agent we propose in our work optimally determines which specific loads to shed and how much, in a coordinated manner, ensuring no locational frequency violations while minimising the total load shed, while the UFLS remains in operation as a last resort.

### 6.3.5 Battery Energy Storage Systems (BESS)

In this study, we consider power systems that include both SGs and BESS. As frequency control responsibilities are increasingly shared between traditional SGs and converter-interfaced devices, the resulting diversity in control schemes introduces modelling challenges that motivate the need for more adaptive, data-driven approaches. Battery Energy Storage Systems (BESS) are energy storage systems capable of offering system frequency support services during emergencies, thanks to power electronics for fast frequency response applications [23,84]. Such devices can provide a convenient way for rapid power output changes —quickly minimising system imbalances after disturbances and improving the system’s frequency response. Shown in Fig. 6.3 is the BESS model used in this chapter, consisting of measurement devices, a battery, a charge controller, and a converter.

The *Frequ. Controller* uses the measured frequency of the grid,  $f_{\text{grid}}$ , to determine the converter’s reference active power,  $P_{\text{ref}}$ , which is necessary for droop control,  $k_{\text{droop}}$ . The *PQ Controller* regulates the active,  $P$ , and the reactive,  $Q$ , power outputs of the converter. Specifically, the controller aligns the converter’s current active power output,

Chapter 6. Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

$P_{\text{meas}}$ , to  $P_{\text{ref}}$ .  $i_{d_{\text{ref}}}$  and  $i_{q_{\text{ref}}}$  are the current phasors transformed in  $dq$ -reference frame, a transformation that enables  $i_{d_{\text{ref}}}$  and  $i_{q_{\text{ref}}}$  to equate to the active and reactive power output of the converter, respectively. This effectively enables the controller to reduce the converter's  $P_{\text{meas}}$  output if the frequency of the grid increases and vice versa.  $P_{\text{meas}}$  is increased if the frequency of the grid deteriorates. The controller controls the converter's  $Q$  output similarly, by increasing it if the measured voltage at the coupling point,  $U_{\text{meas}}$ , is below the nominal value and reducing it if  $U_{\text{meas}}$  exceeds the nominal value. The *Charge Controller* regulates the battery's charging while the *Bat-Model* represents the battery model of the BESS unit, including its State Of Charge (SOC). The BESS droop setting,  $k_{\text{droop}}$ , is the sensitivity of the BESS to frequency fluctuations, determining the ramp rate of the BESS' power output in response to observed fluctuations as follows:

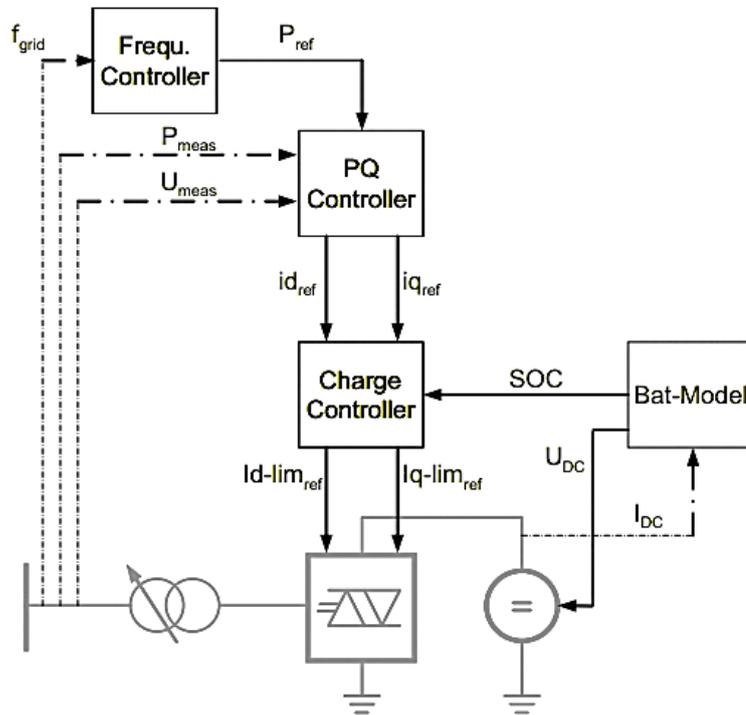


Figure 6.3: Battery Energy Storage System (BESS) Model [4]

$$\Delta P(t)^{\text{BESS}} = \begin{cases} 0, & \text{if } \Delta f_{\text{grid}}(t) < k_{\text{droop}} \cdot f_0 \\ \Delta f_{\text{grid}}(t) k_{\text{droop}} P_{\text{rated}}^{\text{BESS}}, & \text{otherwise} \end{cases} \quad (6.14)$$

Therefore, the active and reactive powers of the BESS enable the frequency and the AC voltage of the system to be controlled independently [4].

## 6.4 Case Studies and Results

Table 6.1: NN and RL Algorithms Architecture and Hyperparameter Settings

<b>NN Model Parameter</b>	<b>NN-1 (IEEE 39-Bus)</b>	<b>NN-2 (Texas 2000-Bus)</b>
Learning Rate	0.001	0.001
Batch Size, $\mathcal{B}$	64	64
Number of Layers	3	3
Number of Neurons	64	64
Activation, $\Theta$	ReLU	Tanh
<b>RL Model Parameter</b>	<b>RL-1 (IEEE 39-Bus)</b>	<b>RL-2 (Texas 2000-Bus)</b>
Learning Rate	$1e^{-4}$	$1e^{-3}$
Batch Size, $\mathcal{B}$	128	128
Experience Replay	$1e^5$	$1e^5$
Number of Layers	3	3
Number of Neurons	100:128:64	100:256:128
Activation, $\Theta$	ReLU & Sigmoid	ReLU & Sigmoid
Episodes, $\mathcal{E}$	200	250
Initial Exploration	0.5	0.5
Exploration Decay	0.9	0.9
Min. Exploration	0.01	0.01
Epsilon decay, $\epsilon_{\text{decay}}$	1000	500
$\epsilon_{\text{start}} - \epsilon_{\text{end}}$	0.9 - 0.05	0.9 - 0.05

In this chapter, the performance of the proposed RL-based AUFLS approach is evaluated in two test networks: the modified IEEE 39-bus network (a smaller network) and the Texas 2000-bus network (a larger network). This selection demonstrates how existing RL methods can converge to optimal results in low-dimensional settings but struggle to scale effectively to higher-dimensional systems. The proposed physics-informed RL approach intends to address these limitations. The modified IEEE 39-bus

network has 19 PQ load buses, while the Texas 2000-bus network has 1,515 PQ load buses. To predict RoCoF across these buses, the modified IEEE 39-bus employs a single NN with 19 regression targets, referred to as *NN-1*. In the Texas 2000-bus network, we adopt a decentralised ML approach, rather than training a single global model with 1,515 regression targets. This design choice is motivated by the fact that smaller models are easier to train (independently or in parallel) and fine-tune, due to fewer parameters. This also makes it easier to retrain the NN models for specific locations, referred to as *NN-2*, after network changes. As such, we limit the regression targets to 20 for each *NN-2* model, making a total of 75 models. The remaining 15 buses are represented by an additional NN, resulting in a total of 76 inference models. For RL agents, we train the agents, referred to as *RL-1* in the modified IEEE 39-bus network and *RL-2* in the Texas 2000-bus test network, over a maximum of 200 and 250 episodes, respectively. Episodes may end early if satisfactory performance is achieved, following an early termination strategy to facilitate training [85, 86]. Each episode consists of the agent interacting with the environment (in our case, a phasor time domain simulation), receiving rewards, and updating its policy based on the actions taken. Multiple episodes ensure that the agent iteratively refines its policy through exploration until it converges to an optimal strategy. We adopt a 70%-30% dataset split for training and testing, respectively. Both actor and critic networks have three layers with ReLU activation functions, while the output layer uses *sigmoid* activation.

The proposed RL agent operates on three and eight clusters (with scenario-specific composition) in the modified IEEE 39-bus and Texas 2000-bus networks, respectively, based on the criteria in Chapter 6.3.2. In contrast, other RL agents, without access to the system’s coherence detection information, operate on the static default/geographically defined clusters—three in the modified IEEE 39-bus network [3] and eight in the Texas 2000-bus network [27]. The parameters of all the ML models are summarised in Table 6.1, tuned using GridSearchCV—a tool that automates the process of finding the best combination of hyperparameters for a given ML model by exhaustively searching over a user-defined parameter grid [57, 58]. Operational scenarios are generated using the AC-OPF in MATPOWER [66], with MATLAB Interior Point Solver (MIPS) as a

Table 6.2: Performance Evaluation of ML Model for RoCoF Prediction

Neural Network	Metric (Hz/s)	IEEE 39-Bus	Texas 2000-Bus
NN	Minimum RMSE	0.0048	0.0023
	Mean RMSE	0.0066	0.0043
	Maximum RMSE	0.0124	0.0078

solver.

Specific procedure for data generation is described in Chapter 6.4.1 and Chapter 6.4.2. RMS-TDS simulations are conducted on the converged solutions from the AC-OPF in DIgSILENT PowerFactory, considering large events, i.e.  $N - 2$  SG outages (the two largest generating SGs in the network). This remains fixed across all case studies. The RL agent(s), therefore, interact with the RMS-TDS environment (as a power system simulator) to establish an optimal AUFLS policy. All simulations were carried out on an 11<sup>th</sup> Gen Intel (R) Core (TM) i7-11700 @ 2.50 GHz with 16 GB installed RAM.

