Enhanced Frequency Control in Low-Inertia Power Systems Using Learning-Based Estimation and Prediction

PhD Thesis

Lu Shen

Department of Electronic and Electrical Engineering

University of Strathclyde, Glasgow

May 28, 2026

Abstract

The growing penetration of Inverter-Based Resources (IBRs) has significantly transformed power system dynamics, leading to reduced system inertia and increasingly complex system behaviour. Consequently, traditional frequency control practices based on offline studies and relatively fixed contingency assumptions are becoming increasingly inadequate for maintaining frequency stability under rapidly changing system conditions. In particular, existing approaches that rely on analytical models and offline simulations are limited in their ability to capture complex and volatile system dynamics and to adapt in a timely manner. Furthermore, frequency control actions are predominantly reactive, typically being triggered only after frequency deviations have already occurred. While this is generally adequate for moderate frequency events, it may prove insufficient during major power imbalance disturbances. Their effectiveness further depends on fast and accurate measurements of key system parameters, particularly inertia, Rate of Change of Frequency (RoCoF), and power imbalance size, which are difficult to estimate reliably in low-inertia power systems.

Recent advances in learning-based techniques have created new opportunities for addressing a range of frequency stability challenges arising from the large-scale integration of IBRs. However, existing research has largely focused on direct control implementations or narrowly scoped prediction tasks, with comparatively limited attention given to the development of proactive, situation-aware, and practically deployable decision-support methodologies. As a result, a clear gap remains in translating learning-based modelling advances into effective real-time operational support for low-inertia power systems. This thesis addresses this gap by leveraging the real-time estimation and prediction capabilities of learning-based methods to enhance frequency control in low-inertia power systems, explicitly complementing and strengthening existing control and operational practices rather than replacing them. Using realistic data and tools available to system operators, the proposed approaches enhance frequency control across multiple stages by improving risk-informed scheduling, disturbance estimation, and adaptive emergency response.

In this thesis, frequency control is conceptualised as three interconnected stages, i.e. pre-event scheduling, the power imbalance event stage, and post-event frequency containment. Learning-based techniques are applied across all three stages to improve the accuracy and effectiveness of frequency control, thereby establishing a more comprehensive and practical estimation and prediction capabilities for future power systems with low-inertia and increasingly complex operating conditions.

In the pre-event scheduling stage, this work proposes a risk-aware prediction component to capture probabilistic frequency nadir behaviour across diverse operating conditions. Instead of relying on deterministic or single-point nadir predictions, the proposed approach explicitly quantifies prediction uncertainties, particularly those arising from real-time inertia measurements. These uncertainty-informed intervals are directly integrated into the optimisation process, which enables the model to ensure that scheduling decisions account for the most severe potential prediction errors. Consequently, the approach facilitates a balance between minimising procurement costs and maintaining stringent frequency stability requirements, while preserving security margins even when prediction accuracy is degraded.

During the power imbalance event stage, which corresponds to the initial transient following a disturbance, the rapid and reliable estimation of RoCoF and power imbalance size is critical for supporting timely frequency control actions. A Discrete Wavelet Transform (DWT)-Support Vector Regression (SVR) RoCoF estimator is developed to balance estimation accuracy and computational speed under noisy and fast-varying measurement conditions. The estimator addresses limitations of conventional RoCoF measurement techniques, which are often sensitive to transients and measurement noise. In addition, an Extreme Gradient Boosting (XGBoost) power imbalance size estimation method is introduced to overcome the limitations of traditional swing equation-based approaches that rely on accurate system inertia information and reliable RoCoF measurements. The proposed method infers the active power imbalance directly from frequency measurements and system operating conditions. These estimation modules provide essential situational information during contingency events to support more effective subsequent frequency control decisions.

To further safeguard the system during disturbances, the post-event frequency containment stage is addressed through the real-time prediction of frequency behaviour and an adaptive emergency control scheme it enables. Based on the predicted frequency behaviour and system state information, the controller adaptively adjusts load shedding thresholds and stages in real-time to contain frequency deviations, which avoid frequency decline beyond critical thresholds during severe disturbances.

The three stages are linked through a progressive, learning-based decision-support structure across different time scales, i.e. pre-event scheduling is guided by probabilistic assessments of frequency nadir risk, disturbance estimation during the initial transient provides timely and detailed system state information, and post-event frequency containment relies on adaptive control decisions informed by predicted frequency behaviour. The proposed methodologies across all three stages are validated through extensive real-time simulations, Hardware-in-the-Loop (HiL) experiments and real world recorded frequency event data. The results demonstrate that learning-based estimation and prediction enhance frequency control performance in low-inertia power systems by improving scheduling robustness, real-time disturbance awareness, frequency nadir containment compared with conventional approaches.

Acknowledgements

I would like to express my sincere gratitude to everyone who has supported and guided me throughout this PhD research.

First and foremost, I am deeply grateful to Professor Qiteng Hong, who has consistently provided insightful direction and devoted considerable time to improving my work. His exceptional academic insight, patience, and consistently thoughtful feedback were invaluable throughout every stage of my PhD studies.

I would also like to sincerely thank Professor Campbell Booth for his support and for providing the resources and environment to conduct this research.

My appreciation also goes to my colleagues and friends, whose encouragement and companionship have made this journey enjoyable.

Finally, I wish to express my deepest gratitude to my parents. Their unconditional support, understanding, and belief in me have been my greatest source of strength, and this work would not have been possible without them.

Contents

List of Abbreviations

AI	Artificial Intelligence
AGC	Automatic Generation Control
ANN	Artificial Neural Network
BESS	Battery Energy Storage System
CHP	Combined Heat and Power
CoI	Center of Inertia
DC	Dynamic Containment
DER	Distributed Energy Resource
DFFR	Dynamic Firm Frequency Response
DFT	Discrete Fourier Transform
DM	Dynamic Moderation
DR	Dynamic Regulation
DWT	Discrete Wavelet Transform
ECDF	Empirical Cumulative Distribution Function
EFR	Enhanced Frequency Response
FFR	Fast Frequency Response

List of Abbreviations

GB	Great Britain
HiL	Hardware-in-the-Loop
HVDC	High Voltage Direct Current
IBR	Inverter-Based Resource
KF	Kalman Filter
LFDD	Low Frequency Demand Disconnection
LoM	Loss of Main
LS	Least Squares
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MFR	Mandatory Frequency Response
ML	Machine Learning
MOS	Model Output Statistics
MPPT	Maximum Power Point Tracking
MPIW	Mean Prediction Interval Width
NESO	National Energy System Operator
OOB	Out-of-Bag
PCC	Point of Common Coupling
PICP	Prediction Interval Coverage Probability
PLL	Phase-Locked Loop
PMU	Phasor Measurement Unit

List of Abbreviations

PV	Photovoltaic
R^2	Coefficient of Determination
RBF	Radial Basis Function
RES	Renewable Energy Resource
RF	Random Forest
RMSE	Root Mean Square Error
RoCoF	Rate of Change of Frequency
RTDS	Real-Time Digital Simulator
SEL	Schweitzer Engineering Laboratories
SFFR	Static Firm Frequency Response
SFFR	System Frequency Response
SG	Synchronous Generator
SHAP	SHapley Additive exPlanations
SoC	State of Charge
SVR	Support Vector Regression
TNAE	Threshold-Normalised Absolute Error
UDP	User Datagram Protocol
UFLS	Under Frequency Load Shedding
VSM	Virtual Synchronous Machine
XGBoost	Extreme Gradient Boosting
ZCD	Zero-Crossing Detection

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Context

1.1.1 Evolvement of Global Energy Landscape

Power systems worldwide are currently undergoing profound structural transformation, which is driven primarily by the urgent need for decarbonising the energy supply and the continuous growth in electricity demand [1]. The ambitious decarbonisation targets have accelerated the large-scale deployment of Renewable Energy Resources (RESs) globally [2]. The IEA Sustainable Development Scenario suggests that the average annual share of renewable energy in total generation will need to reach 45% by 2040 in order to achieve the associated carbon reduction targets [3]. In Great Britain (GB), the majority of renewable capacity growth is constituted by wind and solar photovoltaics (PV) [4]. According to [5], renewable generation surged from under 5% in 2004 to a historic record of 50.4% in 2024. This expansion is primarily due to the rapid increase in wind generation, which reached an all time high in 2024 to account for 29.2% of the total generation mix, while solar output continues to grow steadily. The sustained high rate of renewable capacity deployment is expected to continue throughout this decade [6]. Furthermore, these new generation capacities are being connected to the grid at a wide range of locations, which is transforming the power system's structure from centralised to increasingly distributed architectures. Large-scale renewable plants integrate with the high-voltage AC and DC transmission networks, while increasing numbers of small installations, known as Distributed Energy Resources (DERs), are embedded in distribution networks or installed behind the meter at customer premises [3]. These RESs are typically intermittent in nature and spread out geographically, so they make the generation output more volatile and variable than previously.

In parallel with the unprecedented transformation in the generation patterns, electricity consumption worldwide has more than tripled since 1980 and is expected to grow even faster in the coming decades as end uses historically powered by fossil fuels become increasingly electrified [7]. Electrification replaces fuel-based equipment with electric technologies across sectors such as transport, heating and industry, thereby transferring more energy consumption onto the grid [8]. Consequently, power systems are expected to experience higher peak demand and faster changes in electricity consumption during daily operation and across seasons.

1.1.2 High Penetration of IBRs

The massive transformation of the energy landscape has led to the rapid proliferation of power electronic inverters across all levels of the grid. Conventional power systems have been dominated by fossil fuel-based thermal generation using Synchronous Generators (SGs) for electricity production and grid interfacing. In contrast, renewable generation, e.g. PV and modern wind turbines, typically require inverters to be connected to the AC grid. Furthermore, Battery Energy Storage Systems (BESS), High Voltage Direct Current (HVDC) links, electric vehicle charging infrastructure, and heat pumps also introduce large numbers of controllable inverters to the system [9]. These technologies are collectively referred to as IBRs. Their continued deployment is progressively displacing conventional synchronous generation, thus reducing the number of online rotating machines and accelerating the transition from SG-dominated to IBR-dominated power systems [6], [10]. As a result, power system behaviour is increasingly shaped by software-based control rather than by inherent physical response.

SGs possess intrinsic physical properties that are advantageous for maintaining system stability. In particular, the aggregate kinetic energy stored in their rotating masses, which directly links to system inertia, is inherently provided by SGs and plays a critical role in maintaining system frequency stability. In contrast, IBRs connect to the grid through power electronics that decouple the energy source from the AC frequency. As a result, they lack a synchronously rotating mass that can directly exchange kinetic energy with the system during disturbances. Without a physical electromechanical coupling, IBRs do not inherently provide inertia. Their frequency support is therefore limited to synthetic or Fast Frequency Response (FFR), which contributes to frequency stability but does not replace the rotational inertia contributed by online synchronous machines [11]. Consequently, as the number of online synchronous machines decreases, the available system inertia decreases as well, which will lead to significantly reduced inertia support during disturbances [12].

Furthermore, with the energy resources and loads becoming more diverse, uncertain, and volatile, the operating characteristics are being fundamentally reshaped and becoming increasingly more complex. System planning and operation has been evolving in order to in order to maintain the same level of reliability and security during the energy transition process. Among the various measures that are required, faster and new types of frequency control strategies have been explored and deployed. These may involve coordination of BESS, hybrid plants combining generation and storage, and active resources from flexible loads and prosumers [13], [14]. In turn, the growing integration of frequency responsive resources further increases the complexity of system behaviour, particularly during contingency events.

1.1.3 Challenge of Reduced System Inertia

System frequency serves as a common reference across the entire network, and is an indicator of real-time balance between generation and demand. Frequency stability is one of the most critical aspects for secure operation of power systems, both during steady states with normal demand variations, and more importantly during contingency events, e.g. sudden loss of generation or demand [15].

When a power imbalance occurs, the stored kinetic energy in online synchronous machines is released instantaneously to counteract the deviation [16], [17]. The increasing penetration of IBRs will significantly reduce system inertia, leading to higher initial RoCoF and a faster frequency deviation toward critical limits. This tightens the time window for effective frequency control actions, making it more difficult for conventional frequency control schemes to respond effectively [18]–[20].

This dynamic shift creates significant operational risks for grid stability. When frequency violates critical thresholds, Low Frequency Demand Disconnection (LFDD) schemes may be activated automatically to shed load and prevent uncontrolled frequency collapse [21], [22]. At the same time, high RoCoF levels may trigger Loss of Mains (LoM) protection in DERs, causing the sudden disconnection of embedded generation and exacerbating the original power imbalance [23], [24]. Rapid frequency excursions can also impose mechanical and thermal stress on synchronous machines and industrial processes, potentially leading to protective shutdowns. Since most networks are not designed to withstand such sequential outages, these effects can accumulate and escalate into large-scale disturbances or even system-wide blackouts [25].

As the stabilising influence historically provided by synchronous machines continues to decrease, modern power systems becomes increasingly susceptible to frequency instability. This evolving operating environment presents urgent needs and higher requirements on the operational schemes and planning standards. In particular, dedicated fast frequency services, such as synthetic inertia and FFR, are required to ensure frequency remains within safe operational limits in grids with high IBR penetration [26].

1.1.4 Challenge of Increasingly Complex System Behaviour

With the increasing integration of IBRs, the overall system behaviour is also becoming more complex. Conventional system frequency regulation primarily relies on turbine governors of SGs, whose dynamic characteristics are well known and relatively straightforward to model. However, in IBR-dominated systems, frequency response is provided by a more diverse set of resources with different control mechanisms. These resources include conventional SGs, such as gas turbines, hydropower plants, and Combined Heat and Power (CHP) units, as well as other resources, e.g. demand-side response and BESSs.

IBRs exhibit diverse frequency response behaviours determined by their specific energy sources and internal control architectures. These control algorithms differ widely across manufacturers and may even vary within the same device class. For instance, some IBRs are configured to provide synthetic inertia [27], while others implement might use threshold-based services such as firm frequency response or droop-based services such as Dynamic Containment (DC) in the GB system [28]. In IBR-dominated systems, the timescales of frequency dynamics and the required response times can range from seconds to hundreds of milliseconds, as reduced system inertia and the fast IBR controls can cause interactions to occur over very short timescales [29].

The diversity in frequency response characteristics leads to significantly more complex system behaviour during disturbances. When different control mechanisms, e.g. droop control, synthetic inertia, and threshold-based services, operate independently without a coordination architecture, it becomes extremely challenging to construct an accurate mathematical model to describe the frequency behaviour of the system network. Furthermore, the heightened system complexity and the prevalence of fast-acting transients present significant hurdles for real-time monitoring. Non-linear interactions among multiple control loops make it difficult to reliably observe system behaviour and extract key parameters during the immediate post-event period. Consequently, accurately quantifying and scheduling the required frequency response becomes increasingly challenging. This complicates the implementation of emergency frequency control, e.g. the Under Frequency Load Shedding (UFLS) schemes, and makes it harder to ensure that frequency remains within safe operational limits.