#### 6.4.1 The Modified IEEE 39-Bus Test Case

The modified IEEE 39-bus network includes three BESS (described in Chapter 6.3.5) and one CIG integrated at buses 14, 18, 28 and 15, respectively (see Fig. 6.4). These positions remain fixed throughout the simulations. This configuration has been selected to create two distinct regions: high and low inertia, allowing a clear demonstration of the impact of high CIG integration on locational frequency dynamics. We vary the following three system variables to generate operational scenarios: the loading in the system, the number of online SGs and the CIG output. The system loading is uniformly scaled across all the system loads over a range of 16 values, ranging from 0.6 p.u. to 1.025 p.u. in increments of 0.027 p.u. Four SGs, i.e., SG 4, SG 5, SG 6 and SG 7 of Area 2, are assumed to consist of four equal-sized units,  $u$ , where  $u = [1, \dots, 4]$  (6.15), which are incrementally displaced by the CIG. The updated rating of the SG,  $SG_{MVA,new}$ , is therefore determined by the number of remaining units,  $SG_{MVA,old}$ , thereby inversely scaling the CIG (6.16). The CIG penetration equation introduces scaling factors,  $s$ , where  $s = [0, -0.05, 0, 0.05]$ , and  $r$ , where  $r = [1, 1.4]$  to decouple the

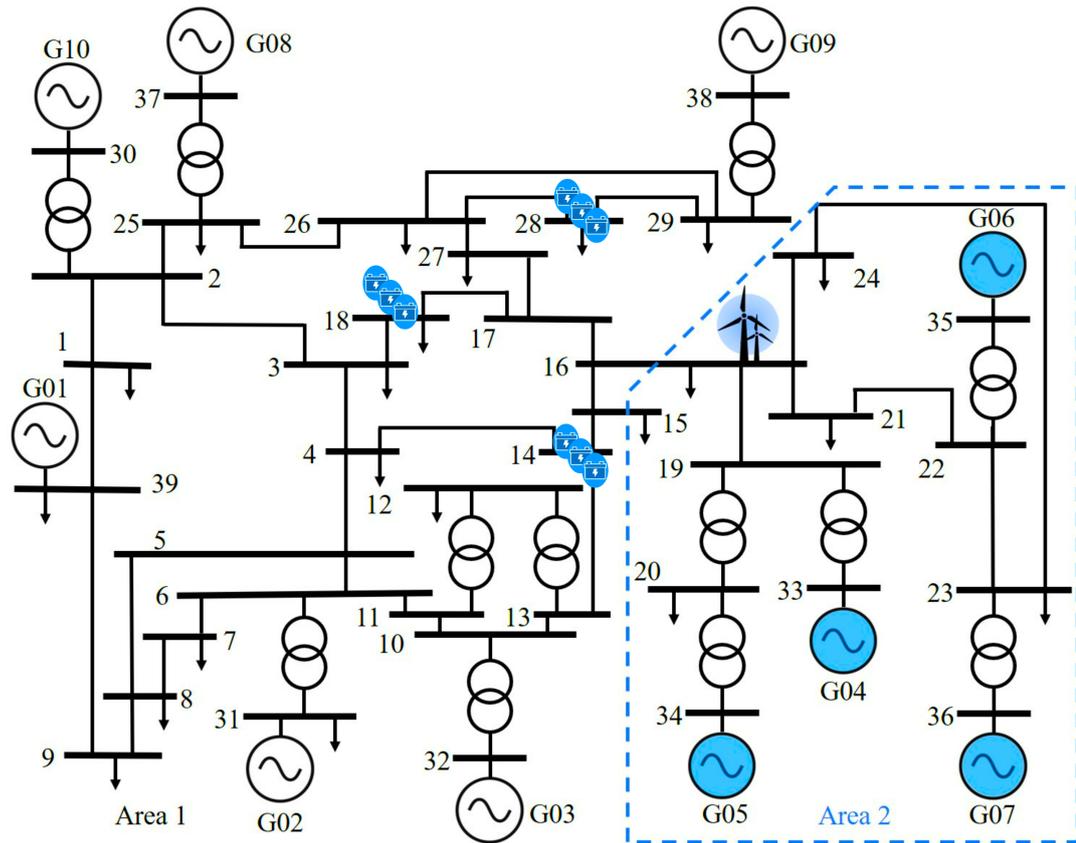


Figure 6.4: The modified IEEE 39-bus network highlighting Area 2, which is a low inertia region with the CIG location at Bus 16 and three BESS units at Buses 14, 18, and 28

Table 6.3: Standard 9-Stages Under-Frequency Load Shedding Scheme for the IEEE 39-Bus Network

UFLS Stage	Threshold (Hz)	Load (%)
<b>1</b>	58.56	5
<b>2</b>	58.50	5
<b>3</b>	58.44	10
<b>4</b>	58.32	7.5
<b>6</b>	58.08	7.5
<b>7</b>	57.84	7.5
<b>8</b>	57.60	5
<b>9</b>	57.36	5

direct linear relationship between SGs displacement and CIG penetration [3]. All of these parameters are uniformly sampled, resulting in 1,536 scenarios. Of these, 6% failed to converge in the AC-OPF process, resulting in 1,452 feasible scenarios.

$$SG_{MVA,new} = u(SG_{MVA,old}/4), \quad u = [1, \dots, 4] \quad (6.15)$$

$$CIG_{MVA} = r \left( \frac{(5-u)SG_{MVA,old}}{4} + s(SG_{MVA,old}) \right) \quad (6.16)$$

The instantaneous penetration of CIG is between 100 MW to 1000 MW, and the ratings of the BESS remain fixed at 90 MW, representing a maximum of about 44% of the system's generation. The CIG is modelled using the Type IV Wind Generator (WTG) utilising the Western Electricity Coordinating Council (WECC) control model [73]. This model is coupled to the grid through a fully rated converter. In contrast, the BESS is modelled using the standard DIgSILENT BESS model [4, 17]. The UFLS scheme has nine stages, adapted from the GB system. As is shown in Table 6.3, 60% of the total system load (6,254.2 MW and 1,387.1 MVAR) is available for the load-shedding strategy [47], and the nadir stability boundary is set at 0.95 p.u., which is 57 Hz in a 60 Hz power system.

### i) Evaluation of the RL Model Training Progress

The training progress of the RL model is evaluated using four metrics: *physics-shield-based rewards*, *cumulative reward*, *amount of load shed*, and *system frequency*

*deviation.* The PS is implemented to eliminate non-essential actions by defining plausible load-shedding regions, as described in Section 6.3.3. Fig. 6.5 summarises the training progress, and it is seen that based on the PS's dictates, at the beginning of the training, the agent's policy is still being developed, leading to random actions garnering both rewards (positive values) and penalties (negative values). Rewards are given when the total load shed remains within the shield's limits, and vice versa. As training progresses and the agent refines the policy, it learns to consistently execute actions that continuously attract positive values. Similarly, the agent's overall cumulative reward reflects this process. Initially, the agent is consistently penalised until it learns an accurate policy and maximises the cumulative reward through optimal actions. To improve efficiency, we use the early termination strategy, which terminates training once the agent demonstrates consistent execution of optimal actions, thereby reducing unnecessary training time [85, 86]. In this study, the RL agent is integrated within the existing UFLS relays to enhance the overall reliability of the control strategy by utilising the UFLS as the last resort. The load-shedding performance of the agent during training can also be seen in the figure (Fig. 6.5 bottom left). The agent executes its actions across three network clusters (determined following fulfilment of the requirements described in Section 6.3.2). During the early training stages, several activations of UFLS relays can be observed, triggered by frequency crossing particular frequency thresholds. This may be attributed to the agent shedding an insufficient load, selecting a sub-optimal combination of load clusters or wrong action execution time (or a combination of these). As the training progresses, the penalties steer the agent to optimise its actions and prevent the UFLS from intervening.

In this chapter, the RL agent is integrated within the existing UFLS relays to enhance the overall reliability of the control strategy by utilising the UFLS as the last resort. The load-shedding performance of the agent during training can also be seen in the figure (Fig. 6.5, bottom left). The agent executes its actions across three network clusters. During the early training stages, several activations of UFLS relays can be observed, triggered by frequency crossing particular frequency thresholds. This may be attributed to the agent shedding an insufficient amount of load, selecting a sub-

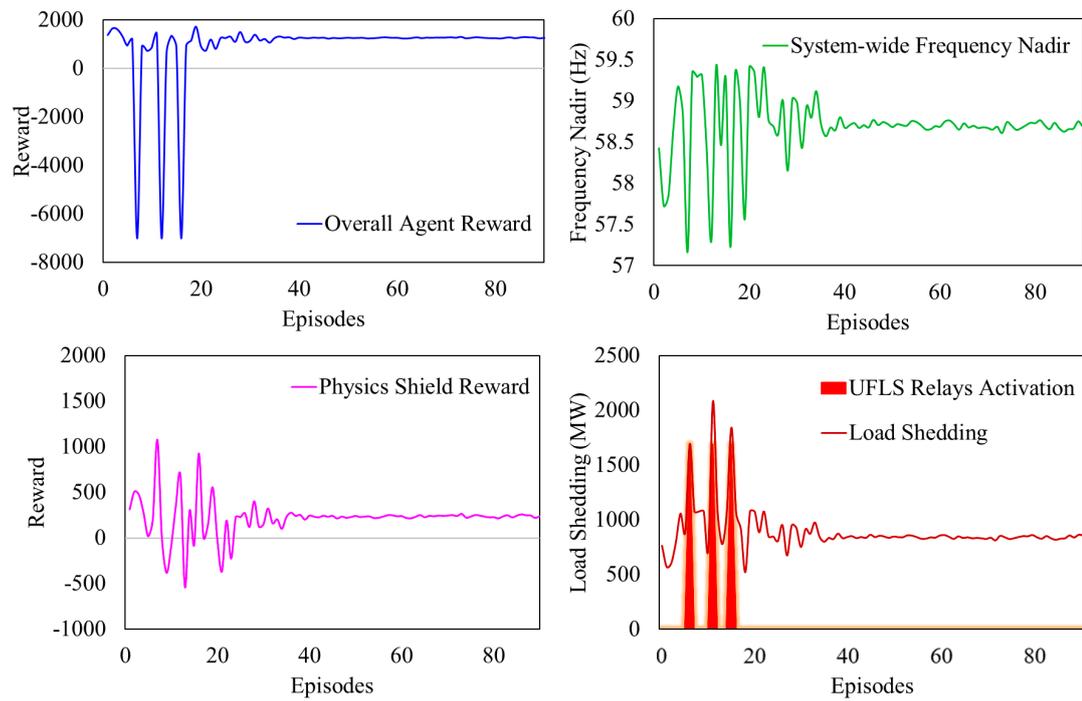


Figure 6.5: AUFLS training progress of the RL physics-informed agent in the Modified IEEE 39-Bus Network

Table 6.4: Performance Evaluation of ML Model for RoCoF Prediction in the IEEE 39-Bus Network

RMSE (Min)	RMSE (Mean)	RMSE (Max)	$R^2$ (Train / Val.)
0.0022	0.0053	0.0089	0.9976 / 0.9661
<b>Train Time (Secs.): 45.52</b>			

optimal combination of load clusters or wrong action execution time (or a combination of these). As the training progresses, the penalties steer the agent to optimise its actions and avoid traditional UFLS interventions.

**ii) Evaluating the Effectiveness of the Physics-Shielded AUFLS RL**

A summary of the quantitative performance of the NN in predicting RoCoF across all PQ load buses in the IEEE 39-bus network is given in Table 6.4. Across all the PQ load buses, the NN’s predictive performance, in terms of RMSE, achieved a minimum of 0.0022 Hz/s, a mean of 0.0053 Hz/s and a maximum of 0.0089 Hz/s. The  $R^2$  values observed during training and validation were 0.9976 and 0.9661, respectively, indicating an accurate explanation of the variance in the data. The training time required was 45.52 seconds. The predictions are therefore used for coherence detection by the K-Means clustering algorithm.

The efficacy of the fully trained proposed RL agent (*NN-Physics-Shielded RL*) is compared against other approaches: the conventional UFLS scheme (with predetermined settings according to the description in Chapter 6.4.1), a purely data-driven RL agent (without PS and *NN*), and an RL agent with the PS only (*Physics-Shielded RL*) to demonstrate the added effectiveness of *NN*’s coherence detection. Fig. 6.6 shows the total amount of load shedding in the system for each OC in the test dataset by the different approaches (three clusters seen by RL agents). The figure shows that the proposed method, i.e., the NN-Physics-Shielded RL agent, followed by the Physics-Shielded agent, shed the least amount of load compared to the rest. This is quantitatively summarised in Table 6.5, which shows that although all methods were able to maintain system stability (i.e., the frequency nadir remains above the 57 Hz threshold), some methods require more load shedding than others to achieve this. Specifically, the

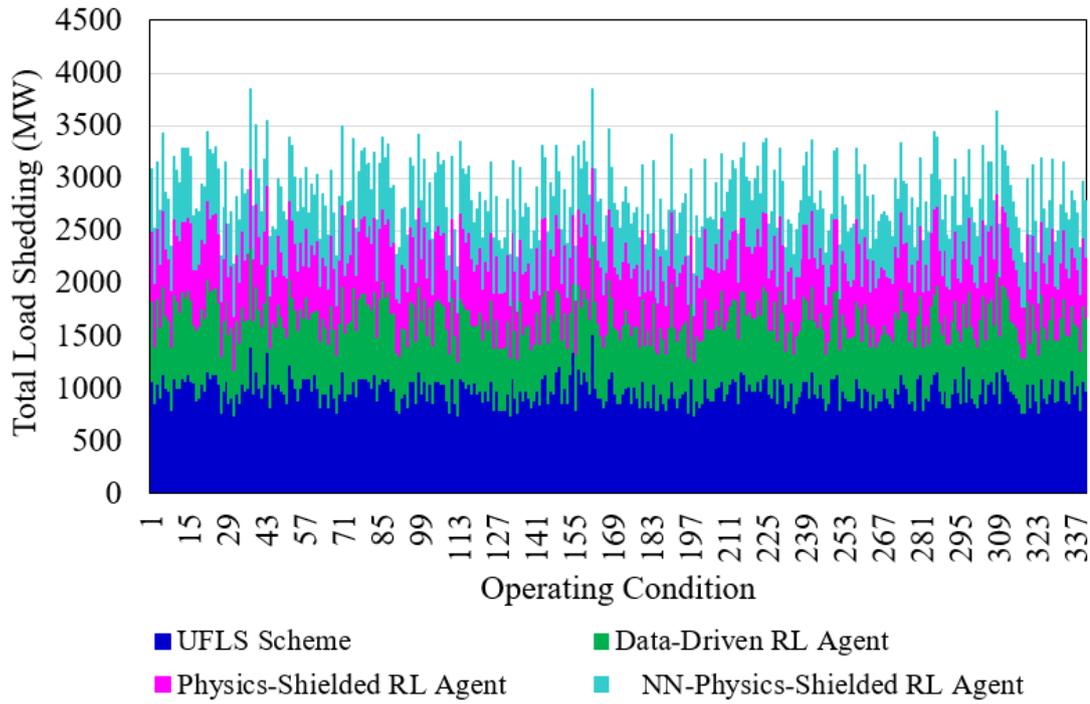


Figure 6.6: Stacked plot of the load shedding schemes’ performance by different methods for emergency frequency control in the Modified IEEE 39-bus network

conventional UFLS scheme requires an average shedding of 971 MW. The data-driven RL agent reduces this requirement by 28% to 692 MW, with a corresponding training time of 190 minutes. The Physics-Shielded RL agent further improves this performance by reducing the requirement by 35% to 627 MW while reducing the training time by 32% to 129 minutes. Additionally, including coherence detection alongside the PS in the NN-Physics Shielded RL agent enhances this performance even more, reducing the shedding by 39% to 588 MW while requiring the least training time of 125 minutes —34% further improvement of the training time required by Data-Driven RL. It should be noted that the ability of RL to control how much load and in which locations it is adaptively shed for each specific case allows for coming closer to the acceptable stability limit. This, in turn, reduces the amount of load shed and, consequently, the associated cost with such remedial actions.

Table 6.5: Mean Load Shedding (MW) for Emergency Frequency Control by Different Methods in the IEEE 39-Bus Network

Frequency Control Model	Mean Nadir	Mean Load Shed	Total Train Time
UFLS Scheme	58.28 Hz	971.40 MW	-
Data-Driven RL	<b><u>58.67 Hz</u></b>	692.47 MW	190 mins.
Physics-Shielded RL	58.43 Hz	627.63 MW	129 mins.
NN-Physics-Shielded RL	58.25 Hz	<b><u>588.63 MW</u></b>	<b><u>125 mins.</u></b>

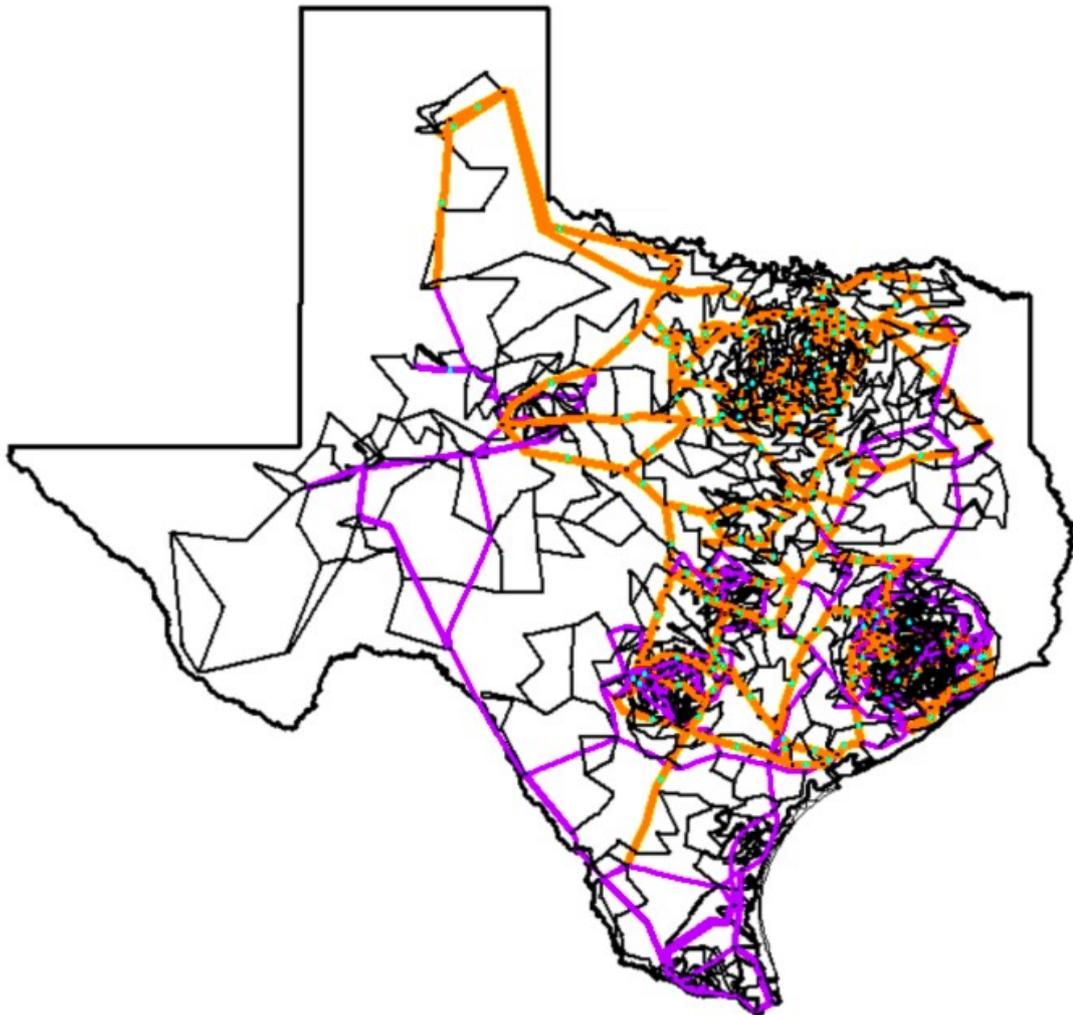


Figure 6.7: One line diagram of the synthetic Texas2000 case. Orange lines indicate 500 kV, pink lines indicate 230 kV, green lines indicate 161 kV, and black lines indicate 115 kV [5]