1.2 Research Motivation

This research is motivated by the increasing challenges of frequency control associated with reduced system inertia and more complex system behaviour as discussed above. Recent advances in learning-based and artificial intelligence have created new opportunities for for addressing a range of frequency stability challenges arising from the large-scale integration of IBRs. However, existing research has largely focused on direct control implementations or narrowly scoped prediction tasks, with comparatively limited attention given to the development of proactive, situation-aware, and practically deployable decision-support methodologies. In response to this gap, this thesis develops learning-based models for enhanced frequency control in low-inertia power systems, explicitly designed to strengthen existing control and operational practices rather than replace them. Using realistic data and tools available to system operators, the proposed approaches enhance frequency control across multiple stages by improving risk-informed scheduling, disturbance estimation, and adaptive emergency response.

In general, frequency control scheme can be conceptualised as three stages, i.e. pre-event scheduling, the power imbalance event stage, and post-event frequency containment, as illustrated in Figure 1.2.1. The pre-event scheduling stage concerns preparing adequate reserves and assessing the system's ability to reflect credible contingencies. The power imbalance event stage focuses on rapidly obtaining the key system information that is required for the immediate post-event dynamics. The final post-event frequency containment stage is responsible for deploying corrective actions that arrest the frequency deviation. The rapid integration of IBRs has led to specific challenges for each of the stages, which are discussed in detail as follows.

Stage I: Pre-Event Frequency Response Service Scheduling Considering Uncertainty

- **Existing Methods:** Frequency service scheduling has been managed through offline planning and static reserve allocations. Stability assessments are typically based on analytical tools derived from the classical swing equation or offline simulation models to evaluate the system response to potential disturbances. Operators generally assume a fixed power imbalance (ΔP) that represents the largest credible loss in the system, and use this as a static assumption to estimate the expected frequency nadir. The resulting nadir prediction guides the scheduling of frequency response services.
- **Limitations:** Offline simulations are slow and computationally demanding, making them unsuitable for near real-time reserve schedule updates required under

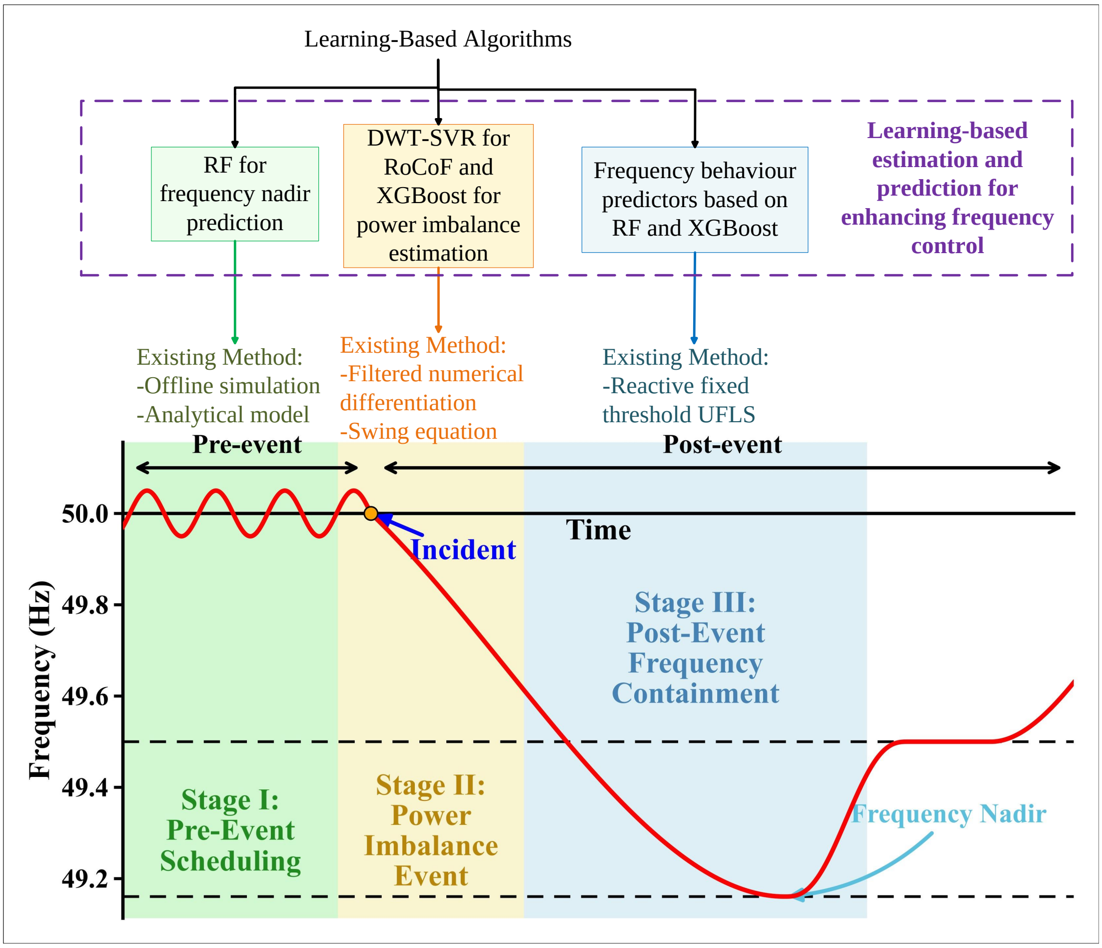

Fig. 1.2.1: Stages of frequency stability and the associated research motivations

more volatile system conditions. In addition, mathematical models for nadir prediction are increasingly unable to capture the complexity of present day system dynamics [30], [31]. Modern frequency response involves multiple ancillary services that are activated through different mechanisms, with distinct thresholds, activation logics and interactions, which cannot be represented reliably using simplified analytical formulations or fixed scenario studies. Conventional planning also assumes fixed dynamic parameters, even though actual inertia and damping vary on short time scales and cannot be captured accurately through offline studies. As a result, these limitations cause static reserve allocations to diverge from the actual nadir risk associated with the current system state, leading to either over-provisioning or inadequate scheduling.

Stage II: Accurate and Rapid Estimation of Disturbance Characteristics During the Initial Transient Period

- **Existing Methods:** Following a disturbance, effective frequency control depends on the fast and accurate monitoring of critical power system quantities. Among these, the RoCoF and power imbalance size are two particularly important parameters. However, achieving both accurate and timely RoCoF measurement remains a longstanding and unresolved challenge. Presently, RoCoF is typically measured using Phasor Measurement Unit (PMU) or obtained through filtered numerical differentiation over a short observation window [32]. Because RoCoF is derived from the second-order derivative of the voltage phase angle, its measurement is inherently sensitive to noise, particularly during disturbances. Existing methods typically rely solely on signal processing techniques, in which filtering is applied to suppress noise. Similarly, the estimation of the active power imbalance size is conventionally performed using the system swing equation, which establishes a physical relationship between the power deficit, the aggregate system inertia, and the measured RoCoF [33]. By combining pre-calculated or assumed inertia values with the real-time RoCoF measurement immediately after a disturbance, system operators can analytically derive the power imbalance size.

- **Limitations:** Conventional RoCoF estimation methods using window-based numerical differentiation are highly sensitive to measurement noise and require filtering, which introduces latency. These limitations are particularly problematic during the very short initial transient following a disturbance, when rapid and accurate RoCoF information is most critical. The 9 August 2019 GB disturbance demonstrated consequences of these limitations, where RoCoF-based protection triggered unintended generation disconnections and a deeper frequency nadir [34]. Estimating the magnitude of the active power imbalance using the swing equation also presents significant challenges, as it depends on accurate knowledge of both RoCoF and system inertia. The limitations of RoCoF measurement have already been highlighted. System inertia, meanwhile, is not always known accurately.

As a result, potentially affecting downstream decision making functions such as loss-of-mains protection, load shedding actions, and broader situational awareness tasks, including power imbalance estimation [35].

Stage III: Emergency Frequency Control During Severe Power Imbalance Events

- **Existing Methods:** In the containment stage, system-level emergency control mechanisms must arrest the frequency decline to maintain stability. Conventional emergency control measures, such as UFLS, typically rely on reactive, fixed-threshold methods with predefined time delays determined through representative offline planning studies [36].
- **Limitations:** These reactive, fixed-threshold schemes do not provide the real-time adaptability required for a system with rapidly evolving dynamics. Because they lack predictive capability, such schemes often face a critical balance between response speed and accuracy. They may either act too late to arrest a sharp frequency decline or alternatively shed excessive load due to a limited understanding of the disturbance severity. Furthermore, static planning fails to capture the complex interactions among fast-acting services, such as FFR and synthetic inertia, which can lead to unpredictable system behaviour when multiple resources operate simultaneously. These issues ultimately reduce the reliability and efficiency of emergency actions in low-inertia grids [37].

1.3 Research Aim and Objectives

These challenges outlined above highlight the need for a holistic approach to frequency control, in which all three critical stages are jointly designed to address emerging issues introduced by increasing integration. Therefore, this PhD research aims to leverage the powerful real-time estimation and prediction capabilities of learning-based techniques to enable more effective frequency control in low-inertia power systems through three interconnected stages: operational scheduling, real-time situational awareness, and emergency control. It seeks to enable more proactive, accurate, and uncertainty-aware decision making by anticipating frequency risk, extracting reliable information during the initial transient period, and adapting corrective actions to evolving system conditions. This will enable frequency control to transition from reactive, model-constrained approaches toward predictive and operationally compatible frequency stability management in future power systems.

To achieve this aim, the specific objectives of the thesis are as follows:

- Develop a pre-event probabilistic nadir prediction based on RF to model the non-linear mapping between system operating conditions and frequency nadirs, enabling the quantification of uncertainty for security-informed frequency response scheduling.
- Establish real-time algorithms that integrate DWT-SVR for multi-resolution RoCoF estimation, and employ XGBoost to estimate power imbalance size directly from frequency data, thereby eliminating the dependency on explicit inertia and RoCoF information.
- Develop a multi-window predictive control algorithm based on RF and XGBoost to determine short-term frequency behaviour and emergency frequency control, leveraging variable observation windows to optimise the balance between response speed and decision reliability.

1.4 Research Contributions

This thesis makes the following key contributions to knowledge:

- **A risk-aware predictive scheduling scheme for frequency response services (Stage I):** The thesis proposes a proactive scheduling scheme that uses Random Forest (RF) nadir predictions to determine the minimum frequency response reserves required to maintain the frequency nadir above a predefined threshold. By embedding a specified risk tolerance into the predictive model, the proposed approach captures variations in operating conditions and enables more secure and economically efficient scheduling than traditional fixed-allocation approaches.

- **DWT-SVR-based multi-resolution RoCoF estimation (Stage II):** A real-time algorithm has been developed to obtain accurate RoCoF measurements using a DWT-SVR approach. In addition, this method utilises a composite scoring mechanism to adaptively select optimal observation windows, effectively balancing estimation accuracy and response latency to minimise the time delays typically associated with low-pass filtering during fast transients. The algorithm consistently outperforms conventional filtering methods and has been validated using simulation datasets and real world power system event records.

- **A real-time XGBoost-based estimator of power imbalance size (Stage II):** A XGBoost-based estimator is developed to quantify power imbalance size using synchronised frequency measurements and operating condition features. This estimator infers the power imbalance size without requiring real-time RoCoF measurements and direct information of the system inertia. It provides reliable confidence intervals for the estimated values and offers a robust foundation for disturbance assessment.

- **A real-time frequency behaviour prediction using XGBoost and RF-based predictors (Stage III):** A multi-window algorithm integrating RF and XGBoost predictors has been developed to provide real-time frequency behaviour prediction following a disturbance. This approach combines an RF-based estimator for frequency nadir and timing with an XGBoost predictor for short-term frequency evolution. Rather than providing a single nadir estimate, the predicted trajectory characterises the expected frequency trend and evolution over time, providing information on frequency nadir depth, timing, and post-nadir frequency evolution to support informed decision making during the post-event frequency containment stage.

- **An adaptive emergency load shedding scheme based on the prediction of frequency behaviour (Stage III):** An adaptive emergency load shedding scheme is proposed that adaptively determines when and how much load to shed in real-time based on predicted frequency trajectories. The scheme mitigates the risks of maloperation and unnecessary load shedding during rapid transients. Real-time HiL evaluations demonstrate that the proposed prediction-supported control enhances system stability while reducing the economic cost of emergency actions.

1.5 Thesis Overview

This thesis is organised as follows:

Chapter 2 presents a comprehensive review of frequency stability and control in low-inertia power systems. This covers the theoretical foundations and conventional industrial practices for frequency control. Existing methods for frequency nadir prediction, RoCoF estimation, and power imbalance quantification are reviewed, with particular emphasis on the limitations of traditional analytical tools and the swing equation under high IBR penetration. Furthermore, the chapter reviews emergency frequency control mechanisms, among which UFLS is widely adopted as a representative implementation, and analyses their evolution from fixed-threshold designs to adaptive approaches. Recent advancements in data-driven approaches to frequency security are also reviewed, based on which the technical gaps that motivate the predictive and adaptive techniques proposed in this thesis are identified.

Chapter 3 investigates pre-event frequency dynamics by proposing a learning-based prediction for frequency nadir prediction and the integrated scheduling of frequency response services. Using pre-event operating conditions and credible contingency sizes as inputs, the proposed models predict the frequency nadir with calibrated uncertainty. These predictions are embedded in a scheduling formulation that allocates fast reserves and synthetic inertia to satisfy security constraints with probabilistic guarantees. The chapter details feature engineering, model selection, and validation on historical disturbances and synthetic scenarios.

Chapter 4 focuses on real-time system monitoring immediately following the contingency. It develops reliable estimators for RoCoF and online inference of the power imbalance size using PMU measurements affected by noise and latency. The chapter investigates filtering strategies, window-selection methods, model-based observers, and physics-informed estimators that combine data features with physical constraints. The resulting monitors provide fast and reliable frequency characterisation, offering substantially improved performance compared with traditional RoCoF and power imbalance size estimation methods.

Chapter 5 develops a real-time learning-based approach for predicting post-event frequency behaviour and enabling adaptive emergency frequency control. By integrating short-horizon prediction with decision making in a multi-window structure, the approach provides forward-looking situational awareness, anticipates critical frequency violations, and determines the minimum required load shedding while accounting for system delays and disturbance severity. The effectiveness of the proposed approach is validated through extensive simulations and real-time HiL studies.

Chapter 6 summarises the contributions of this thesis and discusses potential directions for future research.