Table 6.6: The 6-Stages Under-Frequency Load Shedding Scheme for the Texas 2000-Bus Network [7]

UFLS Stage	Threshold (Hz)	Load (%)
<b>1</b>	59.60	7.4
<b>2</b>	59.40	5.2
<b>3</b>	59.20	5.2
<b>4</b>	59.00	5.2
<b>5</b>	58.70	6.3
<b>6</b>	58.50	3.4

### 6.4.2 The Texas 2000-Bus Synthetic Grid

The synthetic Texas 2000-bus network (Fig. 6.7), modelled on the footprint of the Electric Reliability Council of Texas (ERCOT), operates at four voltage levels (500/230/161/115 kV) and has a total generation capacity of 98 GW, utilising coal, gas, hydro, nuclear, solar, and wind technologies. No additional CIG or BESS units are added, and all parameters, including loads, generator settings, and governor settings, are default [27]. Operational scenarios are sampled from the dataset in [5], which has a total of 8,784 scenarios, generated by independently scaling the load profiles in all eight geographically interconnected grids, within the 29.25% and 100.52% range. We randomly sample 2,000 operating conditions from this dataset to illustrate the proposed approach. The UFLS scheme implemented in the network is shown in Table 6.6, according to the North American Electric Reliability Corporation (NERC), and the stability boundary is 58 Hz [7], and the system is classified as completely unstable if the frequency nadir falls below 58 Hz. Among all  $N - 2$  SG outage contingencies applied, 58% involved generators in the North Central grid, 40% in the Coast grid, and 2% in the South Central grid.

A summary of the quantitative performance of the NNs in predicting RoCoF across all PQ load buses in the Texas 2000-bus network is given in Table 6.7. Across all the PQ load buses, the predictive performance of the NNs, in terms of RMSE, achieved a minimum of 0.0045 Hz/s, a mean of 0.0063 Hz/s and a maximum of 0.0082 Hz/s. Similarly, the minimum  $R^2$  values observed during training and validation across all the NNs were 0.9960 and 0.9941, respectively, indicating a strong fit. The training

Table 6.7: Performance Evaluation of ML Model for RoCoF Prediction in the Texas 2000-Bus Network

RMSE (Min)	RMSE (Mean)	RMSE (Max)	$R^2$ (Train / Val.)
0.0045	0.0063	0.0082	0.9960 / 0.9941
<b>Train Time (Secs.): 150.79</b>			

time required was 150.79 seconds on average per model. Note that training can be parallelised to reduce the overall training time. The predictions are therefore used for coherence detection by the K-Means. Since the NNs eliminate solving DAEs, their predictions are almost instant, enabling seamless adaptive coherence detection, despite the system’s size.

The efficacy of the fully trained proposed RL agent is compared against different control models as summarised in Table 6.8. While all approaches successfully maintained system frequency stability (i.e., kept the system frequency above the 58 Hz threshold), they caused varying levels of service interruption in the system. Specifically, over the simulated week, the conventional UFLS scheme required an average of 598.22 MW, while the Data-Driven RL agent halves this requirement (50%) to 299.09 MW, with a corresponding training time of 700 minutes. This is further improved by the Physics-Shielded agent, which reduces this requirement by 79% to 127.94 MW while reducing the training time by 20% to 540 minutes. The NN-Physics-Shielded agent even further improves this and demonstrates the best performance. The requirement is reduced by 80%, bearing the least value of 90.56 MW. By combining the PS and coherence detection, using the NN’s predicted snapshot of post-disturbance frequency dynamics, the NN-Physics-Shielded agent can efficiently prune the search space —eliminating fruitless actions that other agents might waste time and ‘capacity’, exploring —enabling allocation of more time to discover effective shedding actions. The training time, reduced by 35% to 455 minutes, is also the lowest.

The results in the Texas 2000-bus network showcase good performance and also reasonable training times (about 7.5 hours per week of hourly training data —35% lower than the conventional data-driven RL agent), which highlight the ability of the

Table 6.8: Mean Load Shedding (MW) for Emergency Frequency Control by Different Methods in the Texas 2000-Bus Network

<b>Frequency Control Model</b>	<b>Mean Nadir</b>	<b>Mean Load Shed</b>	<b>Total Train Time</b>
UFLS Scheme	59.53 Hz	598.22 MW	-
Data-Driven RL	<b>59.63 Hz</b>	299.09 MW	700 mins.
Physics-Shielded RL	59.62 Hz	127.94 MW	540 mins.
NN-Physics-Shielded RL	<b>59.63 Hz</b>	<b>90.578 MW</b>	<b>455 mins.</b>

proposed method to scale to large real-world systems, addressing a significant challenge for ML-based applications. Note that for a trained RL agent, the time to provide real-time decisions is near-instantaneous. Note that, compared to the IEEE 39-bus case in Chapter 6.4.1—which is much smaller, the NN-Physics-Shielded RL sheds less load in this case study due to the relative severity of the disturbance. The IEEE 39-bus network has only 10 SGs, and the most severe  $N - 2$  disturbances accounted for up to 35% and 25% of the generation and reserves, respectively; while the Texas 2000-bus network has over 400 online SGs with significant headroom, experiencing a maximum generation and reserve impact of 8% and 14%, respectively. This explains why the average frequency nadir in the Texas 2000-bus network is higher than that in the IEEE 39-bus network—suggesting that a sequence of load shedding events due to relay activations was triggered more frequently in the latter case.

### 6.4.3 Locational Aspect Considerations and the Effectiveness of Shedding Strategies

A key aspect of our proposed method is its ability to identify coherent areas within the network systematically and adaptively. This facilitates scalability in high-dimensional settings and enables more efficient, as well as precise, utilisation of the load clusters, compared to, for example, clustering based on geographic areas. In this study, the proposed RL agent operates on three and eight clusters in the modified IEEE 39-bus and Texas 2000-bus networks, respectively, based on the criteria in Chapter 6.3.2. In contrast, other agents use geographically defined clusters—three in the 39-bus and eight in the Texas 2000-bus network. We analyse the locational aspects of the schemes in

Chapter 6. Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

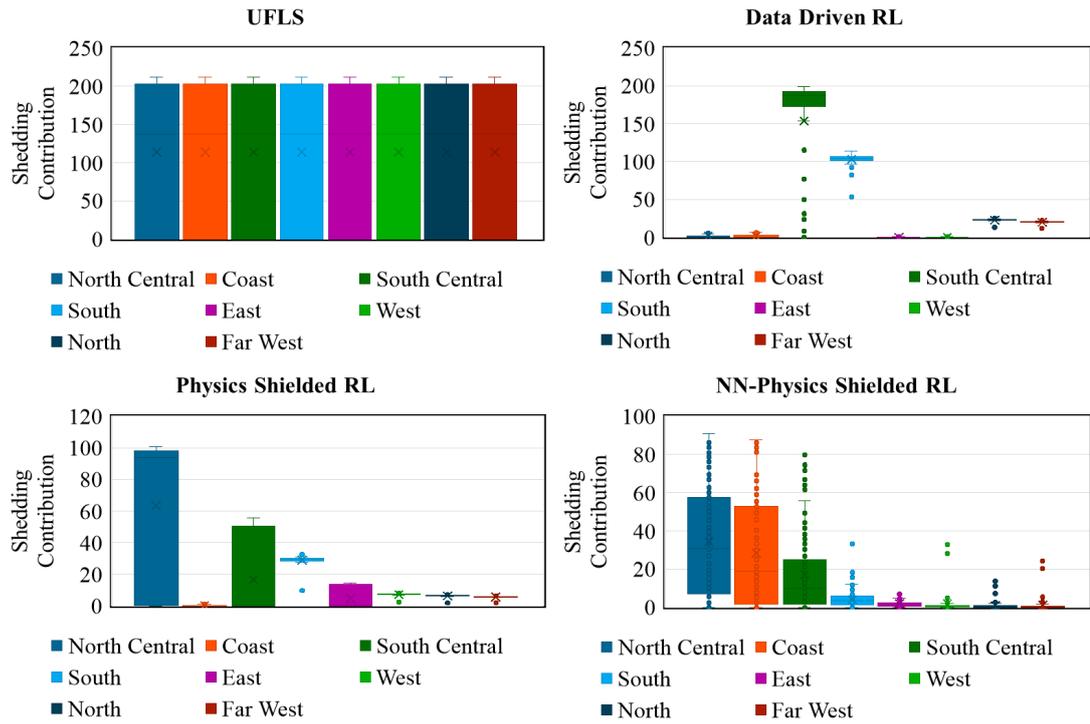


Figure 6.8: Load shedding behaviour of different frequency control models focusing on shedding contribution (in MW) from participating loads in the Texas 2000-bus network

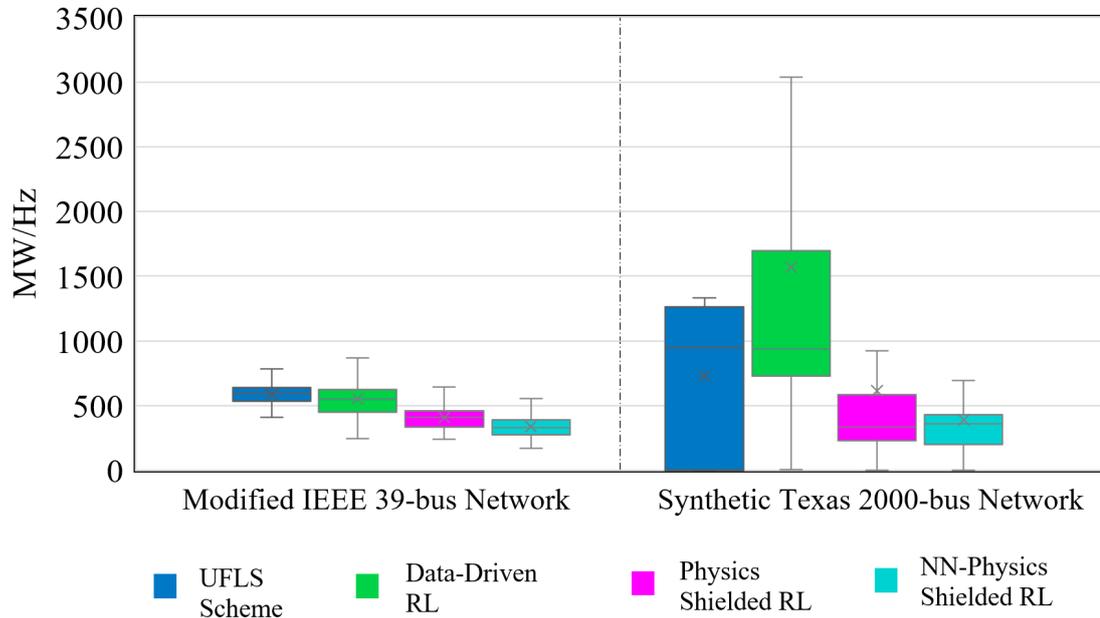


Figure 6.9: Efficiency of the emergency load shedding schemes in the Modified IEEE 39-Bus and the Texas 2000-Bus Networks

Fig. 6.8, based on the Texas 2000-bus network. The UFLS scheme uniformly allocates participating loads across the network to the shedding scheme. In contrast, RL-based methods adaptively adjust the contributions from each cluster based on their policies, minimising where it is perceived as less effective and maximising in those deemed effective. Specifically, the Data-Driven RL and Physics Shielded RL agents, lacking direct access to the real-time dynamic state of the system, develop a limited but noticeable concept of spatial shedding effectiveness, compared to the NN-Physics Shielded RL agent, which has direct access to the detailed system’s dynamics. The Data-Driven RL agent maximises shedding from the South Central and the South grids, the Physics Shielded RL agent maximises from the North Central, South Central and South grids, and the NN-Physics Shielded RL agent maximises the North Central, the Coast and the South Central grids.

The effectiveness of the load shedding policies is evaluated by measuring the load shedding required per Hz, expressed in MW/Hz. This is calculated by dividing the amount of load shed by the frequency deviation for each stable scenario, then averaging across all test cases. This metric indicates how much extra load the scheme sheds, beyond what is minimally necessary to maintain stability. Ideally, a lower value is preferable, as it indicates that the scheme can keep the frequency within acceptable limits, without resorting to aggressive load shedding—even when the frequency nadir approaches, but remains above, the stability boundary. As shown in Fig. 6.9, in the modified IEEE 39-bus network, the UFLS scheme has the highest MW/Hz requirement of 550 MW/Hz. The Data-Driven RL agent slightly improves this by 4.5% to 525 MW/Hz, but has wider whiskers and inter-quartile range, suggesting more aggressive loading shedding for a limited number of scenarios. The Physics-Shielded RL agent significantly improves this by 25% by reducing the value to 400 MW/Hz. Furthermore, the NN-Physics Shielded RL agent achieves the best performance, with an improvement of 40%, resulting in the lowest value of 330 MW/Hz.

In the synthetic Texas 2000-bus network, a high-dimensional system reflecting a common challenge in modern grids, the Data-Driven RL agent struggles and performs worse than the UFLS. This results from a vast search space, which causes the agent to

waste a considerable amount of time exploring non-essential actions. As a result, it has the highest requirement of 1,581 MW/Hz, suggesting more aggressive load shedding than necessary to maintain stability. As shown in Table 6.9, the highest recorded nadir by the agent is 59.97 Hz, exceeding 59.62 Hz by the UFLS. This, coupled with highly inefficient load shedding contributed by the clusters, leads to high MW/( $\Delta$  Hz) values—note how other agents attained similar frequency nadirs but with less load shedding, hence lower MW/( $\Delta$  Hz value). This effectively demonstrates how high dimensionality impacts the effectiveness of purely data-driven models.

Consequently, while the agent appears to perform well on average based on Table 6.8, Fig. 6.9 shows that it can be highly inefficient under certain scenarios. Specifically, by combining the agent’s load shedding behaviour with the corresponding frequency response in the system, the figure reveals shedding inefficiencies reaching as high as 3,000 MW/Hz—due to significant load shedding, resulting in a small frequency change  $\Delta$ Hz, increasing the MW/Hz value. Generally, for purely data-driven RL models in high-dimensional settings, two outcomes are possible: i) *failure to converge*, or ii) *convergence to a sub-optimal solution* (as observed here). Attempts to improve the model would significantly add to the computational requirements. This performance is then followed by the UFLS scheme with 728 MW/Hz, which improves it by 54%. The Physics Shielded agent improves efficiency by 61%, reducing this requirement to 616 MW/Hz. An even further improvement of 75% is achieved by the NN-Physics Shielded agent, leading to the lowest requirement of 393 MW/Hz as the most efficient. This is due to the model’s ability to distribute the load more effectively across coherent locations (based on the current system’s dynamic state). TDS also show that the NN-Physics Shielded RL agent achieved this by targeting locations with the highest frequent disturbances (as indicated in Section 5.5.3, the agent does not have direct access to this information).

#### 6.4.4 Computational Considerations

The computational requirements for training RL models are crucial for their adoption and application in modern large-scale power systems, which are challenged by high

Table 6.9: Comparison of load shedding (MW) and frequency nadir (Hz) between four load shedding schemes

Network	Metric	UFLS	Data-RL	Phys-RL	NN-Phys-RL
IEEE 39-Bus	Min. Load (MW)	731.65	440.03	443.48	<b><u>386.59</u></b>
	Max. Load (MW)	1,509.03	948.80	820.01	<b><u>798.57</u></b>
	Min. Nadir (Hz)	57.43	<b><u>57.89</u></b>	57.65	57.40
	Max. Nadir (Hz)	58.38	<b><u>59.36</u></b>	58.99	59.00
Texas 2000-Bus	Min. Load (MW)	<b><u>0.00</u></b>	126.67	47.07	84.14
	Max. Load (MW)	1,160.99	362.39	204.21	<b><u>109.87</u></b>
	Min. Nadir (Hz)	58.80	<b><u>59.40</u></b>	59.38	59.36
	Max. Nadir (Hz)	59.62	<b><u>59.97</u></b>	<b><u>59.97</u></b>	59.96

dimensionality. The time required for training RL models depends on numerous factors, including the problem’s dimensionality, objectives, model architecture (such as the parameters shown in Table 6.1), problem complexity, and the training setup, among others. Therefore, careful consideration is needed when setting up RL models. As shown by Table 6.5 and Table 6.8, the proposed approach significantly reduces the training time by at least 30%. This value could be even higher if the conventional data-driven RL agent was strictly required to converge to solutions better than the UFLS (i.e. in some cases, the data-driven RL agent still requires the activation of UFLS to limit frequency deviations). The proposed method demonstrates good scalability and is particularly beneficial in high-dimensionality scenarios, where training times can be a huge challenge.

## 6.5 Conclusion

Reinforcement Learning (RL) has demonstrated significant potential in various power systems domain applications, including adaptive Under-Frequency Load Shedding (AU-FLS). However, key challenges encountered by RL-based approaches include the vast search space inherent in power systems, which leads to sub-optimal performance and sampling inefficiencies due to excessive sampling and consequent difficulty in scaling up. This eventually hinders the adoption of RL by power systems practitioners. This

## Chapter 6. Emergency Adaptive Frequency Control Mechanisms with Reinforcement Learning

work addresses these challenges by developing a Physics-Informed Deep Reinforcement Learning (PI-DRL) model for AUFLS, considering the swing equation and adaptive system coherence detection. Consequently, this approach addresses the curse of dimensionality, improves optimality and enables scalability to large networks.

The effectiveness of the proposed method is demonstrated through two case studies: the modified IEEE 39-bus model and the synthetic Texas 2000-bus model. By comparing the proposed approach with the conventional, purely data-driven RL approach and the conventional UFLS, improvements were notable in terms of the optimality of solutions. For instance, close to 80% mean value reduction was observed in the amount of load shed in the 2000-bus model while reducing the training time by at least 35%. Moreover, including the governing physics within the model helps in building trust in RL models among practitioners. The empirical performance improvements and scalability to scale in large, realistic networks (up to 2000 buses) highlight PI-DRL as a promising solution for advancing AUFLS in modern power systems.

## Chapter 7

# Conclusions and Future Work

The transition to zero-emission power systems with high renewable penetration presents significant technical challenges related to capturing, understanding and managing the detailed system dynamics while ensuring secure system operation. This thesis developed a novel set of Machine Learning (ML) models to address these challenges, focusing on situational awareness, as well as preventive and corrective stability management where conventional approaches fall short.

Power systems dynamics are becoming more complex, leading to the frequency dynamics becoming more locational/regional. As a result, conventional approaches, such as Centre of Inertia (COI)-based methods, may fail to capture these locational variations, thereby leading to risks of unforeseen locational frequency violations. To address this challenge and enhance the system's situational awareness, in Chapter 3, we proposed an ML-based technique to capture the system's detailed locational frequency dynamics. The key insight obtained is that ML-based situational awareness offers a dual benefit of high accuracy comparable to that of time domain simulations, and high predictive speed, at a fraction of the time required by conventional analytical approaches —without imposing additional computational overhead. This is because to make such predictions, ML-based models do not require solving the network's Differential Algebraic Equations (DAEs), allowing ample time for system interventions to ensure secure operation.