1.6 Publications

List of Conference Papers

- L. Shen, Q. Hong and C. Booth, “Frequency Behaviour Prediction Based on Machine Learning for Future Low-Inertia Power Systems”, *2024 59th International Universities Power Engineering Conference (UPEC)*, Cardiff, United Kingdom, 2024.

- L. Shen, Q. Hong, J. Han, J. Wang and C. Booth, “Real-time Estimation of Power Imbalance During Contingency Events: A Data-Driven Approach”, *2025 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*, Valletta, Malta, 2025.

List of Journal Papers

- L. Shen, Q. Hong and C. Booth, "Adaptive Load Shedding Based on the Real-Time Prediction of the Frequency Trajectory", *IEEE Transactions on Power Systems* (final internal revision).

- L. Shen, Q. Hong and C. Booth, "RoCoF Estimation with Wavelet Decomposition and SVR in Low-Inertia Power Systems", *IEEE Transactions on Smart Grid* (final internal revision).

References

[1] A. M. Saleh, I. Vokony, M. A. Khan, M. Waseem, and A. N. A. Ahmed, "Power system stability in the era of energy transition: Importance, opportunities, challenges, and future directions," *Energy Conversion and Management: X*, vol. 24, p. 100 820, 2024.

[2] A. ur Rehman, M. J. Sanjari, R. M. Elavarasan, and T. Jamal, "Sustainability-aligned pathways for energy transition: A review of low-carbon energy network solutions," *Renewable and Sustainable Energy Reviews*, vol. 226, p. 116 428, 2026.

[3] "Power systems in transition: Challenges and opportunities ahead for electricity security," International Energy Agency (IEA), Paris, Special report, Oct. 2020. [Online]. Available: `https://iea.blob.core.windows.net/assets/cd69028a-da78-4b47-b1bf-7520cdb20d70/Power_systems_in_transition.pdf`.

[4] National Energy System Operator (NESO), *Britain's energy explained: 2025 review*, 2025. [Online]. Available: `https://www.neso.energy/news/britains-energy-explained-2025-review`.

[5] "Digest of united kingdom energy statistics (dukes) 2025: Chapters 1 to 7," Department for Energy Security and Net Zero, Jul. 2025. [Online]. Available: `https://assets.publishing.service.gov.uk/media/68dbe477ef1c2f72bc1e4c4d/DUKES_2025_Chapters_1-7.pdf`.

[6] "Renewables 2024: Analysis and forecast to 2030," International Energy Agency (IEA), Paris, 2024. [Online]. Available: `https://iea.blob.core.windows.net/assets/17033b62-07a5-4144-8dd0-651cdb6caa24/Renewables2024.pdf`.

[7] "Key world energy statistics 2021," International Energy Agency (IEA), Paris, Sep. 2021. [Online]. Available: `https://iea.blob.core.windows.net/assets/52f66a88-0b63-4ad2-94a5-29d36e864b82/KeyWorldEnergyStatistics2021.pdf`.

[8] U.S. Department of Energy, *What is electrification*, `https://www.energy.gov/sites/default/files/2024-04/04-04-2024-oe-voe-what-is-electrification.pdf`, 2024.

[9] "High penetration of power electronic interfaced power sources and the potential contribution of grid forming converters," ENTSO-E Technical Group on High Penetration of Power Electronic Interfaced Power Sources, 2023. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/Publications/SOC/High_Penetration_of_Power_Electronic_Interfaced_Power_Sources_and_the_Potential_Contribution_of_Grid_Forming_Converters.pdf`.

[10] "Status of power system transformation 2018," International Energy Agency, 2018. [Online]. Available: `https://iea.blob.core.windows.net/assets/ede9f1f7-282e-4a9b-bc97-a8f07948b63c/Status_of_Power_System_Transformation_2018.pdf`.

[11] "Reliability guideline: Bps-connected inverter-based resource performance," North American Electric Reliability Corporation (NERC), Atlanta, GA, Sep. 2018. [Online]. Available: `https://www.nerc.com/comm/RSTC_Reliability_Guidelines/Inverter-Based_Resource_Performance_Guideline.pdf`.

[12] J. Matevosyan, J. MacDowell, N. Miller, *et al.*, "A future with inverter-based resources: Finding strength from traditional weakness," *IEEE Power and Energy Magazine*, vol. 19, no. 6, pp. 18–28, 2021.

[13] K. Strunz, K. Almunem, C. Wulkow, M. Kuschke, M. Valescudero, and X. Guillaud, “Enabling 100% renewable power systems through power electronic grid-forming converter and control: System integration for security, stability, and application to europe,” *Proceedings of the IEEE*, vol. 111, no. 7, pp. 891–915, 2023.

[14] “Unlocking the potential of distributed energy resources: Power system opportunities and best practices,” International Energy Agency (IEA), Paris, Technical Report, 2022. [Online]. Available: `https://www.iea.org/reports/unlocking-the-potential-of-distributed-energy-resources/executive-summary`.

[15] P. Kundur, J. Paserba, V. Ajjarapu, *et al.*, “Definition and classification of power system stability ieee/cigre joint task force on stability terms and definitions,” *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1387–1401, 2004.

[16] U. Markovic, O. Stanojev, P. Aristidou, E. Vrettos, D. Callaway, and G. Hug, “Understanding small-signal stability of low-inertia systems,” *IEEE Transactions on Power Systems*, vol. 36, no. 5, pp. 3997–4017, 2021.

[17] C. Seneviratne and C. R. Ozansoy, “Frequency response due to a large generator loss with the increasing penetration of wind/pv generation: A literature review,” *Renewable and Sustainable Energy Reviews*, vol. 57, pp. 659–668, 2016.

[18] Y. Gu and T. C. Green, “Power system stability with a high penetration of inverter-based resources,” *Proceedings of the IEEE*, vol. 111, no. 7, pp. 832–853, 2023.

[19] European Network of Transmission System Operators for Electricity (ENTSO-E), “Inertia and rate of change of frequency (RoCoF),” Brussels, Version 17, Dec. 2020. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Inertia%20and%20RoCoF_v17_clean.pdf`.

[20] R. Azizipanah-Abarghooee, M. Malekpour, M. Paolone, and V. Terzija, “A new approach to the online estimation of the loss of generation size in power systems,” *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 2103–2113, 2019.

[21] National Grid, “Proposed grid code changes to include low frequency demand disconnection relay settings,” GCRP Paper 06/22, Sep. 2006. [Online]. Available: `https://www.neso.energy/document/23986/download`.

[22] ENTSO-E AISBL, “Technical background for the low frequency demand disconnection (lfdd) requirements,” Tech. Rep., Nov. 2014. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20ER/141215_Technical_background_for_LFDD.pdf`.

[23] J. K. P. Desai and V. Makwana, “Protection of synchronous generator,” in *Power Swing Detection and Generator Out-of-Step Protection Under Renewable Power Source Integration*. Singapore: Springer Nature Singapore, 2023, pp. 25–44.

[24] J. Vieira, W. Freitas, W. Xu, and A. Morelato, “Efficient coordination of rocof and frequency relays for distributed generation protection by using the application region,” *IEEE Transactions on Power Delivery*, vol. 21, no. 4, pp. 1878–1884, 2006.

[25] Energy Emergencies Executive Committee (E3C), “Gb power system disruption on 9 august 2019: Final report,” Her Majesty’s Government – Department for Business, Energy & Industrial Strategy, Tech. Rep., Sep. 2019. [Online]. Available: `https://assets.publishing.service.gov.uk/media/5e0e1fa9e5274a0fa7b4d96a/e3c-gb-power-disruption-9-august-2019-final-report.pdf`.

[26] F. Teng, V. Trovato, and G. Strbac, “Stochastic scheduling with inertia-dependent fast frequency response requirements,” *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1557–1566, 2016.

[27] S. D’Arco and J. A. Suul, “Equivalence of virtual synchronous machines and frequency-droops for converter-based microgrids,” *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 394–395, 2014.

[28] B. Lian, A. Sims, D. Yu, C. Wang, and R. W. Dunn, “Optimizing lifepo4 battery energy storage systems for frequency response in the uk system,” *IEEE Transactions on Sustainable Energy*, vol. 8, no. 1, pp. 385–394, 2017.

[29] F. Milano, F. Dörfler, G. Hug, D. J. Hill, and G. Verbič, "Foundations and challenges of low-inertia systems (invited paper)," in *2018 Power Systems Computation Conference (PSCC)*, 2018, pp. 1–25.

[30] L. Liu, W. Li, Y. Ba, J. Shen, C. Jin, and K. Wen, "An analytical model for frequency nadir prediction following a major disturbance," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2527–2536, 2020.

[31] I. Egido, F. Fernandez-Bernal, P. Centeno, and L. Rouco, "Maximum frequency deviation calculation in small isolated power systems," *IEEE Transactions on Power Systems*, vol. 24, no. 4, pp. 1731–1738, 2009.

[32] A. G. Phadke and J. S. Thorp, *Synchronized Phasor Measurements and Their Applications*, 2nd. Cham, Switzerland: Springer, 2017.

[33] Q. Hong, C. Booth, B. Wang, X. Dong, and L. Ji, "Estimation of power imbalance size with consideration of impact of emulated inertia," in *2019 IEEE 8th International Conference on Advanced Power System Automation and Protection (APAP)*, 2019, pp. 730–734.

[34] "Technical report on the events of 9 august 2019," National Grid Electricity System Operator (ESO), Tech. Rep., Sep. 2019. [Online]. Available: `https://www.ofgem.gov.uk/sites/default/files/docs/2019/09/eso_technical_report_-_final.pdf`.

[35] ENTSO-E, "Frequency stability evaluation criteria, methodology, and requirements," European Network of Transmission System Operators for Electricity (ENTSO-E), Brussels, 2018. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/RGCE_SPD_frequency_stability_criteria_v10.pdf`.

[36] B. Delfino, S. Massucco, A. Morini, P. Scalera, and F. Silvestro, "Implementation and comparison of different under frequency load-shedding schemes," in *2001 Power Engineering Society Summer Meeting. Conference Proceedings (Cat. No.01CH37262)*, vol. 1, 2001, 307–312 vol.1.

[37] IEEE Power System Relaying and Control Committee, Working Group C32, "Protection challenges and practices for interconnecting inverter based resources to utility transmission systems," IEEE Power System Relaying and Control Committee (PSRC), System Protection Subcommittee, Piscataway, NJ, 2020. [Online]. Available: `https://www.pes-psrc.org/kb/report/109.pdf`.

Chapter 2

Literature Review

With the increasing penetration of IBRs in modern power systems, the dynamic behaviour of system frequency has changed significantly. Conventional frequency control approaches, originally developed for high-inertia systems, therefore face growing limitations. This chapter reviews both traditional and recent researches to frequency control, with a focus on the challenges arising from low-inertia operation and the associated need for advanced estimation and predictive methods.

The chapter begins by reviewing frequency control strategies in low-inertia power systems, starting with conventional frequency control schemes and emerging industrial practices, and then examining recent research advancements aimed at enhancing frequency control under reduced inertia. Based on review, highlighting the growing need for improved situational awareness and more effective decision support during frequency disturbances.

The chapter reviews frequency nadir prediction methods, which play a critical role in assessing system security following large disturbances. This includes traditional analytical approaches, their extensions, and data-driven techniques that have limitations in low-inertia environments. This is followed by a review of other key parameters that are also critical during disturbance events, such as RoCoF and power imbalance size. Theoretical foundations based on the swing equation are introduced, followed by a review of existing methods, with attention given to challenges arising from measurement noise, parameter uncertainty, and fast transient dynamics. Finally, emergency frequency control schemes are reviewed with a focus on UFLS, covering conventional fixed-threshold approaches, semi-adaptive and adaptive schemes incorporating power imbalance estimation, and more recent data-driven and predictive UFLS strategies. The chapter concludes by synthesising the key findings of the literature and clarifying how they inform the contributions developed in subsequent chapters of this thesis.

2.1 Frequency Control in Low-Inertia Power Systems

2.1.1 Conventional Frequency Control Schemes

To maintain system frequency within the statutory and operational limits, power systems have historically relied on a structured frequency control framework, in which different control schemes are activated through distinct mechanisms. Although the specific implementation varies across different system operators in different countries, such a framework is commonly organised into multiple control layers operating over different timescales.

Within the conventional frequency control paradigm, the broader hierarchical framework is commonly described as consisting of three functional layers [1]:

1. **Primary Frequency Response (Containment):** This layer activates typically within seconds following a disturbance as shown in Figure 2.1.1. Its objective is to arrest the initial frequency decline and stabilise the system at a new steady state value. The response is typically achieved through droop control, where conventional SGs autonomously adjust their active power output in response to locally measured frequency deviation. Turbine governors provide proportional increases in mechanical power according to the magnitude of the frequency deviation. To meet GB grid standards, this response must be initiated within 2 seconds, and is required to achieve 50% delivery within 15 seconds and full delivery within 30 seconds [1]. This initial frequency containment action is critical because it determines the frequency nadir and prevents the system from triggering automatic load shedding [2].

2. **Secondary Frequency Response (Restoration):** Secondary response acts over timescale from roughly 30 seconds to several minutes. Its purpose is to remove steady state frequency errors and bring frequency back to 50Hz. In many power systems, this function is provided through Automatic Generation Control (AGC), which monitors frequency continuously and issues automated adjustments to participating units [3]. Manual dispatch is used only as a backup if frequency remains outside normal operating limits [4], [5]. In contrast, the implementation of secondary frequency restoration in the GB system differs from this conventional AGC-based approach, where it relies primarily on operator-driven dispatch actions within the balancing mechanism, delivered by contracted service providers rather than through a fully automated AGC scheme. In addition, the GB operational requirements specify that secondary response must reach full delivery within 30 seconds and be sustained for up to 30 minutes [6].

3. **Tertiary Frequency Response (Balancing):** Over timescales from minutes to hours, tertiary reserves and re-dispatch actions are deployed to maintain long-term frequency balance and ensure readiness for subsequent disturbances. Although this layer is defined in many international frameworks [7], it is not explicitly identified in GB operational practice.

While these three layers represent the general theoretical definition, in the GB context, the primary frequency control mechanism is specifically achieved through the Mandatory Frequency Response (MFR). MFR provides an automatic adjustment of active power in response to measured frequency deviations to maintain the system frequency within statutory limits (i.e. 49.5Hz – 50.5Hz) and within tighter operational bands (i.e. 49.8Hz – 50.2Hz) [8]. As such, MFR forms the backbone of the broader hierarchical framework, ensuring that the different control layers operate harmoniously to contain, restore, and sustain frequency stability following a disturbance [9].

Figure 2.1.1 illustrates the coordinated operation of these layers from initial containment through to restoration.