System operators often encounter challenges of diminished system understanding

due to the increasing complexity associated with the high share of renewables. To address this challenge, in Chapter 4, we introduced an ML-based model for explaining locational frequency dynamics in power systems. An additive feature attribution method, using SHapley Additive exPlanations (SHAP), was proposed to explain the system and enhance its understanding by identifying statistically critical features influencing the stability metric(s) of interest. Furthermore, rather than simply explaining the insights derived, we proceeded to verify these insights within the power system context. The key insight obtained from this work is that by combining domain knowledge with the explanations or insights derived from the ML model, operators are equipped with advanced decision-making tools necessary for effective targeted preventive actions capable of enhancing the stability of the system.

To ensure reliable frequency stability management in security-constrained optimisation, accuracy and computational efficiency are prerequisites. To tackle this dual challenge, in Chapter 5, we proposed a formulation of a neural network-constrained Optimal Power Flow (OPF) problem that considers locational frequency dynamics where conventional analytical methods face significant challenges. Specifically, if detailed analytical formulations are applied, a significant computational hindrance is created due to solving DAEs, which slows the convergence time taken by the solver. Furthermore, if simplified models are applied, the security of the system is affected due to the risks of introducing errors, and if the system is over-secured, the cost of generation increases. The key insight obtained from the proposed approach is that a well-trained ML model can be formulated as an accurate and efficient constraint within power system optimisation models, leading to superior performance. This results from the ML's ability to directly account for these detailed locational dynamics, including the 'hard-to-model' relationships that existing analytical methods struggle to accurately model, without the need to solve the network's DAEs during the optimisation. Such an approach is well-suited for both real-time (or close to real-time) and offline applications in scenario screening, where applying detailed conventional methods would make it computationally prohibitive.

Considering that all power systems are susceptible to disturbances during their

operation, real-time control actions are an integral part of system operation. Efficient control actions are scenario-specific, optimising system security without incurring any excessive costs. This is unlike the conservative static approach adopted by conventional UFLS, which utilises pre-defined thresholds to shed load to maintain stability, leading to inefficient actions. This challenge is more pressing now than ever, as the accuracy of these thresholds could be affected by the introduction of stochastic renewables into the system, as well as evolving load profiles. To address this challenge, in Chapter 6, we proposed an adaptive real-time frequency control for emergencies, through physics-informed RL. The proposed approach incorporates the physics governing the dynamic system to attain superior performance over existing AUFLS RL methods, through a Physics Shield (PS). The key insights derived are that by integrating the legacy protection schemes with the proposed AUFLS RL model as supplementary controllers, the optimality of the overall protection against frequency violation is improved, by minimising the scale of system interruptions while maintaining stability. Additionally, embedding some of the governing physics governing the dynamic system offers the advantage of faster training of the RL agent(s) —a crucial and necessary characteristic to foster scalability to large networks.

### **7.1 Future Work**

The research presented in this thesis has focused on applying ML methods to enhance the stability of power systems, focusing on frequency stability. While this thesis addresses some of the key challenges faced by conventional methods in situational awareness, and preventive as well as corrective actions to maintain frequency stability, several potential research directions have been identified for further research as follows:

#### **7.1.1 Power System Stability Situational Awareness and Stability Management**

ML methods are data-driven techniques whose accuracy depends on the quality of available training data. As demonstrated by the results in Chapter 3, these models

are capable of developing highly accurate functions that map inputs to outputs of a dynamic system. However, potential directions for future research lie in the guarantees of their predictions. This stems from the fact that these models are usually validated within the learned distributions; as such, their predictive accuracy can be unpredictable when applied to scenarios beyond their training domain. Thus, this highlights a key area for future research, such as verifying and/or improving the robustness of such ML models beyond the training domain. Despite the initial steps made in this thesis, further improvements are needed, such as investigating the learning of the physics governing the dynamic system (see [Pub. H] as an example), to go beyond learning the relationships captured in the datasets. This could potentially ensure reliable ML prediction performance, a crucial aspect among practitioners to foster adoption and application in the safety-critical power system infrastructure. Furthermore, while Chapter 3 focuses on developing fast and efficient models for enhancing situational awareness, a valuable direction for future research lies in leveraging real-time co-simulation environments. Tools such as OPAL-RT offer valuable Hardware In the Loop (HIL) capabilities, enabling real-time validation and scalability to large-scale system studies, such as in the real GB network. Furthermore, integration with platforms such as ePHASORSIM or HYPERSIM could facilitate deployment readiness assessments.

In addition, the research conducted in this thesis employed RMS-TDS to capture the frequency dynamics following disturbances. However, as modern grids become increasingly dominated by CIGs, emergent frequency behaviours —particularly those driven by fast electromagnetic interactions —necessitate modelling within the EMT Electro-Magnetic Transients (EMT) domain. This shift makes such modelling more crucial than before to accurately account for these dynamics. Consequently, exploring the effectiveness and adaptability of the proposed ML-based techniques within EMT environments presents a valuable path for future research, especially in improving situational awareness in power systems.

Regarding the application of ML models to enhance the understanding of power system dynamics, future research could explore the simultaneous implementation of multiple SHAP-derived constraints to improve the system’s stability, rather than indi-

vidual implementation. This could go beyond the initial steps made in this thesis by providing a more structured and coordinated approach to constraint implementation by effectively balancing the influence of multiple features to enhance the overall effectiveness of such insights, without causing antagonistic effects. Additionally, causal models could also be explored to uncover underlying causal relationships within the system, beyond the correlative relationships generated by additive feature attribution methods. Such an approach could enable the identification of causal features and support the design of more precise interventions to enhance system stability.

### **7.1.2 Preventive Frequency Stability Management in Security-Constrained Optimisation**

ML models with nonlinear activation functions offer greater predictive accuracy as they are capable of capturing the system’s nonlinearities. However, introducing nonlinear constraints within the optimisation problem makes the problem much harder to solve and significantly increases the optimisation time. Consequently, as demonstrated in Chapter 5, security-constrained optimisation models often rely on the linearisation of these inference models as constraints to efficiently capture and account for the detailed system stability dynamics.

However, effectively linearising nonlinear inference models can be quite challenging, as it requires balancing accuracy and computational efficiency, without focusing on one at the expense of the other. It is therefore crucial to preserve the level of accuracy needed, while maintaining the necessary computational simplicity, to allow reliable, practical application of the models in close to real-time timescales. Therefore, building on the foundations laid in this thesis, a promising direction for future research is in identifying an optimal integration point that effectively embeds complex and accurate inference models within the optimisation framework, without imposing significant computational overhead. This may involve employing more advanced optimisation tools such as CROWN, an efficient verification algorithm that facilitates bound propagation by relaxing nonlinear activation functions using adaptive linear bounds.

Furthermore, although neural networks exhibit remarkable expressive power, they

remain powerful black-box models. In contrast, Kolmogorov-Arnold Networks (KANs) have recently emerged as suitable alternatives to traditional neural network architectures, due to their ability to decompose multivariate functions into a finite composition of simple functions. KANs have demonstrated superior accuracy in certain contexts, while offering full interpretability through their functional matrices. Consequently, an interesting avenue for future research involves extracting symbolic expressions from KANs and embedding them as stability constraints within optimisation frameworks, thereby enhancing both transparency and precision in security-constrained optimisation.

### 7.1.3 Adaptive Frequency Control Mechanisms for Emergencies

The physics-informed AUFLS technique proposed in this thesis, in Chapter 6, presents a significant improvement over existing AUFLS RL methods, particularly in terms of computational efficiency and optimality of the generated solutions. Specifically, while the current implementation already minimises excessive sampling by RL agents in the AUFLS domain by a significant margin, further improvements are possible. To this end, going beyond the proposed physics and exploring increasingly rich physical insights can further constrain the RL to only explore within a search space of essential actions, thereby improving the scalability and optimality of AUFLS RL solutions even more. Thus, there is effectively no limit to how much physics the agent can account for during its training, as long as such modifications do not introduce implementation inefficiencies/complexities. Additionally, introducing physics-informed loss functions offers additional advantages by embedding domain knowledge directly into the learning objective(s). Ultimately, these promising directions improve the sampling efficiency and mitigate the significant computational requirements associated with RL training, which currently remains a bottleneck for application in large-scale power systems with large datasets.

Furthermore, employing physics-informed surrogate models, rather than detailed phasor domain power system simulators as the RL agent's environment, can significantly accelerate the training process. By eliminating the need to solve complex dif-

## Chapter 7. Conclusions and Future Work

ferential equations through the simulators within the training loop, surrogate models, as an efficient dynamics predictor, enable rapid simulation of the system responses, thereby facilitating applications in modern large-scale power systems challenged with high-dimensionality. Finally, inspired by the success of Google’s AlphaZero in mastering complex decision spaces such as chess, the RL agent’s exploration strategy could be further refined. Rather than relying on random sampling, a Monte Carlo-based exploration mechanism may offer more structured and efficient sampling to improve convergence and robustness of the learned policy.

# Bibliography

- [1] Energy Institute Statistical Review of World Energy, “Statistical review of world energy,” 2023. [Online]. Available: [https://www.energyinst.org/\\_data/assets/pdf\\_file/0004/1055542/ELStat\\_Review\\_PDF\\_single\\_3.pdf](https://www.energyinst.org/_data/assets/pdf_file/0004/1055542/ELStat_Review_PDF_single_3.pdf)
- [2] J. Rocabert, A. Luna, F. Blaabjerg, and P. Rodríguez, “Control of power converters in ac microgrids,” *IEEE Transactions on Power Electronics*, vol. 27, no. 11, pp. 4734–4749, 2012.
- [3] R. I. Hamilton, P. N. Papadopoulos, W. Bukhsh, and K. Bell, “Identification of important locational, physical and economic dimensions in power system transient stability margin estimation,” *IEEE Transactions on Sustainable Energy*, vol. 13, no. 2, pp. 1135–1146, 2022.
- [4] DIgSILENTGmbH. (2010) Battery energy storing systems (bess). [Online]. Available: <https://www.digsilent.de/en/downloads.html>
- [5] H. Li, J. H. Yeo, A. L. Bornsheuer, and T. J. Overbye, “The creation and validation of load time series for synthetic electric power systems,” *IEEE Transactions on Power Systems*, vol. 36, no. 2, pp. 961–969, 2021.
- [6] R. Preece, “A probabilistic approach to improving the stability of meshed power networks with embedded hvdc lines,” Ph.D. dissertation, University of Manchester, 2013. [Online]. Available: <https://research.manchester.ac.uk/en/studentTheses/a-probabilistic-approach-to-improving-the-stability-of-meshed-pow>
- [7] North American Electric Reliability Corporation. (2021) Reliability guideline recommended approaches for ufls program: Design with increasing

## Bibliography

- penetrations of ders. [Online]. Available: [https://www.nerc.com/comm/RSTC\\_Reliability\\_Guidelines/Recommended\\_Approaches\\_for\\_UFLS\\_Program\\_Design\\_with\\_Increasing\\_Penetrations\\_of\\_DERs.pdf](https://www.nerc.com/comm/RSTC_Reliability_Guidelines/Recommended_Approaches_for_UFLS_Program_Design_with_Increasing_Penetrations_of_DERs.pdf)
- [8] UK Parliament. (2024) The uk’s plans and progress to reach net zero by 2050. [Online]. Available: <https://commonslibrary.parliament.uk/research-briefings/cbp-9888/#:~:text=Download%20full%20report-,The%20UK%20is%20committed%20to%20reaching%20net%20zero%20by%202050,warming%20and%20resultant%20climate%20change>
- [9] ESO. (2024) Britain’s electricity explained: July 2024. [Online]. Available: <https://www.nationalgrideso.com/document/323766/download>
- [10] UK Government, “Clean power 2030 action plan: A new era of clean electricity,” 2024. [Online]. Available: <https://assets.publishing.service.gov.uk/media/677bc80399c93b7286a396d6/clean-power-2030-action-plan-main-report.pdf>
- [11] Electric Reliability Council of Texas, Inc, “Inertia: Basic concepts and impacts on the ERCOT grid,” 2021. [Online]. Available: [http://www.ercot.com/content/wcm/key\\_documents\\_lists/141324/Inertia\\_Basic\\_Concepts\\_Impacts\\_On\\_ERCOT\\_v0.pdf](http://www.ercot.com/content/wcm/key_documents_lists/141324/Inertia_Basic_Concepts_Impacts_On_ERCOT_v0.pdf)
- [12] A. Birchfield, “Inertia adequacy in transient stability models,” *11th Bulk Power Systems Dynamics And Control Symposium (IREP)*, July 25-30, 2022.
- [13] E. F. A. K. Tuohy, P. Dattaray and E. Lannoye, “Implications of reduced inertia levels on the electricity system: Technical report on the challenges and solutions for system operations with very high penetrations of non-synchronous resource,” 2019.
- [14] L. Badesa, F. Teng, and G. Strbac, “Conditions for regional frequency stability in power system scheduling—part i: Theory,” *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5558–5566, 2021.

## Bibliography

- [15] —, “Conditions for regional frequency stability in power system scheduling—part ii: Application to unit commitment,” *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5567–5577, 2021.
- [16] J. Liu, C. Wang, J. Zhao, and T. Bi, “Rocof constrained unit commitment considering spatial difference in frequency dynamics,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1111–1125, 2024.
- [17] DIgSILENTGmbH. (2023) Powerfactory v2023. [Online]. Available: <https://www.digsilent.de/en/downloads.html>
- [18] Y. Xie, J. Zheng, F. Mei, G. Taylor, and A. Gao, “An efficient approach for regional photovoltaic power forecasting optimization based on texture features from satellite images and transfer learning,” *Applied Energy*, vol. 385, p. 125505, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306261925002351>
- [19] M. Khodayar, G. Liu, J. Wang, and M. E. Khodayar, “Deep learning in power systems research: A review,” *CSEE Journal of Power and Energy Systems*, vol. 7, no. 2, pp. 209–220, 2021.
- [20] S. Püschel-LØvengreen and P. Mancarella, “Frequency response constrained economic dispatch with consideration of generation contingency size,” in *2018 Power Systems Computation Conference (PSCC)*, 2018, pp. 1–7.
- [21] D. Ortiz Villalba, J. Llanos, A. Muñoz Jadán, R. Moreno, C. Rahmann, and B. Pal, “Optimizing system operation with nadir considerations via simulations of detailed system dynamic responses,” *Electric Power Systems Research*, vol. 212, 07 2022.
- [22] H. Ahmadi and H. Ghasemi, “Security-constrained unit commitment with linearized system frequency limit constraints,” *IEEE Transactions on Power Systems*, vol. 29, no. 4, pp. 1536–1545, 2014.

## Bibliography

- [23] G. Misyris, D. Ramasubramanian, P. Mitra, and V. Singhvi, “Locational aspect of fast frequency reserves in low-inertia systems – control performance analysis,” 07 2022.
- [24] Q. Wang, F. Li, Y. Tang, and Y. Xu, “Integrating model-driven and data-driven methods for power system frequency stability assessment and control,” *IEEE Transactions on Power Systems*, vol. 34, no. 6, pp. 4557–4568, 2019.
- [25] M. Aghahassani, E. D. Castronuovo, P. Ledesma, and F. Milano, “Extended frequency divider formula with inclusion of der control dynamics,” in *2023 IEEE Power Energy Society General Meeting (PESGM)*, 2023, pp. 1–5.
- [26] F. Milano and Ortega, “Frequency divider,” *IEEE Transactions on Power Systems*, vol. 32, no. 2, pp. 1493–1501, 2017.
- [27] Texas AM University. (2024) Activsg2000: 2000-bus synthetic grid on footprint of texas. [Online]. Available: <https://electricgrids.engr.tamu.edu/electric-grid-test-cases/activsg2000/>
- [28] J. Zhao, L. Mili, and F. Milano, “Robust frequency divider for power system online monitoring and control,” *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 4414–4423, 2018.
- [29] J. Zhao, Y. Tang, and V. Terzija, “Robust online estimation of power system center of inertia frequency,” *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 821–825, 2019.
- [30] H. R. Chamorro, A. D. Orjuela-Cañón, D. Ganger, M. Persson, F. Gonzalez-Longatt, V. K. Sood, and W. Martinez, “Nadir frequency estimation in low-inertia power systems,” in *2020 IEEE 29th International Symposium on Industrial Electronics (ISIE)*, 2020, pp. 918–922.
- [31] E. Nukic and T. Konjić, “Center of inertia frequency estimation using deep learning algorithm,” *BH Electrical Engineering*, vol. 15, 12 2021.

## Bibliography

- [32] X. Wang, Y. Ba, R. Liu, H. Wang, R. Ju, W. Shi, N. Zou, and W. Li, “Model-data integration driven based power system frequency response model,” in *2021 IEEE International Conference on Artificial Intelligence and Computer Applications (ICAICA)*, 2021, pp. 107–112.
- [33] A. Venzke and S. Chatzivasileiadis, “Verification of neural network behaviour: Formal guarantees for power system applications,” *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 383–397, 2021.
- [34] R. I. Hamilton, J. Stiasny, T. Ahmad, S. C. Chevalier, R. Nellikkath, I. Murza-khanov, S. Chatzivasileiadis, and P. N. Papadopoulos, “Interpretable machine learning for power systems: Establishing confidence in shapley additive explanations,” *ArXiv*, vol. abs/2209.05793, 2022.
- [35] R. I. Hamilton and P. N. Papadopoulos, “Using shap values and machine learning to understand trends in the transient stability limit,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1384–1397, 2024.
- [36] J. Kruse, B. Schäfer, and D. Witthaut, “Revealing drivers and risks for power grid frequency stability with explainable ai,” *Patterns*, vol. 2, no. 11, p. 100365, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666389921002270>
- [37] N. Nguyen-Hong and Y. NAKANISHI, “Frequency-constrained unit commitment considering battery storage system and forecast error,” in *2018 IEEE Innovative Smart Grid Technologies - Asia (ISGT Asia)*, 2018, pp. 1171–1176.
- [38] L. Zhang, Q. Zhang, H. Fan, H. Wu, and C. Xu, “Big-m based milp method for scuc considering allowable wind power output interval and its adjustable conservativeness,” *Global Energy Interconnection*, vol. 4, no. 2, pp. 193–203, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2096511721000396>
- [39] K. Li, D. Liu, P. Gu, W. Qiu, J. Fang, and X. Ai, “Linearized frequency deviation based frequency-constrained unit commitment with support from wind farm,” in

## Bibliography

- 2022 IEEE 5th International Electrical and Energy Conference (CIEEC)*, 2022, pp. 1679–1685.
- [40] X. Zhao, H. Wei, J. Qi, P. Li, and X. Bai, “Frequency stability constrained optimal power flow incorporating differential algebraic equations of governor dynamics,” *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 1666–1676, 2021.
- [41] M. Rezaeian Koochi, S. Esmaili, and G. Ledwich, “Taxonomy of coherency detection and coherency-based methods for generators grouping and power system partitioning,” *IET Generation, Transmission Distribution*, vol. 13, 06 2019.
- [42] Q. Wang, C. Zhang, Y. Lü, Z. Yu, and T. Yi, “Data inheritance-based updating method and its application in transient frequency prediction for a power system,” *International Transactions on Electrical Energy Systems*, vol. 29, p. e12022, 04 2019.
- [43] T. Xia, N. Zhang, W. Li, E. Du, Y. Su, C. Fang, and C. Kang, “Efficient embedding of neural network-based stability constraints into power system dispatch,” *IEEE Transactions on Power Systems*, vol. 39, no. 3, pp. 5443–5446, 2024.
- [44] M. Tuo and X. Li, “Selectively linearized neural network based rocof-constrained unit commitment in low-inertia power systems,” in *2023 North American Power Symposium (NAPS)*, 2023, pp. 1–5.
- [45] M. A. Toro-Mendoza, J. Segundo-Ramírez, A. Esparza-Gurrola, N. Visairo-Cruz, C. A. N. Guitiérrez, and C. Pérez-negrón, “Toward adaptive load shedding remedial action schemes in modern electrical power systems,” *IEEE Access*, vol. 11, pp. 111 011–111 033, 2023.
- [46] S. Bustamante, J. W. gonzález, G. J. López, and H. A. Cardona, “Ufls and smart load for frequency regulation in electrical power system: A review,” *IEEE Access*, vol. 11, pp. 110 967–110 984, 2023.