However, such a conventional frequency control framework is increasingly chal-

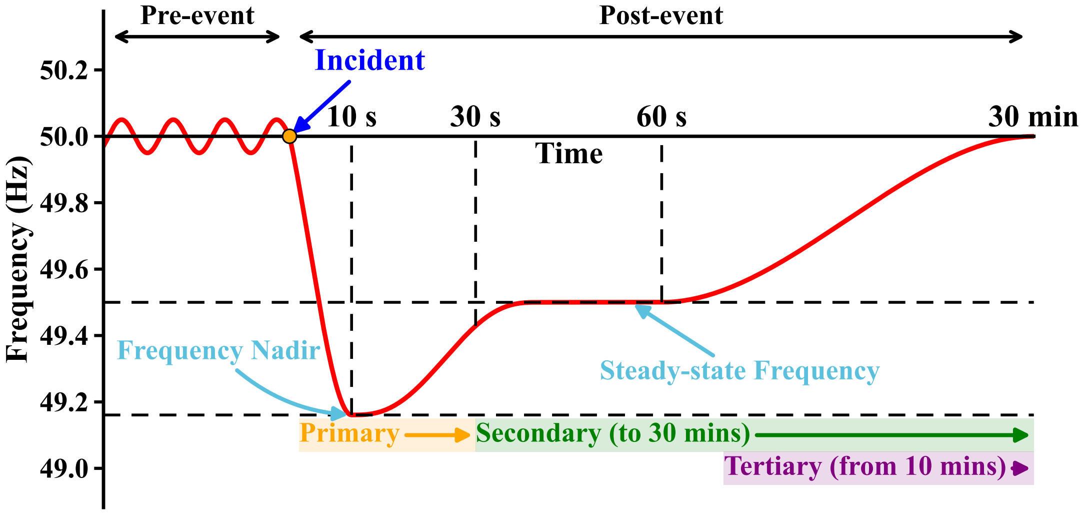

Fig. 2.1.1: System frequency evolution and conventional frequency control layers and their corresponding timescales following a contingency [10]

lenged under low-inertia operating conditions. The reduced capability to withstand fast frequency changes makes it difficult to maintain acceptable nadir and settling performance [11].

2.1.2 Emerging Industrial Practices in Frequency Control

Fast Frequency Response (FFR)

With the continued decrease of system inertia, many system operators have introduced FFR in order to maintain the effective control of frequency [12]. Different countries' definition and requirements for FFR are also different, but in general, FFR refers to frequency control mechanisms that are faster than the conventional primary response [13]. Typically, the activation and full delivery time is within one second instead of multiple seconds, and it is often delivered by different sources from SGs, mostly frequently BESSs [14].

In the GB, FFR was introduced by the National Energy System Operator (NESO) as an essential ancillary service to support real-time frequency management. With the decline of traditional system inertia, FFR enabled both generators and demand-side providers to adjust power output or consumption within short time windows [15].

Historically, the FFR framework consisted of two distinct products: Dynamic Firm Frequency Response (DFFR) and Static Firm Frequency Response (SFFR) [16]. However, the standard 10 seconds full-delivery window associated with DFFR became inadequate under low-inertia conditions, because the frequency nadir often occurs before such a time scale following a contingency. As a result, DFFR was phased out in 2023 to facilitate the establishment of the new suite of pre-fault dynamic frequency response products [17].

While the FFR framework was the primary mechanism for years, the need for faster response led to specialised interventions. Parallel to the evolution of the FFR framework, Enhanced Frequency Response (EFR) was introduced as an early form of sub-second dynamic service designed for non-synchronous technologies in 2016 [18], e.g. BESS. EFR required full delivery within 1 second following a frequency deviation and adopted a proportional response characteristic that is more aggressive than conventional linear droop, enabling rapid active power compensation during the initial containment phase. The service provided two operational configurations, service 1 (wide band) and service 2 (narrow band), which mainly differ in sensitivity settings and deadband widths to protect BESS from excessive cycling due to very small frequency fluctuations.

As illustrated in Figure 2.1.2, EFR follows a normalised reference envelope. The key technical parameters and set points for these envelopes are presented in Table 2.1.1. For intermediate frequency values, the corresponding power output is obtained by linear interpolation between the specified reference points. In addition, service 1 includes a wider insensitivity band around 50 ± 0.05 Hz, while service 2 operates with a narrower and more sensitive band of 50 ± 0.015 Hz.

Although EFR demonstrated the feasibility of rapid frequency control, it is no longer procured as an independent service as it was meant to be trial phase for FFR. The technical principles and fast response requirements initially defined for EFR have instead been incorporated into the modern family of dynamic response products now used in system operation.

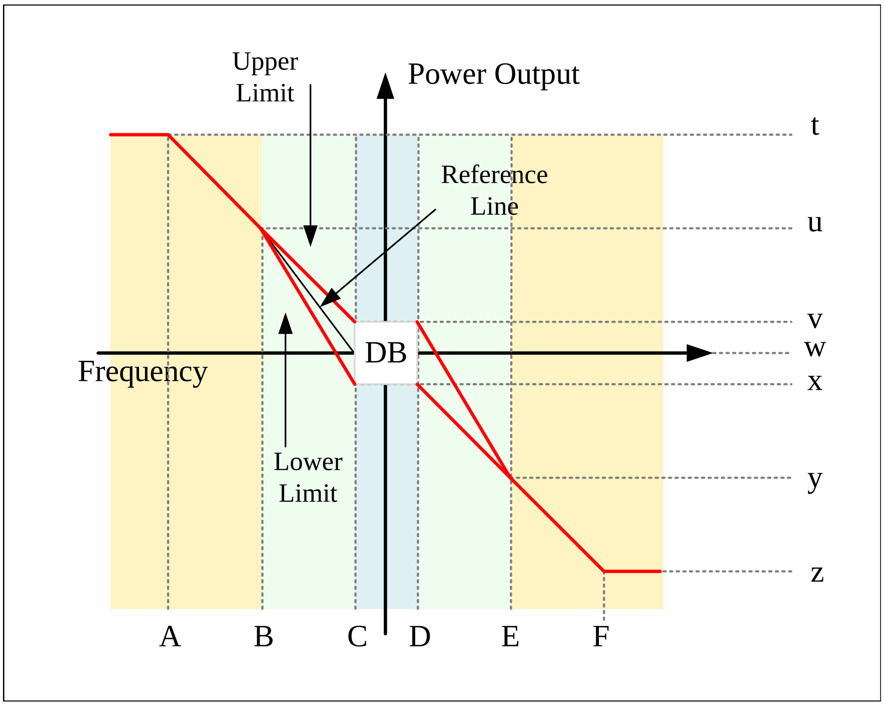

Fig. 2.1.2: EFR dynamic response envelope [19]

Table 2.1.1: Technical set points for EFR service envelopes [19]

Frequency Reference Setpoints			**Power Reference Setpoints**		
Point	**Service 1 (Hz)**	**Service 2 (Hz)**	**Point**	**Service 1 (%)**	**Service 2 (%)**
A	49.50	49.50	t	100	100
B	49.75	49.75	u	44.44444	48.45361
C	49.95	49.985	v	9	9
-	-	-	w	0	0
D	50.05	50.015	x	-9	-9
E	50.25	50.25	y	-44.44444	-48.45361
F	50.50	50.50	z	-100	-100

Current Frequency Ancillary Services in the GB

Based on previously experiences and learnings from the various ancillary services adopted in the past decade, presently, a suite of dynamic response services have been adopted by NESO, including DC, Dynamic Moderation (DM) and Dynamic Regulation (DR) [20].

The control logic and performance requirements of these services are suited to specific types of frequency deviations, as illustrated in Fig. 2.1.3 and summarised in Table 2.1.2. They are developed based on earlier services, e.g. EFR, and replace legacy products to provide continuous, rapid, and proportional active power adjustments to support real-time frequency stability [21].

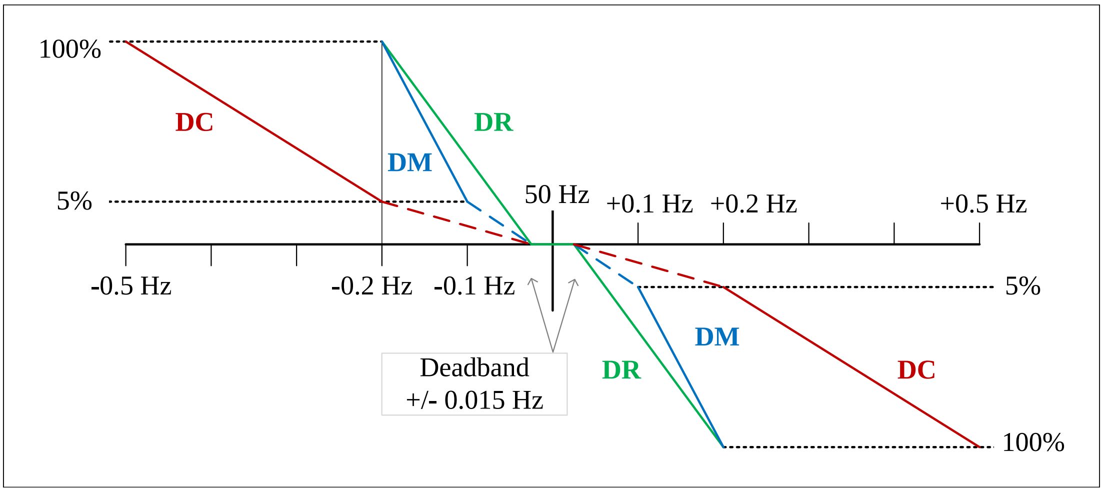

Fig. 2.1.3: Control characteristics of modern dynamic response services DC, DM and DR [22]

Table 2.1.2: Key performance specifications of DC, DM and DR

Specification	Description	DC	DM	DR
Initiation time	Time between a frequency deviation and the start of the service response	0.5 s	0.5 s	2 s
Full delivery time	The time required for the service to reach its rated power after a disturbance	1 s	1 s	10 s
Delivery duration	Minimum time for which energy constrained resources must sustain output	15 min	30 min	60 min

- **Dynamic Containment (DC):** DC is designed to provide very fast frequency containment once frequency deviations exceed a predefined deadband, acting as a primary mechanism to limit frequency excursions following a disturbance. It's main targets are contingency events, e.g. the sudden loss of large generation or demand.

The control follows a service envelope in which the active power adjustment is small for minor deviations and increases steeply as frequency approaches statutory limits. BESS are the principal providers and must reach full output within 1 second after a deviation. The prompt response helps limit the initial frequency drop and mitigates the reduced inertia associated resulted from increasing IBR integration.

- **Dynamic Moderation (DM):** DM is a pre-fault mitigation service that keeps system frequency within the operational range from 49.8 Hz to 50.2 Hz. It acts as a bridge between routine regulation provided by DR and the fast containment provided by DC. DM follows a steep proportional control profile that activates once the frequency deviates beyond the ± 0.015 Hz deadband, providing fast acting pre-fault delivery for particularly volatile periods.
- **Dynamic Regulation (DR):** DR is a continuous service that corrects small imbalances using proportional control within a narrow band of ±0.1 Hz. While DC and DM prioritise containment and near-term mitigation, DR maintains overall energy balance and corrects longer-term trends. DR reacts more slowly than DC and DM, but requires sustained delivery. Providers must comply with State of Charge (SoC) management to ensure a minimum of 60 minutes of continuous response.

While DC, DM and DR provide continuous dynamic regulation across complementary timescales, SFFR remains a critical component of the post-fault frequency recovery process due to its cost effectiveness and simplicity [23]–[25]. According to the most recent NESO guidance, the SFFR service contains two main requirements:

- **Triggering:** SFFR is a non-dynamic relay-based service that operates automatically when frequency reaches a predefined threshold, typically 49.7 Hz. For testing, the trigger is often set at 49.65 Hz with a tolerance of ±0.01 Hz. The minimum entry capacity is 1 MW, provided by a single resource or an aggregation.
- **Delivery:** Providers must achieve full active power within 30 seconds of the trigger

> and sustain output for at least 30 minutes. During this period, the standard deviation of active power error must remain below 2.5% of the contracted capacity to ensure reliable performance.

The combined operation of these services during a contingency is shown in Fig. 2.1.4. Compared with the traditional framework that relied on relatively slower reserves with delivery times up to 30 seconds, the modern suite operates on faster timescales. The rapid responses of DC and DM within 0.5 to 1 second limit the initial frequency decline and avoid unnecessary activation of sensitive LoM protection. In parallel, SFFR provides a discrete backup at 49.7 Hz with a delivery duration of 30 minutes, and DR supports the continuous recovery of the frequency trajectory.

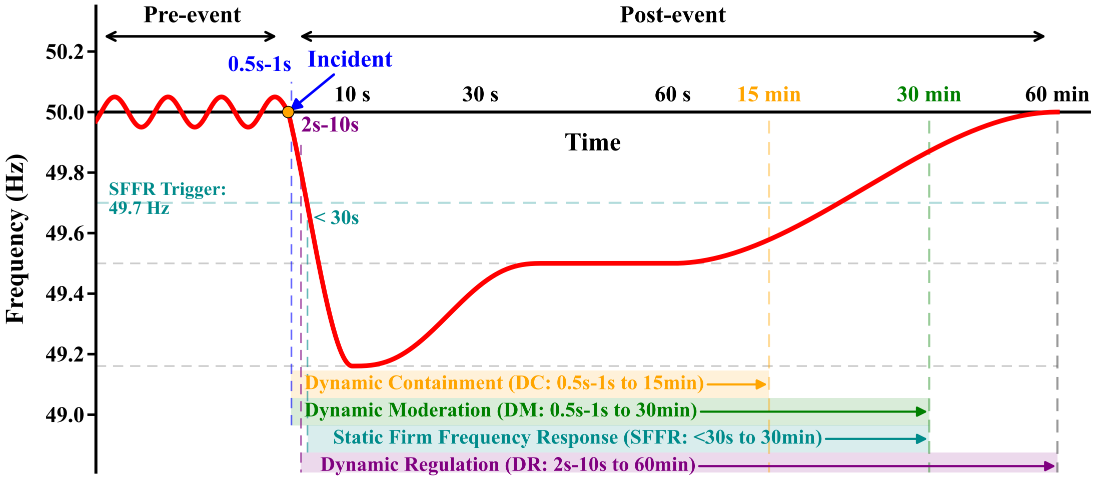

Fig. 2.1.4: System frequency evolution and modern dynamic response services timescales for DC, DM, DR and SFFR

2.1.3 Research Advancements in Frequency Control for Low-Inertia Systems

To address the various challenges associated with frequency control, extensive research work has been conducted in the past decade, which can be categorised into four main areas, i.e. inertia emulation, coordinated frequency support from diverse resources (often within microgrids), advanced demand side response, and data-driven frequency response modelling and control.

Inertia Emulation

Since one of the key challenges in future frequency control is due to IBRs not being able to naturally contribute to system inertia, a major area of research has focused on controlling IBRs to emulate inertia [26].