## Bibliography

- [47] S. Gordon, C. McGarry, J. Tait, and K. Bell, “Impact of low inertia and high distributed generation on the effectiveness of under frequency load shedding schemes,” *IEEE Transactions on Power Delivery*, vol. 37, no. 5, pp. 3752–3761, 2022.
- [48] A. Husham, I. Kamwa, and H. Suprême, “Koopman model predictive control based load modulation for primary frequency regulation,” *IET Generation, Transmission & Distribution*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:266108061>
- [49] S. S. Silva Jr and T. M. Assis, “Adaptive underfrequency load shedding in systems with renewable energy sources and storage capability,” *Electric Power Systems Research*, vol. 189, p. 106747, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378779620305502>
- [50] C. Chen, M. Cui, F. Li, S. Yin, and X. Wang, “Model-free emergency frequency control based on reinforcement learning,” *IEEE Transactions on Industrial Informatics*, vol. 17, no. 4, pp. 2336–2346, 2021.
- [51] D. Lekshmi J, Z. H. Rather, and B. C. Pal, “Online estimation of disturbance size and frequency nadir prediction in renewable energy integrated power systems,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1126–1137, 2024.
- [52] J. Li, S. Chen, X. Wang, and T. Pu, “Load shedding control strategy in power grid emergency state based on deep reinforcement learning,” *CSEE Journal of Power and Energy Systems*, vol. 8, no. 4, pp. 1175–1182, 2022.
- [53] H. Chen, J. Zhuang, G. Zhou, Y. Wang, Z. Sun, and Y. Levron, “Emergency load shedding strategy for high renewable energy penetrated power systems based on deep reinforcement learning,” *Energy Reports*, vol. 9, pp. 434–443, 2023, 2022 2nd International Joint Conference on Energy and Environmental Engineering. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352484723002640>

## Bibliography

- [54] C. Wang, S. Mei, Q. Dong, R. Chen, and B. Zhu, “Coordinated load shedding control scheme for recovering frequency in islanded microgrids,” *IEEE Access*, vol. 8, pp. 215 388–215 398, 2020.
- [55] X. Chen, M. Zhang, Z. Wu, L. Wu, and X. Guan, “Model-free load frequency control of nonlinear power systems based on deep reinforcement learning,” *IEEE Transactions on Industrial Informatics*, pp. 1–9, 2024.
- [56] Y. Zhao, T. Liu, and D. J. Hill, “A multi-agent reinforcement learning based frequency control method with data-enabled predictive control guided policy search,” in *2022 IEEE Power Energy Society General Meeting (PESGM)*, 2022, pp. 1–5.
- [57] Sklearn. (2019) Scikit learn user guide release 0.21.3., 2019. [Online]. Available: [https://scikit-learn.org/0.21/\\_downloads/scikit-learn-docs.pdf](https://scikit-learn.org/0.21/_downloads/scikit-learn-docs.pdf)
- [58] O. Alimi, K. Ouahada, and A. Abu-Mahfouz, “A review of machine learning approaches to power system security and stability,” *IEEE Access*, vol. PP, pp. 1–1, 06 2020.
- [59] L. W. L. Duchesne, E. Karangelos, “Recent developments in machine learning for energy systems reliability management,” vol. 108, no. 9, pp. 1656–1676, 2020.
- [60] A. Nath, R. Sen Biswas, and A. Pal, “Application of machine learning for on-line dynamic security assessment in presence of system variability and additive instrumentation errors,” 09 2019.
- [61] S. M. Miraftebzadeh, M. Longo, F. Foiadelli, M. Pasetti, and R. Igual, “Advances in the application of machine learning techniques for power system analytics: A survey,” *Energies*, vol. 14, p. 4776, 08 2021.
- [62] P. N. Papadopoulos and J. V. Milanović, “Probabilistic framework for transient stability assessment of power systems with high penetration of renewable generation,” *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 3078–3088, 2017.

## Bibliography

- [63] S. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *31st Conference on Neural Information Processing Systems (NIPS 2017)*, 12 2017.
- [64] S. Lundberg, G. Erion, H. Chen, A. DeGrave, J. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, “From local explanations to global understanding with explainable ai for trees,” *Nature Machine Intelligence*, vol. 2, 01 2020.
- [65] L. Scott. (2018) Deepexplainer. [Online]. Available: <https://shap-lrjball.readthedocs.io/en/latest/generated/shap.DeepExplainer.html>
- [66] R. Zimmerman and C. Murillo-Sanchez. (2020) Matpower users manual version 7.1. [Online]. Available: <https://matpower.org/docs/MATPOWER-manual-7.1.pdf>
- [67] A. D’Amour, “On multi-cause causal inference with unobserved confounding: Counterexamples, impossibility, and alternatives,” 02 2019.
- [68] M. Lu, H. Nagarajan, R. Bent, S. D. Eksioglu, and S. J. Mason, “Tight piecewise convex relaxations for global optimization of optimal power flow,” in *2018 Power Systems Computation Conference (PSCC)*, 2018, pp. 1–7.
- [69] S. Gopinath and H. L. Hijazi, “Benchmarking large-scale acopf solutions and optimality bounds,” in *2022 IEEE Power & Energy Society General Meeting (PESGM)*, 2022, pp. 1–5.
- [70] R. I. Hamilton and P. N. Papadopoulos, “Using shap values and machine learning to understand trends in the transient stability limit,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1384–1397, 2024.
- [71] J. J. Q. Yu, D. J. Hill, A. Y. S. Lam, J. Gu, and V. O. K. Li, “Intelligent time-adaptive transient stability assessment system,” *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 1049–1058, 2018.
- [72] J. Gonzalez, P. N. Papadopoulos, J. V. Milanović, G. Peskir, and J. Moriarty, “Risk-constrained minimization of combined event detection and decision time for

## Bibliography

- online transient stability assessment,” *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4564–4572, 2021.
- [73] WECC Renewable Energy Modeling Task Force. (2014) Wecc solar plant dynamic modeling guidelines. [Online]. Available: <https://www.wecc.org/Reliability/WECC%20Solar%20Plant%20Dynamic%20Modeling%20Guidelines.pdf>
- [74] Z. J. Zhang, P. T. Mana, D. Yan, Y. Sun, and D. K. Molzahn, “Study of active line flow constraints in dc optimal power flow problems,” in *2020 SoutheastCon*, 2020, pp. 1–8.
- [75] S. Ghamizi, A. Ma, J. Cao, and P. Cortes, “Opf-hgmn: Generalizable heterogeneous graph neural networks for ac optimal power flow,” 03 2024.
- [76] Y. Liu, S. You, J. Tan, Y. Zhang, and Y. Liu, “Frequency response assessment and enhancement of the u.s. power grids toward extra-high photovoltaic generation penetrations—an industry perspective,” *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 3438–3449, 2018.
- [77] C. Tjandraatmadja, R. Anderson, J. Huchette, W. Ma, K. Patel, and J. P. Vielma, “The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification,” in *2020 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada, 2020*.
- [78] C. Zhao and X. Li, “Linearization of relu activation function for neural network-embedded optimization: Optimal day-ahead energy scheduling,” 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2310.01758>
- [79] S. K. Jayaweera, *Markov Decision Processes*, 2015, pp. 207–268.
- [80] T. Zhou, Y. Wang, Y. Xu, Q. Wang, and Z. Zhu, “Applications of reinforcement learning in frequency regulation control of new power systems,” in *2022 International Conference on Cyber-Physical Social Intelligence (ICCSI)*, 2022, pp. 501–506.

## Bibliography

- [81] E. H. Sumiea, S. J. Abdulkadir, H. S. Alhussian, S. M. Al-Selwi, A. Alqushaibi, M. G. Ragab, and S. M. Fati, “Deep deterministic policy gradient algorithm: A systematic review,” *Helvion*, vol. 10, no. 9, p. e30697, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405844024067288>
- [82] A. Mehrzad, M. Darmiani, Y. Mousavi, M. Shafie-Khah, and M. Aghamohammadi, “A review on data-driven security assessment of power systems: Trends and applications of artificial intelligence,” *IEEE Access*, vol. 11, pp. 78 671–78 685, 2023.
- [83] F. Zelaya-Arrazabal, T. Thacker, H. Pulgar-Painemal, and Z. Guo, “Supplementary primary frequency control through deep reinforcement learning algorithms,” in *2023 North American Power Symposium (NAPS)*, 2023, pp. 1–6.
- [84] J. Ekomwenrenren, E. Farantatos, M. Patel, A. Haddadi, and L. Zhu, “Data-driven fast frequency control,” in *11th Bulk Power Systems Dynamics And Control Symposium (IREP)*, July 25-30, 2022.
- [85] Rognes Solheim, B. Annfelt Høverstad, and M. Korpås, “Deep reinforcement learning applied to monte carlo power system reliability analysis,” in *2023 IEEE Belgrade PowerTech*, 2023, pp. 01–08.
- [86] J. Li, S. Chen, X. Wang, and T. Pu, “Load shedding control strategy in power grid emergency state based on deep reinforcement learning,” *CSEE Journal of Power and Energy Systems*, vol. 8, no. 4, pp. 1175–1182, 2022.