The fundamental principle of inertia emulation (also known as synthetic or virtual inertia) is to artificially replicate the kinetic energy response of conventional SGs through power electronics [27]. Most proposed approaches are formulated based on the swing equation and regulate the active power output of IBRs as a function of the measured RoCoF. By injecting or absorbing power proportionally to the frequency derivative, the controller provides an emulated inertial response during the initial moments following a disturbance, helping to arrest the initial frequency decline and improve the resulting frequency nadir.

In variable-speed wind energy systems, frequency support is obtained by extracting part of the kinetic energy stored in the rotating blades. When a disturbance occurs, the controller will reduce the rotor speed to release this stored energy to the grid. The overall process is managed by coordinated inertial, primary, and secondary control loops designed to provide effective frequency support while limiting mechanical stress and allowing the rotor speed to ultimately recover [28].

To realise such an emulated inertial response in practice, a specific virtual inertia control loop is typically implemented at the converter level. A representative structural configuration of this mechanism is illustrated in the control block diagram in Fig. 2.1.5. As shown in the figure, the grid frequency is first measured and processed through a low-pass filter to attenuate measurement noise. A derivative operator then estimates the RoCoF, which is scaled by a virtual inertia constant to compute a proportional active power reference. This supplementary control signal is injected into the electrical reference node of the converter, enabling rapid modulation of the turbine's active power output in response to frequency disturbances [29].

Beyond wind turbines, inertia emulation is predominantly implemented in BESSs, where a range of advanced control topologies have been adopted [30]. Early implementations have largely adopted grid-following control schemes, in which Phase-Locked

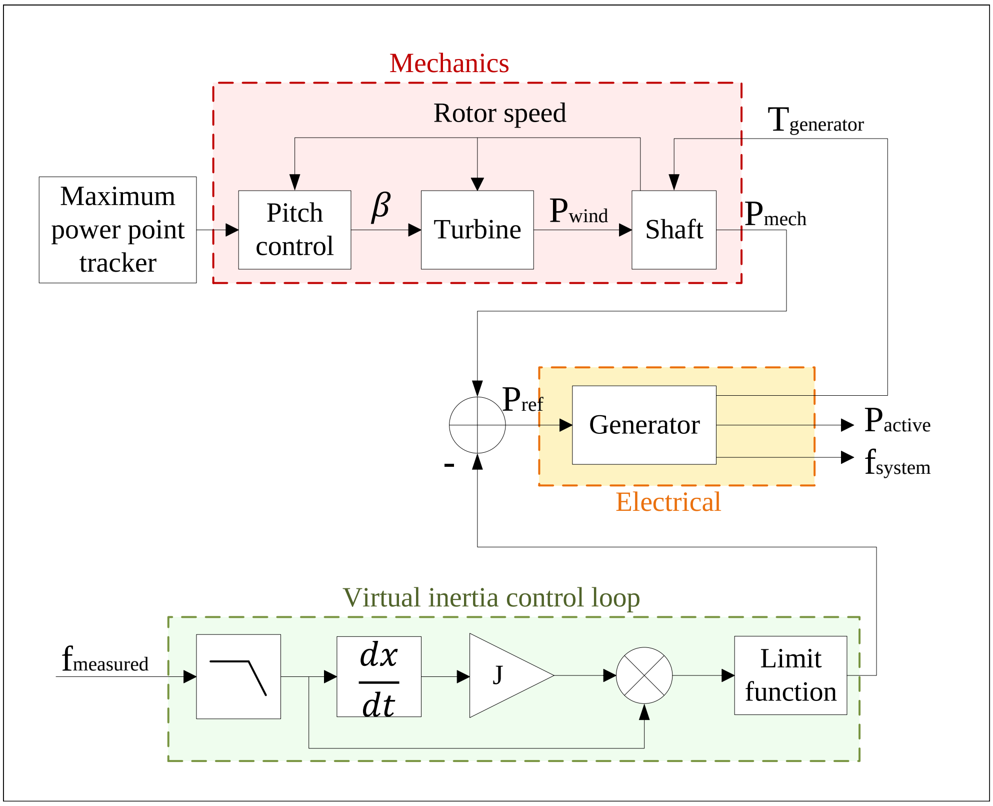

Fig. 2.1.5: Control block diagram of a wind turbine with virtual inertia model [29]

Loops (PLLs) are employed to continuously estimate the grid frequency and derive the RoCoF [31]. A virtual inertia controller then scales this derivative signal by a predefined inertia constant to generate a proportional active power reference. However, numerical differentiation of frequency measurements inherently amplifies noise, which necessitates substantial low-pass filtering to ensure stable operation. This filtering unavoidably introduces phase delays, thereby compromising the rapid response required for effective inertia emulation in low-inertia systems. To mitigate these limitations, more recent research has increasingly focused on grid-forming control frameworks, most notably virtual SG [32]. By embedding swing equation-based dynamics directly within their inner control loops, virtual SGs regulate their internal frequency and voltage angle autonomously and operate as voltage sources. This control paradigm enables inertial and damping responses to be delivered without explicit RoCoF estimation, thereby reducing sensitivity to measurement noise and filtering delays associated with grid-following inertia emulation [33]–[35].

Although BESS offer superior flexibility and faster response capabilities compared to mechanical prime movers, their practical deployment for inertia emulation is subject to strict hardware and operational constraints. To address the severe electrical stress imposed on power electronic interfaces during rapid power surges, researchers have proposed various mitigation strategies, such as incorporating dynamic current limiting algorithms directly into the virtual inertia control loops to prevent thermal overload [36]. Furthermore, because BESS are inherently energy limited resources, maintaining the SoC during sustained frequency deviations is a critical research focus. Recent literature has explored adaptive virtual inertia frameworks and advanced SoC recovery strategies. For instance, adaptive control schemes can dynamically scale the virtual inertia and damping coefficients based on real-time SoC. This allows the system to inject substantial power immediately following a disturbance, and subsequently reduce the output as the battery discharges, thereby preventing premature energy depletion [37].

Despite these advancements in hardware protection and adaptive control, inertia emulation techniques still exhibit several inherent structural limitations that cannot be fully resolved at the local device level. The performance of many control design would still depend heavily on accurate RoCoF measurements, which are highly sensitive to transient noise and spatial variability across the network. As the control loop illustrated in Fig. 2.1.5, the essential low-pass filtering required to manage this noise inevitably incurs phase delays, which severely compromises the timely activation needed in low-inertia systems. Moreover, inertia emulation predominantly operates as a decentralised, local control action. It relies on fixed mathematical models and does not inherently possess system-wide situational awareness to account for complex operating conditions or coordinate multiple geographically distributed resources. As a result, while local inertia emulation can alleviate the immediate symptoms of a disturbance, it does not on its own provide a comprehensive solution to the broader, uncertain frequency control challenges arising in modern IBR-dominated power systems.

Coordinated Frequency Support from Different Resources

Future power systems are expected to include an increasingly diverse set of resources that can contribute to frequency regulation. Therefore, another key research area has focused on the coordinated control of different types of resources, e.g. distributed generation, energy storage, and controllable loads, to provide effective collective frequency support [38]. Among other arrangement of implementations, such coordinated control is often implemented in microgrids, which aggregates DERs into an overall controllable resource. Within such coordinated frameworks, frequency regulation is realised through rapid and coordinated adjustments of active and reactive power setpoints across participating IBRs, enabling different resources to collectively respond to frequency deviations [39].

In grid-connected operation, collaborative frequency support from a microgrid requires coordinated participation of resources with fundamentally different dynamic characteristics and operational constraints. Energy storage systems can provide rapid and accurate power injections to arrest fast frequency deviations, whereas controllable loads and dispatchable distributed generation typically respond on slower timescales but can sustain adjustments for longer durations [40]. Effective frequency regulation therefore relies on coordination strategies that allocate fast transient support to storage units, while slower resources compensate energy imbalances and restore operating points over extended periods. Such coordinated allocation differs fundamentally from conventional SG-based reserve provision, as distributed resources are often variable, energy constrained, and affected by local operational limits. Consequently, frequency support from microgrids is not achieved by the individual action of single units, but through coordinated control mechanisms that explicitly account for the complementary capabilities of multiple resources [28].

A representative example of coordinated frequency support in microgrids is presented in [41]. In this work, virtual inertia-based control and primary frequency control are jointly coordinated through adaptive participation coefficients, enabling different frequency response mechanisms to contribute in a complementary manner. Rather than relying on independent local control actions, the proposed strategy allocates frequency support responsibilities across multiple resources within the microgrid, thereby improving frequency nadir and dynamic performance under disturbance conditions. This study illustrates how coordinated control of heterogeneous frequency response mechanisms at the microgrid level can enhance collective frequency support while respecting the operational characteristics of IBRs.

Despite their potential to enhance frequency regulation, coordinated frequency support strategies in microgrids face several systemic limitations. First, effective coordination often heavily relies on fast and reliable communication networks between local controllers and supervisory layers. Communication delays, measurement noise, and partial observability of system states can severely degrade the timely activation of frequency responses. Second, continuously optimising the allocation of support responsibilities among energy constrained resources introduces significant computational complexity. Real-time coordinated strategies often struggle to proactively anticipate impending disturbances or predict long-term renewable variability. As a result, while coordinated frequency support from microgrids can significantly improve local flexibility, it does not fully eliminate the need for robust, predictive decision making mechanisms and higher-level system coordination in IBR-dominated power systems.

Advanced Demand Side Response

The demand side is increasingly recognised as a vital flexible contributor to system frequency control. Advanced demand side response enables frequency regulation through continuous and fine grained modulation of power consumption, allowing responsive loads to complement generation based control actions while respecting local operational constraints [42]. By exploiting the inherent flexibility of electrical loads, demand side response provides an alternative approach of delivering fast frequency support in systems with declining synchronous generation and increasing IBRs.

To make effective use of this flexibility, extensive research has investigated control strategies for different types of loads that can actively respond to frequency deviations. A major area of focus is the aggregation of thermostatically controlled loads, including heating, ventilation and air conditioning systems, heat pumps, and domestic refrigeration units. Due to their inherent thermal inertia, the power consumption of these loads can be adjusted over short time horizons without violating operational limits, making them well suited for frequency regulation. Both centralised aggregation approaches [43] and decentralised, frequency based stochastic control strategies [44] have been proposed to enable large populations of thermostatically controlled loads to provide primary and secondary frequency support. In addition, electric vehicles have been widely studied as flexible demand resources, where charging power is adjusted dynamically in response to system frequency deviations [45].

A representative approach to implementing advanced demand side response for frequency control is presented in [44]. In this framework, a local frequency based control logic is embedded within individual appliances, eliminating the need for a central aggregator or extensive communication infrastructure. Each device measures the local grid frequency and adjusts its internal temperature setpoint within a predefined deadband. These incremental setpoint modifications alter the duty cycle and switching probability distribution of a large population of thermostatically controlled loads, resulting in a rapid aggregate active power response analogous to droop-based control. The study demonstrates how autonomous, frequency responsive behaviour at the appliance level can effectively contribute to primary frequency support in power systems.

Although high contribution of demand response in frequency regulation, several operating concerns still exist. The aggregate response of large load populations is influenced by behavioural diversity, heterogeneous device dynamics, and recovery effects following prolonged activation, which complicates the prediction of their collective frequency response. Furthermore, the effectiveness of demand response depends critically on appropriate activation levels and accurate system awareness, as excessive utilisation may induce over frequency while insufficient deployment can significantly reduce frequency control benefits. These challenges are particularly pronounced in multi-area power systems or clusters of microgrids, where uncertainty in disturbance magnitude complicates the coordinated deployment of demand side actions.

Data-Driven Frequency Response Modelling and Control

The efficacy of data-driven frequency control relies on the ability to translate high-resolution synchrophasor measurements into accurate, real-time system representations [46], [47]. By reducing dependence on detailed offline physical parameters, measurement-based modelling captures time-varying dynamics such as shifting system inertia and damping levels, which are critical for predicting indicators like the frequency nadir in IBR-dominated grids [48], [49].

Based on such real-time situational awareness, research has shifted from static, fixed-parameter strategies toward advanced intelligent optimisation. Conventional frequency control schemes, such as fixed-parameter virtual inertia or static droop control, are typically designed for specific operating conditions and lack adaptability to time-varying system dynamics. For example, fuzzy logic controllers have been employed to adaptively tune virtual inertia and droop parameters of distributed BESS based on real-time measurements, resulting in improved frequency stability in microgrid applications [50]. In addition, data-driven UFLS frameworks have been proposed that use real-time frequency measurements to estimate key indicators such as frequency nadir, system inertia, and power imbalance, while explicitly accounting for measurement uncertainties and communication delays. This enables targeted load shedding actions that enhance frequency security while minimising unnecessary load interruptions [51].

An example of this integration is the predictive load frequency control framework developed in [52]. In this work, real-time operational data are used to capture effective frequency dynamics without relying on precise physical models. This data-driven representation enables a predictive control scheme to anticipate frequency deviations and adjust actions proactively rather than reactively. This approach demonstrates how data-driven models can be directly embedded into control loops, achieving adaptive frequency regulation.

Existing data-driven approaches remain limited in their ability to systematically exploit real-time measurement information for system-level decision making. Many studies primarily focus on local parameter estimation, while higher-level coordination and frequency control strategies are still governed by predefined rules or static assumptions.

As a result, the information contained in high-resolution measurement data is not fully utilised to assess system-wide frequency security margins or to support proactive control decisions under extreme uncertainty.

2.2 Frequency Nadir Prediction

2.2.1 Traditional Analytical Approaches

Accurate estimation of the frequency nadir is essential for power system operators to maintain sufficient inertia and reserve capacity. However, the reliability of such estimation is fundamentally rooted in the precise measurement of frequency during transients. By utilising advanced digital signal processing techniques such as self tuning algorithms and improved recursive frameworks, the accuracy and noise immunity of these measurements can be significantly enhanced to provide reliable inputs for prediction models [53], [54]. Traditional methods typically adopt analytical approaches to address this issue, where low-order linear models and aggregation techniques are employed to achieve computationally efficient estimation suitable for planning and short-term analysis [55]. However, as mentioned by [56], such reduced-order representations cannot fully capture fast system dynamics, making pre-event nadir prediction challenging.

To facilitate frequency nadir prediction without relying on full non-linear multi-machine simulations, [57] proposed simplified dynamic equivalent models based on the delay model and the Canonical model. The delay model approximates the initial frequency drop by substituting fast actuator dynamics with an effective time delay, whereas the Canonical model employs shared exponential basis functions to algebraically aggregate multiple steam turbine units. These two methods provide a low-order linear representation that effectively preserves key transient characteristics during power imbalance events. Similarly, [55] developed a System Frequency Response (SFR) model that aggregates reheat-dominated systems into a single equivalent. This formulation retains dominant inertia, reheat, and damping effects to provide explicit analytical formulas for the RoCoF and the nadir. Beyond static estimation, [58] integrated frequency related stability constraints directly into the economic dispatch framework.

By incorporating dispatch decisions with a validated dynamic model and employing sequential quadratic programming for the optimisation, they demonstrated on the Irish system that frequency security (containing the nadir ≥ 48.6 Hz) could be maintained while reducing operating costs.

Although these frequency nadir estimation models are are analytically simple, the aggregate parameters required by them are often only approximately known in advance, which introduces significant estimation uncertainty in pre-fault applications. The aggregate parameters describing inertia, damping, control response, and load behaviour are typically only known with limited accuracy in advance and may fail to capture non-linear phenomena such as control saturation and the threshold-triggered activation of FFR services. In modern power systems, these parameters are highly time varying due to frequent changes in generation mix and weather conditions, such that even small errors in the assumed parameter values can result in large deviations in the predicted frequency nadir. Therefore, while traditional nadir calculations remain valuable for their clear physical interpretation and use in planning studies, their predictive accuracy prior to a disturbance is fundamentally constrained by parameter uncertainty.

2.2.2 Extended Analytical Approaches

Extended analytical approaches have been developed to preserve explicit analytical formulations while providing a more accurate representation of IBR-dominated operating conditions [59], [60]. These methods embed the dynamics of IBRs, FFR controllers and hybrid AC-DC interfaces directly within analytical formulations, allowing stability metrics to be expressed explicitly as functions of controller parameters such as virtual inertia, droop gains and RoCoF-based settings.

Several studies have introduced analytical expressions representations for inverter based FFR. [61] approximated the initial frequency trajectory with a parabolic form and employed polynomial fitting for various governor responses, deriving algebraic equations that significantly reduce computational costs compared to high-order simulations. Building on this, [62] proposed a second-order SFR framework specifically for IBRs.

Beyond single-area models, the application of these analytical tools has been further extended to hybrid architectures. Specifically, [63] modelled voltage source converter based HVDC links within a second-order structure to obtain analytical form expressions for the frequency trajectory. This study identified two distinct frequency nadirs arising from the temporal decoupling between fast converter responses and slower synchronous dynamics. Their sensitivity analysis revealed a critical interaction where increasing control gains to improve the first nadir can inadvertently worsen the second, highlighting the complexity of tuning HVDC parameters for system safety. To address the scale of such hybrid grids, [64] employed parameter identification and the inverse Laplace transform to obtain time-domain equations for systems with ultra-HVDC links. This approach enables the explicit calculation of four key performance indices including RoCoF, the frequency nadir, the nadir time, and the recovery rate. Similarly, [65] extended this security assessment to the Jeju power system by developing a model that captures the dynamic interaction between HVDC control mechanisms and system frequency. These developments demonstrate that analytical models can match the accuracy of detailed non-linear simulations while providing rapid estimation for planning and sizing resources in networks with a high penetration of IBRs.

To address probabilistic reserve planning, [66] developed a full-response analytical model for FFR studies. This approach introduces an aggregation scheme to represent multiple resources within a single tractable framework that produces complete post-event frequency trajectories. The model analytically computes key metrics across various disturbance types, and its probabilistic formulations generate distributions of frequency indices rather than single point estimates, facilitating reserve sizing under uncertainty without relying on Monte Carlo simulations.

In summary, compared with time-domain simulations or numerical integration of high-order dynamic models, analytical formulations enable much faster evaluation of frequency metrics such as the frequency nadir. However, as analytical models are extended to represent increasingly complex system structures, including IBR-dominated networks, FFR control, and hybrid systems, their predictive reliability becomes progressively more difficult to maintain. Such models typically rely on linearised representations and assume that aggregated physical parameters can be obtained with sufficient accuracy pre-event. In practice, achieving close agreement between assumed model parameters and the actual system conditions at the time of a fault is challenging. Variations in operating conditions and control behaviour introduce uncertainties that are difficult to capture accurately using fixed analytical representations. In addition, the use of linear approximations restricts the ability of analytical models to represent abrupt and non-smooth control. As a result, the applicability of purely analytical approaches is increasingly constrained in modern power systems. These limitations motivate the adoption of data-driven methodologies, which infer frequency security directly from operational measurements and reduce reliance on predefined physical parameters.

2.2.3 Data-Driven Approaches

To address the limitations of analytical approaches, data-driven approaches learn mappings from pre-fault operating features to the post-fault frequency nadir. These approaches learn the complex relationship between system conditions and frequency paths directly from data. This reduces the need for exact physical models and makes model setup easier. Algorithms such as RF and gradient boosting have shown a strong ability to work well with new data and handle measurement noise [67]. Their ability to capture complex non-linear interactions makes them well suited for power systems with varying inertia levels and diverse disturbance scenarios.

Existing studies typically train offline models using historical data or event records. They are then used to predict the nadir across many expected system conditions, which support frequency control ahead of time. Online adaptive schemes update model parameters with continuous measurements to provide real-time predictions that reflect current operating conditions. Typical input data include total system load, available reserves, estimated inertia, the size and location of the disturbance, and regional renewable generation [68]. These features describe the frequency response potential of the system just before a disturbance happens. Several studies have shown that using only the information available before a disturbance can produce accurate nadir predictions. This allows operators to take preventive actions before a fault even occurs. For real-time applications, [69] introduces nadir prediction as a supervised learning task using synchrophasor data, with generator bus voltage phase angles as inputs and the Center of Inertia (CoI) nadir as the target. To generate training data, the study simulates low-inertia conditions on a benchmark system. Several algorithms are compared, with hyperparameters tuned by Bayesian optimisation. Among these, decision trees were identified as the most effective model.

Neural network models have also been explored, where[70] proposes a feedforward neural network for online prediction of the frequency nadir after severe power imbalances. The network is trained offline using time series generated from dynamic simulations. These simulations cover a wide range of operating conditions, disturbance size and load voltage characteristics. By using the initial frequency and voltage paths instead of simple aggregate models, the neural network captures the complex interactions between inertia, load effects, and frequency control. Results show that high prediction accuracy is maintained even when inertia is only roughly estimated.

Data-driven approaches for frequency nadir prediction still face several important challenges. First, prediction accuracy can degrade when real world operating conditions differ significantly from the data used for training, particularly when the system encounters previously unseen scenarios. This limits the robustness of purely data-driven models under rapidly changing system conditions. Second, obtaining the accurate real-time parameters required for these models is often unfeasible. Since many input features are based on uncertain estimates rather than direct measurements, these inaccuracies are reflected in the final prediction.

2.3 Estimation of RoCoF and Power Imbalance Size

2.3.1 Role of RoCoF and Power Imbalance Estimation in Frequency Control

In low-inertia power systems, maintaining frequency security requires the rapid and accurate assessment of disturbance severity within a short time window, which is typically a few hundred milliseconds following the event. Real-time assessment relies on two closely related physical quantities, i.e. RoCoF and the corresponding active power imbalance. RoCoF represents the rate of frequency deviation and provides a fast indication of dynamic stress following a power imbalance disturbance [71]. Their relationship is governed by the classical aggregated swing equation, which states that the initial RoCoF is approximately proportional to the net active power imbalance divided by the equivalent system inertia [56]. Therefore, for a given inertia, larger generation losses produce higher RoCoF values, whereas systems with lower inertia will also lead to increased RoCoF. This makes the monitoring of these variables vital for situational awareness.

Accurate and fast estimation of both RoCoF and power imbalance is essential for the effective operation of emerging frequency control framework [72]. These two quantities are necessary for different aspects of control: RoCoF provides the trigger for corrective actions, while the imbalance size informs the required response volume.

- **RoCoF for detection and triggering:** RoCoF estimates are primarily used for fast disturbance detection and classification. They act as the activation trigger for FFR services and emergency control schemes, such as UFLS.
- **Power imbalance size for quantification:** While RoCoF indicates the speed of the frequency decline, power imbalance size estimation quantifies the magnitude of the required mitigating action [73]. This enables emergency frequency control schemes to determine the precise volume of load to be shed, preventing both under shedding (which risks system instability) and over shedding (which causes unnecessary service interruptions).

2.3.2 Theoretical Foundation: Swing Equation

The swing equation can be used to quantify the relationship between frequency dynamics and power system disturbances, as it captures how the stored kinetic energy of synchronous machines responds to active power imbalances. In conventional power systems, SGs transform mechanical power supplied by prime movers into electrical power. These machines comprise a fixed stator and a rotating rotor. The rotor produces a magnetic field, and its rotational speed is coupled with the power system frequency according to the fundamental relationship [74]

$$n = \frac{120f}{p} \tag{2.3.1}$$

where n is the rotor speed (rpm), f is the electrical frequency (Hz), and p denotes the number of poles. As p is fixed for a given machine, the rotor speed is proportional to system frequency. Due to its mass, the rotor provides rotational inertia that resists speed changes, thereby mitigating fast variations in system frequency [75]. In an interconnected grid, the aggregated inertia of online synchronous machines forms a kinetic energy reserve that supports frequency stability and helps limit severe frequency deviations during disturbances [76].

At the machine level, the mass moment of inertia of the rotor is given by

$$J = \int r^2 \, dm = r^2 m \tag{2.3.2}$$

where m denotes the rotor mass (kg) and r denotes an effective radius (m). Consequently, the kinetic energy (KE) stored in the rotor when it rotates at mechanical angular speed ω_m (rad/s) is expressed as

$$KE = \frac{1}{2} J \omega_m^2 \tag{2.3.3}$$

which shows that the stored kinetic energy increases with both the rotor inertia and the square of its rotational speed.

Mechanical angular speed ω_m is related to the rotor speed n by

$$\omega_m = \frac{2\pi n}{60} \tag{2.3.4}$$

Based on these definitions, the inertia constant H of an SG is defined as the kinetic energy stored in the rotor at rated speed ω_{sm}, normalised by the apparent power rating S. This constant quantifies the theoretical-time the generator could deliver its rated power using the rotor's stored energy

$$H = \frac{KE_{rated}}{S} = \frac{\frac{1}{2} J \omega_{sm}^2}{S} \tag{2.3.5}$$

For a system with n online SGs, the aggregated inertia constant on the total base $\sum S_i$ is given by

$$H_{\text{total}} = \frac{\sum_{i=1}^{n} S_i H_i}{\sum_{i=1}^{n} S_i} \tag{2.3.6}$$

where S_i and H_i are the apparent power rating and inertia constant of the i^{th} machine, respectively. The quantity H_{total} characterises the initial inertial response of the interconnected system.

When a disturbance occurs (e.g. a generator trip or a sudden load change), an imbalance arises between the mechanical torque T_m and the electromagnetic torque T_e. For an area represented by an equivalent machine, Newton's second law for rotational motion governs the dynamics

$$T_m - T_e = J_{eq} \frac{d^2\theta}{dt^2} \tag{2.3.7}$$

where J_{eq} is the equivalent rotational inertia of the area and θ is the mechanical rotor angle measured with respect to a stationary reference.

Using the relation $P = T\omega$ and assuming that the rotor speed remains close to the synchronous speed ω_{sm}, the mechanical and electrical torques are approximated as

$$T_m = \frac{P_m}{\omega_{sm}} \qquad T_e = \frac{P_e}{\omega_{sm}} \tag{2.3.8}$$

with mechanical input power P_m and electrical output power P_e. Substituting these in to (2.3.7)

$$\frac{P_m - P_e}{\omega_{sm}} = J_{eq} \frac{d^2\theta}{dt^2} \tag{2.3.9}$$

To formulate the system-level swing equation, the equivalent moment of inertia J_{eq} is expressed via the aggregated inertia constant H_{total} relative to the system power base S_{base}. Here, S_{base} is defined as the total rated capacity of the system, given by $S_{base} = \sum S_i$, which serves as the common MVA base for all per-unit calculations. The relationship is defined as

$$J_{eq} = \frac{2H_{total} S_{base}}{\omega_{sm}^2} \tag{2.3.10}$$

Let $P_{m,pu} = P_m/S_{base}$ and $P_{e,pu} = P_e/S_{base}$ denote the per-unit power on the system base. Dividing (2.3.9) by S_{base} and inserting J_{eq} from (2.3.10) gives the classical per-unit swing equation in terms of the absolute rotor angle θ

$$\frac{2H_{total}}{\omega_{sm}}\frac{d^2\theta}{dt^2} = P_{m,pu} - P_{e,pu} \tag{2.3.11}$$

where H_{total} is the system-wide aggregated inertia constant expressed on the S_{base} power base.

For stability analysis, the absolute rotor angle $\theta(t)$ is expressed as the sum of the synchronous position $\omega_{sm}t$ and the rotor angle deviation $\delta_m(t)$. Since the synchronous speed ω_{sm} is constant, the acceleration terms are identical

$$\frac{d^2\theta}{dt^2} = \frac{d^2}{dt^2}(\omega_{sm}t + \delta_m) = \frac{d^2\delta_m}{dt^2} = \frac{d\Delta\omega}{dt} \tag{2.3.12}$$

where $\Delta\omega$ is the deviation of the rotor speed. Replacing θ in (2.3.11) with δ_m leads to the swing equation in terms of speed deviation

$$\frac{2H_{total}}{\omega_{sm}}\frac{d\Delta\omega}{dt} = P_{m,pu} - P_{e,pu} = \Delta P_{pu} \tag{2.3.13}$$

Relating the mechanical speed to the electrical frequency through $\omega_{sm} = 2\pi f_n$ and $\Delta\omega = 2\pi\Delta f$, the equation can be expressed in terms of the system frequency f

$$\frac{2H_{total}}{2\pi f_n}\cdot\left(2\pi\frac{df}{dt}\right) = \frac{2H_{total}}{f_n}\frac{df}{dt} = \Delta P_{pu} \tag{2.3.14}$$

Eq. (2.3.14) defines the relationship between RoCoF and power imbalance during the immediate post-event window, assuming inertia-dominated dynamics prior to governor intervention.

Consequently, if the maximum RoCoF at the occurrence of a disturbance is denoted by $RoCoF_{max}$, the corresponding initial power imbalance can be estimated as

$$\Delta P_{pu} = \frac{2H_{total}}{f_n}\,RoCoF_{max} \tag{2.3.15}$$

2.3.3 RoCoF Estimation

Fundamental Frequency Extraction

The most fundamental way to RoCoF estimation is to firstly obtain an instantaneous frequency from local voltage or current waveforms and subsequently calculate its time derivative [77]. Common implementations include Zero-Crossing Detection (ZCD), PLL algorithms, and synchrophasor-based methods in PMUs, as well as dedicated functions in conventional protection relays.

ZCD is one of the most common and computationally efficient methods for measuring the frequency of periodic waveforms [78]. The method determines the time interval between consecutive zero crossings or estimates the waveform period [79]. To reduce sensitivity to phase noise and waveform distortions in applied settings, the algorithm typically counts the number of reference clock cycles over multiple fundamental periods rather than a single cycle. While extending the measurement window effectively improves steady-state accuracy by averaging out local variations, it unavoidably slows down the update rate of the frequency estimate. This presents a significant challenge as frequency measurements for synchronisation require both high accuracy and fast response simultaneously [80], [81]. For RoCoF estimation that differentiates discrete frequency samples, such extended averaging obscures rapid post disturbance dynamics. Consequently, ZCD is generally unsuitable for accurate RoCoF estimation when millisecond level resolution is required.

PLL-based techniques are widely used in grid-interfaced converters and measurement equipment to track the fundamental component of the voltage waveform [82]. In a synchronous reference frame, the measured voltage is mapped into rotating axes [83], and a feedback controller adjusts an internal oscillator to align the estimated phase with the grid voltage. Consequently, the instantaneous frequency is derived from the time derivative of this estimated phase. Compared with ZCD, the PLL provides a continuous and smoothed estimate of the local system frequency. For RoCoF estimation, this continuous frequency output is processed through numerical differentiation, typically combined with filtering to attenuate noise. With appropriately calibrated control parameters, a PLL can balance noise suppression with dynamic response. However, performance remains sensitive to parameter settings. Severe voltage sags, unbalance, and harmonic distortion may bias the frequency estimate, while bandwidth limitations can introduce tracking delay during rapid deviations, resulting in inaccurate RoCoF during the critical interval immediately after a disturbance. Let $\hat{f}(k)$ be the discrete-time frequency at sample k with a sampling interval T_s. The simplest approach applies a backward difference approximation [84]

$$\mathrm{RoCoF}(k) = \frac{\hat{f}(k) - \hat{f}(k-1)}{T_s} \tag{2.3.16}$$

Improved accuracy can be achieved by the central difference

$$\mathrm{RoCoF}(k) = \frac{\hat{f}(k+1) - \hat{f}(k-1)}{2T_s} \tag{2.3.17}$$

These schemes are simple to implement and computationally light, so they are commonly adopted in protection relays and measurement devices. However, direct differencing is highly sensitive to measurement noise and quantisation errors in $\hat{f}(k)$ and can severely amplify wideband disturbances, rendering RoCoF estimates unreliable under harmonics, interharmonics and waveform distortion.

As a more robust alternative, Least Squares (LS) fitting replaces point to point differencing with local linear regression over a sliding window. The implementation begins by using a frequency estimator to generate a high rate sequence of instantaneous frequency samples [85], applies median filtering to suppress spikes, and then fits a straight line to the most recent samples, where the fitted slope is reported as RoCoF. The window length is optimised to balance noise suppression against the response time. Because LS-based estimators process multiple samples within the window, they are significantly less sensitive to measurement noise while providing smoother and more stable RoCoF estimates.

Synchrophasor-based algorithms in PMUs provide relatively more advanced frequency estimates by computing complex phasors over sliding windows and reporting frequency and, in some cases, RoCoF as derived quantities. Standardised procedures and performance classes promote consistency across devices, with IEEE C37.118 specifying, for example, latency of $\leq 100\,\mathrm{ms}$ and a steady state RoCoF error of $\leq 0.01\,\mathrm{Hz/s}$ [86], [87]. Despite these requirements, accurate RoCoF reporting remains challenging because differentiation of the measured phase substantially amplifies noise [88]. This is because it involves the second derivative of the measured phase, so even small amounts of noise in the phase signal can lead to substantial amplification in the calculated RoCoF [89]. Although PMUs are designed to report RoCoF with high speed and low latency, their outputs remain sensitive to this measurement noise and usually require substantial filtering to achieve the desired accuracy [90], [91]. Furthermore, most PMU algorithms assume a narrow band sinusoidal model that fails to capture the broadband spectrum observed during transients or sudden generation and load imbalances. As a result, they still face conflicting design requirements regarding window length, spectral leakage suppression, and dynamic tracking capability. Windowing and filtering introduce latency that may smooth or distort high RoCoF values, while oscillatory and non stationary conditions can produce significant estimation errors. Consequently, different implementations employ different window lengths, filters, and RoCoF definitions, leading to inconsistent RoCoF values for the same event and complicating the coordination of protection and control schemes [88].

To address these limitations, advanced methods based on the Discrete Fourier Transform (DFT) aim to solve accuracy problems caused by short observation windows. For example, an enhanced interpolated DFT (e-IpDFT) algorithm is presented in [92]. Standard DFT results are limited to fixed frequency steps, meaning the real frequency in power system often lies between these steps. To address this, the e-IpDFT method applies a specific window function (e.g. Hann) and uses a correction formula to precisely calculate the fundamental phasor [93]. Mathematically, the negative part of the spectrum can overlap with the positive part when the window is short, leading to measurement errors. By removing this specific interference, the algorithm significantly reduces errors while allowing the use of short windows for fast reporting. Building on similar principles, an iterative interpolated DFT approach [94] operating on short windows (typically three cycles) mitigates interference from the negative frequency image and nearby interharmonics, thereby enabling phasor, frequency and RoCoF estimation with high accuracy under distortion.

To better capture rapidly changing signals, an approach using the Taylor-Fourier multi-frequency model [95] treats the signal as a superposition of dynamic spectral components. Rather than assuming that the signal frequency remains constant within the observation window, the Taylor-Fourier multi-frequency model explicitly captures time-varying spectral components by jointly estimating the fundamental frequency, harmonics, and their time derivatives. Incorporating derivatives directly into the model enhances the separation of true dynamics from measurement noise and provides stable estimates under oscillatory and non stationary waveforms.

In summary, while techniques ranging from simple numerical differentiation to advanced spectral interpolation and dynamic signal modelling have significantly improved RoCoF estimation, their performance remains constrained by the balance between noise suppression and response speed. The extensive windowing required to ensure numerical stability and accuracy inevitably introduce reporting latency, leading to inconsistent RoCoF values that can complicate the coordination of fast protection and control schemes.

RoCoF Estimation Using Filtering Techniques

Filtering-based methods for RoCoF estimation advance beyond window-based and LS approaches by using historical information. Instead of treating each observation window as an independent event, these methods model the phasor, frequency, and RoCoF as continuous states within a dynamic system. Typically, a Kalman Filter (KF) provides an optimal recursive framework that predicts the next state based on the previous estimate and subsequently corrects it using incoming measurements. This recursive structure ensures that each new estimate is physically consistent with the system's previous trajectory, effectively mitigating the impact of transient outliers and quantisation noise.

A representative two-stage approach integrating an IpDFT with a linear KF is reported in [96]. In the first stage, a short window IpDFT acts as the preliminary process, providing low-latency, but potentially noisy frequency estimates. These estimates then act as the input for a dynamic model, where frequency and RoCoF are linked by a simple physical relationship. The KF recursively combines the model's prediction with the IpDFT measurements to produce a smoothed RoCoF estimate. Compared with direct numerical differencing, this cascaded filtering structure significantly suppresses noise amplification while preserving the ability to detect rapid frequency events.

To address more complex dynamics and waveform distortion, extended KF estimators have been proposed in [97]. Unlike the two-stage methods as discussed above, these approaches employ a non-linear dynamic phasor model where amplitude, phase, frequency, and RoCoF are all defined as internal state variables. The algorithm mathematically relates the raw voltage waveform directly to these states, allowing the extended KF to perform prediction and correction at every sampling instant. This enables the simultaneous filtering and tracking of signal variations while removing harmonics and interference. Expanding on this framework, Taylor-based extended KF schemes include higher-order terms to maintain accuracy under unbalanced or multi-frequency conditions [98]. In these frameworks, RoCoF is obtained directly as a filtered state output without secondary differentiation.

In summary, filtering-based methods provide an effective intermediate solution between simple window-based estimators and fully model-based observers. These methods produce smoother and more reliable RoCoF values while maintaining relatively low latency. However, their performance strongly depends on how well the filter settings match actual system behaviour, which can vary significantly under different operating conditions. In addition, because filtering inherently smooths estimated trajectories based on assumed system behaviour, abrupt and irregular frequency changes following severe disturbances may be partially attenuated or delayed. As a result, tracking accuracy can degrade during rapid, non-uniform frequency dynamics.

RoCoF Estimation based on Data-Driven Methods

With the rapid development of data-driven methods, power system studies increasingly use historical records to train algorithms for RoCoF-related estimation tasks. These approaches aim to learn complex mappings between system operating conditions and frequency dynamics, enabling fast inference and supporting real-time decision making [99]. Compared with conventional techniques, data-driven models can capture non-linear relationships and interactions that are difficult to represent explicitly.

Several studies have explored data-driven regression frameworks that treat RoCoF as a primary prediction target to bypass the latencies of traditional measurements. In [100], a framework is proposed using pre-event operating features including total system load, spinning reserve, aggregated kinetic energy, and estimated inertia, which allows immediate RoCoF prediction without requiring post-event measurements. Similarly, a recent study in [101] utilises ensemble learning models, including RF and XGBoost, to map pre-disturbance conditions to the expected maximum RoCoF. While both studies demonstrate that algorithms such as XGBoost achieve high accuracy and consistency compared with traditional linear models, [101] specifically highlights the resilience of this approach under high renewable penetration. By shifting from reactive differentiation to proactive mapping, these data-driven models offer a potential method for RoCoF-based protection in low-inertia systems.

Existing data-driven approaches to RoCoF prediction still present several important limitations. First, their predictive performance is highly dependent on the quality and representativeness of the training dataset, which may limit their ability to generalise to previously unseen disturbances. Second, models designed to achieve high prediction accuracy often rely on rich input features or more elaborate model structures, which may reduce their suitability for use during the critical early post-disturbance period when information is limited and rapid response is required. Third, most approaches depend on detailed system operating conditions, such as inertia and reserve levels, which are difficult to estimate accurately in real-time and are themselves subject to uncertainty. These limitations highlight the need for estimation frameworks that can provide fast and reliable disturbance characterisation while reducing reliance on uncertain system parameters.

2.3.4 Power Imbalance Estimation

As mentioned previously, many modern monitoring and control functions require accurate estimation of power imbalance sizes. Presently, one of the most widely used methods is the model-based approach, which derives power imbalance based on the swing equation using the measured RoCoF and system inertia. This practice presumes that system inertia is accurately known at the time of the event. However, online inertia estimates are often uncertain in low-inertia systems. This difference between the assumed and actual system inertia can introduce an estimation error in the inferred power imbalance.

While various existing studies focus on the estimation of power imbalance, they often assume a constant system inertia, neglecting the instantaneous loss of inertia when a generator trips. To bridge this gap, [102] introduces an approach that explicitly accounts for post-event inertia by subtracting the lost generator's contribution from the total pre-event inertia. This method employs a reduction coefficient identified online using smoothed RoCoF measurements from PMUs, enabling the simultaneous estimation of lost generation and inertia reduction. However, this approach remains sensitive to the accuracy of pre-event inertia data and requires accurate RoCoF measurements. In practical low-inertia systems, where RoCoF signals are often contaminated by significant noise and transient, may limit the estimator's reliability.

Despite these methodological advancements, the accuracy of existing power imbalance estimation is still dominated by the quality of its two key parameters, i.e. RoCoF and system inertia. As indicated in (2.3.15), the maximum RoCoF is a critical variable for accurate estimation of the power imbalance. As also discussed in the previous section on RoCoF estimation, RoCoF measurements are highly sensitive to noises and transients during disturbances [103]. To mitigate these measurement errors, low-pass filtering is typically applied. However, this process unavoidably distorts the signal with measurement error and latency. Fig. 2.3.1 presents an example, which illustrate that the filter can introduce not only a response delay but also a significant reduction in peak magnitude. Consequently, the filtered output (blue dashed line) fails to capture the true peak of the actual RoCoF (green solid line), directly leading to an underestimation of the power imbalance. Therefore, selecting appropriate filter parameters involves a difficult balance between smoothing measurement noise and maintaining accuracy.

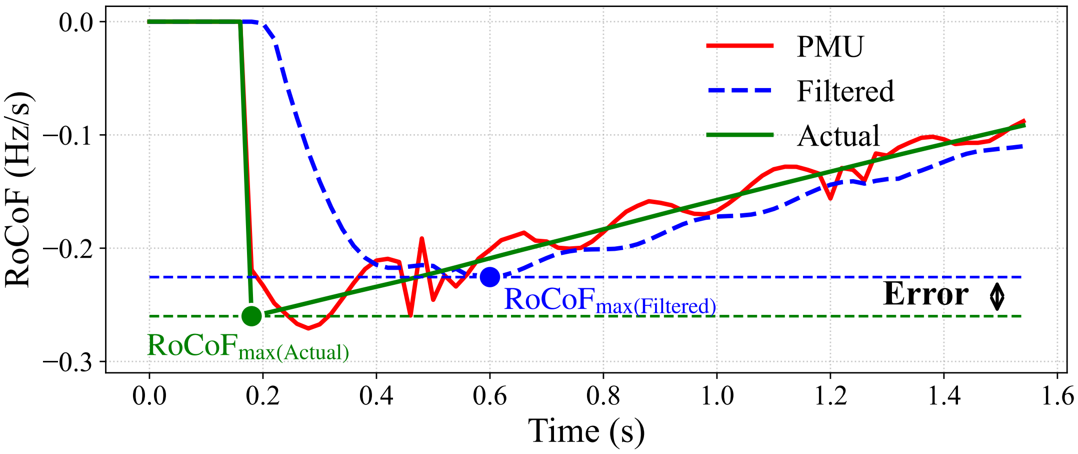

Fig. 2.3.1: Comparison between actual, filtered and PMU-measured RoCoF

The second key parameter required in Eq. (2.3.15) is the system inertia H. Obtaining an accurate value of system inertia in real-time remains a significant challenge, especially in large power systems with high penetration of variable RES. These resources may also employ diverse control strategies, e.g. inertia emulation, droop, etc., their dynamic responses differ significantly from those of synchronous machines [104]. Furthermore, the inherent intermittency of RES introduces uncertainty into real-time inertia estimation, which ultimately degrades the accuracy of the power imbalance size estimation [105].

In summary, existing literature demonstrates that accounting for post-event inertia, using optimisation-based parameter inference, and compensating for emulated inertia can significantly improve estimation accuracy. However, there still have several limitations. First, these methods rely heavily on high-quality early-stage RoCoF measurements and precise time synchronisation, which are difficult to guarantee in real world scenarios. In addition, the estimation process remains highly sensitive to signal filtering, as the smoothing required to stabilise RoCoF measurements inevitably introduces delays and attenuates peak values, leading to underestimation of the power imbalance. These limitations motivate approaches that reduce dependence on early-stage RoCoF accuracy while meeting the low-latency requirements of real-time control.

2.4 Emergency Frequency Control Schemes

Ensuring the reliable operation of modern power systems increasingly depends on effective defence measures that can respond across multiple timescales. The growing volatility introduced by IBRs indicates means that system frequency may change more rapidly following major disturbances, which necessitates a hierarchical emergency frequency control framework. Accordingly, emergency frequency control actions are typically structured according to both the severity of the disturbance and the required speed of response. In GB, the first line of defence consists of fast acting reactive frequency response services such as DC and FFR. These are complemented by limited frequency sensitive mode and traditional primary frequency response from SGs, which provide additional stabilising support through power modulation.

As system inertia continues to decline, the RoCoF resulting from significant power imbalances can reach levels that exceed the operational limits of compensatory measures. Under these conditions, the rapid frequency decline may outpace the response capabilities of primary control and FFR, rendering them insufficient to maintain system stability independently. In such extreme scenarios, UFLS serves as the essential final layer within the defence hierarchy. While earlier measures act by adjusting generation to restore balance, UFLS preserves system integrity by mandating the immediate disconnection of demand once the frequency violates critical operational limits such as the 48.8 Hz threshold in the GB [106]. As the last safeguard prevents system widespread outages, the speed and coordination of UFLS are of central importance. In contrast with emergency actions such as intentional islanding or generator power reduction, which involve complex coordination and operational risks, UFLS provides a decentralised and rapid stabilising effect that is well suited to fast frequency events. As a result, recent research has increasingly concentrated on improving UFLS strategies to address the challenges of IBRs-dominated grids. Existing methods are generally categorised into fixed-threshold, semi-adaptive, adaptive schemes [107], and emerging approaches that incorporates predictive elements to further enhance overall system stability [108].

2.4.1 Conventional UFLS with Fixed Thresholds

Traditional UFLS remains the most widely deployed approach in power systems. In such schemes, relays operate locally, relying on the frequency measured at their installation point. The measurement is continuously compared against a set of predefined under-frequency thresholds, known as shedding stages. When the frequency crosses a pre-defined threshold, the relay trips a prescribed portion of the load for that stage after an intentional-time delay. At the system level, the scheme specifies the total amount of load to be shed per stage, while substations allocate this amount according to a load priority list. Because tripping is implemented in discrete blocks and coordinated across multiple substations, the exact target shedding amount can only be approximated during actual operation. The effectiveness of this approach depends on the proper setting of these thresholds and delays, which are typically determined by the transmission system operator.

The number of stages and their setpoints are defined by the transmission system operator and vary widely. For example, the Slovenian electric power system employs four stages [109], whereas the Latvian system uses as many as twenty [110]. Nevertheless, most studies report that three to six stages are typical [111]. At the interconnection level, for example within ENTSO-E, guidance documents are typically available to assist transmission system operators in parameterising traditional UFLS schemes [112]. Such parameterisation defines the number of stages, the corresponding frequency thresholds, the amount of load to be disconnected at each stage, and the intentional-time delays applied prior to tripping.

The increased penetration of IBRs and decreased system inertia have made the limitations of fixed-threshold UFLS schemes increasingly evident. Among others, one of the key issue is RoCoF is not typically used for decision making. In reality, a fast frequency drop usually signals a large power imbalance, while a slow drop typically means a small power imbalance. However, traditional schemes only look at the absolute frequency value and treat these two situations exactly the same. This often leads to inadequate or even incorrect actions. For example, if the frequency declines gradually, the scheme may shed an excessive amount of load, causing a sharp rebound in frequency.

As a result, generators may trip due to over-speed protection and cause a secondary blackout. Conversely, if the frequency declines very rapidly, the fixed time delay may prove excessively long, resulting in slower actions that fail to effectively contain the frequency decline. Consequently, the frequency may continues to decrease rapidly, leading to a blackout.

These risks are clearly illustrated by a real event in the Slovenian power system on 27 April 2007 [108]. Following a disturbance, the frequency dropped below the thresholds of a four-stage traditional plan, which triggered multiple stages of load shedding. Because the scheme was fixed and could not adapt, the total amount of disconnected load did not match the actual power imbalance. Consequently, the frequency rose sharply above 52.5 Hz, triggering the over-frequency protection on several generators. This loss of generation led to a second frequency collapse and eventually caused a total system blackout.

Even when traditional UFLS succeeds in saving the system, the result is rarely efficient. Since the logic ignores the RoCoF and the size of the power imbalance, the amount of shed load rarely matches what the system actually needs. As a result, the system recovery often involves large frequency swings and unnecessary power loss. These problems highlight the clear need for advanced methods, i.e. more adaptive or predictive approaches are required to better align the emerging faster system dynamics due to the IBR integration.

2.4.2 Semi-Adaptive UFLS with Fixed Thresholds

To overcome the limitations of conventional schemes, semi-adaptive UFLS represents a viable alternative. As characterised by [109], [113], these methodologies represent a compromise between traditional fixed logic and fully adaptive approaches. They preserve the standard multi-stage frequency thresholds, but integrate dynamic adjustments based on the local RoCoF. The fundamental concept is to estimate the severity of the power imbalance in real-time, where the detection of a high RoCoF indicates a severe event. Consequently, the relay accelerates the shedding process by reducing the time delay or increasing the amount of shed load. Conversely, under low RoCoF conditions, the relay follows the nominal settings to avoid over-shedding. This allows the shedding actions to align the speed of their responses with the actual needs of the system. Therefore, semi-adaptive UFLS bridges the gap between the static schemes prescribed in many countries' standards (e.g. [114]), which rely on fixed settings, and the fully adaptive methods proposed in the literature (e.g. [115]), which determine the exact load shedding amount online.

A primary advantage of this strategy is its operational independence. As detailed in [116], the logic avoids complex online parameter estimation and needs no external communication channels. This design choice reduces the risk of failure since the device does not depend on external signals and simplifies validation, as each unit can be tested independently. However, because decisions are based solely on local measurements, the scheme has limited capability to distinguish between small and large power imbalances. To overcome this, industry practice relies on three main techniques that introduce incremental adaptability. The first method is RoCoF-based classification. In this approach, the relay measures the frequency slope and assigns it to a discrete severity level. Based on this level, it selects the shedding amount and time delay from a predefined list. The second technique is first-stage reinforcement. This logic focuses on providing a strong initial response. When the relay detects a severe fault, it temporarily increases the load shed at the first stage. However, the settings for subsequent stages remain unchanged to prevent excessive disconnection. The third method involves changing the shedding sequence. During a rapid frequency decline, the logic moves a larger group of load to the first stage to provide immediate support. Consequently, the smaller groups are shifted to later stages.

Early studies, e.g. [117], also investigated the use of the initial frequency slope to adjust the speed of the first stage. If the slope is steep, the response becomes faster and stronger. This method improves performance without needing complex system models. More recent work has refined this approach. Specifically, [118] uses offline simulations to analyse many possible outage scenarios. Engineers can then optimise the threshold settings and decision rules beforehand. During operation, the relay simply follows these optimised rules based on the local signal.

Overall, semi-adaptive UFLS improves the response speed and accuracy of the protection compared to conventional fixed-stage schemes. It achieves this while maintaining the practical benefit of local and independent operation. However, these schemes still have limitations. They rely on a single measurement point at the installation location. In complex grids with low-inertia, the local RoCoF might not perfectly reflect the system-wide disturbance severity. If the local signal is distorted by oscillations or noise, the relay may still take incorrect actions. Therefore, the risk of system power imbalance or instability remains a concern in challenging operating conditions.

2.4.3 Adaptive UFLS Based on Power Imbalance Estimation

Adaptive UFLS schemes based on power imbalance size and CoI frequency represent a significant evolution from semi-adaptive concepts. Rather than relying on local measurements and fixed stages, these methods utilise information reflecting the system-wide response. They estimate the magnitude of the active power imbalance and track the evolution of the overall system frequency to determine the exact amount of load to be disconnected. The core principle is that global frequency behaviour, such as CoI frequency and its rate of change, provide a more reliable measure of severity than a single local measurement. Unlike semi-adaptive schemes that merely adjust fixed-thresholds, this class of methods calculates the shedding quantity online as a function of the estimated system-level imbalance. Consequently, the decision relies on the global system response rather than isolated local logic. This leads to a closer alignment between the generation loss and the amount of shed load.

Within this class of schemes, the control logic relies on two fundamental components: an online estimation of the total disturbance magnitude and the CoI frequency [119]. This signal represent as the primary input for the UFLS decision process. Following a significant disturbance, the controller analyses the frequency of CoI and its rate of change. By using the swing equation, this analysis estimates the active power imbalance across the entire grid. Based on this estimate, the scheme computes the exact amount of load shedding required. The objective is to keep the frequency nadir above a safety limit and restore power balance within an acceptable time. Two implementation patterns are commonly observed. Single-shot adaptive schemes calculate a single shedding action from the initial estimate and execute it immediately to stop the frequency decline. Alternatively, iterative adaptive schemes perform an initial shedding step followed by further corrections. Each subsequent action is updated according to the ongoing trajectory of the CoI frequency. In all cases, the decision depends on the estimated response of the total system rather than a fixed mapping from local thresholds to predefined shedding amounts.

In [120], it proposes an adaptive UFLS approach that can be implemented locally without relying on communication infrastructure. The scheme uses a locally estimated approximation of the CoI RoCoF to infer the magnitude of the generation loss and determine the required shedding amount. To achieve accuracy, an inflection point detector is applied to local frequency measurements. This filter removes inter-area oscillations, using a precise approximation of the global CoI RoCoF. With periodically updated system inertia, each relay converts this global estimate into a specific generation loss. It then sheds a share of load proportional to its assigned responsibility, ensuring the total shed matches the estimated event size. Load shedding is executed in a single adaptive step once the threshold is violated. This replaces fixed multi-stage settings based on assumed worst-case conditions with a dynamic response.

The optimisation of these adaptive steps is further refined in [121], which proposes distributing load shedding over multiple stages based on the initial frequency deviation. By applying the swing equation, the controller updates the shedding amount in real-time according to the evolving frequency response. This approach allows primary frequency control to be utilised as much as possible before additional load is disconnected, thereby reducing the risk of excessive load shedding.

Focusing on real-time implementation, [51] introduces a predictive UFLS scheme that uses analytical estimation. Instead of waiting for frequency to cross fixed limits, this logic uses PMU measurements to predict the future frequency path. By estimating the power imbalance and system inertia from a small set of measurements, the controller decides if the frequency will drop below safety levels. This allows for an early response, stopping the frequency drop more effectively than traditional methods that only react after the drop happens.

While adaptive UFLS schemes based on estimated disturbance magnitude and CoI frequency offer the potential for targeted shedding, several implementation challenges remain. First, these approaches depend on wide-area measurements and communication infrastructure. During severe disturbances, when protection reliability is critical, these channels may suffer from latency. Such disruptions can reduce the accuracy of CoI estimation, leading to delayed or incorrect control actions. Second, performance relies on accurate representations of aggregate inertia and system parameters. However, in grids with high penetration of IBRs, these quantities fluctuate with operating conditions and are difficult to estimate precisely. Consequently, modelling errors can directly result in excessive shedding or an insufficient response. Third, many schemes are validated using simplified system equivalents or a limited set of scenarios. This creates uncertainty about their resilience under structural changes or extreme events.

2.4.4 Data-Driven Approaches for Predictive UFLS

Predictive UFLS schemes aim to solve the main problems of semi-adaptive and adaptive methods. Semi-adaptive schemes lack the precision to match the actual severity of different events because they rely on fixed stages. While physics-based adaptive schemes improve this by using system-wide data, they depend heavily on accurate parameters and reliable communication. In low-inertia systems, these requirements are hard to meet because frequency drops too fast and operating conditions change too often for simple models to stay accurate. These challenges drive the use of data-driven predictors [122].

A related work focuses on estimating the required amount of UFLS using offline learning. For example, [123] proposes a data-driven method for small island power systems in which regression trees are trained using datasets generated from offline dynamic simulations. The trained model provides an estimate of the total UFLS required for a given operating condition and disturbance scenario and is primarily intended to support scheduling-stage planning and reserve allocation. While this approach reduces reliance on detailed analytical models and enables fast estimation once trained, it does not incorporate real-time measurements to adapt load shedding actions after a disturbance has occurred. As such, it supports pre-contingency preparation rather than real-time predictive UFLS control.

As an early example of supervised learning applied to UFLS, [124] develops an Artificial Neural Network (ANN)-based method for an isolated power system. A database of simulated contingency events is constructed offline under varying loading levels and spinning reserve conditions. Using this dataset, a multilayer ANN was then trained with four inputs (total generation, total load, reserve, and initial frequency decay rate) and one output (the required total shedding). During operation, the ANN provides an immediate estimate of the shedding amount. This approach replaces fixed UFLS settings with a learned mapping from the system operating conditions and early frequency dynamics to the required load shedding action.

Reinforcement learning has also been applied to data-driven emergency frequency control and UFLS. For example, [125] proposes a Q-learning-based approach for microgrid emergency control, where a control policy is trained offline using simulation data to map measured system states, such as frequency, voltages, load levels, and DER status, to appropriate load shedding and voltage control actions.

For large-scale power systems, [126] extends this concept using a distributional soft actor-critic framework, enabling risk-aware load shedding by learning a distribution of possible outcomes for each control action. By accounting for uncertainty in system dynamics and disturbance severity, the learned policy determines UFLS actions that maintain frequency within safe limits.

The reviewed predictive UFLS approaches indicate potential advantages over conventional schemes. However, several critical gaps remain. First, many methods depend on detailed system models or offline-generated datasets, and their performance under evolving operating conditions and rare disturbances has not been sufficiently validated. Furthermore, a large proportion of existing studies are evaluated primarily through offline simulations, with limited evidence of their feasibility and reliability in real-time protective applications. Implementation challenges related to computational latency, online decision speed, and compatibility with existing protection and control infrastructure are often overlooked. As a result, the practical deployability of many predictive

UFLS strategies in low-inertia power systems remains uncertain.

2.5 Research Gaps - the Need for Advanced Estimation and Prediction

Extensive research has advanced frequency control in low-inertia power systems through inertia emulation, collaborative microgrid support, demand-side response, and data-driven frequency response modelling. However, despite their demonstrated benefits, these approaches share a critical limitation in that they are predominantly reactive. Inertia emulation and Virtual Synchronous Machine (VSM) schemes typically follow predefined control laws without explicit awareness of disturbance severity, while coordinated and demand side-based approaches often rely on local, event-triggered activation. Even data-driven methods frequently focus on control gain adaptation initiated only after substantial frequency deviations have been observed. Consequently, these strategies have limited capability to proactively shape the frequency response before the frequency nadir is reached.

From a system operation perspective, this limitation persists even as modern dynamic frequency services are increasingly procured through commercial ancillary service mechanisms. Such market-based frameworks allow system operators to adjust reserve procurement according to predicted inertia levels and weather-driven variability. However, the effective technical realisation of these services remains constrained by estimation accuracy and latency. Measurement noise, communication and processing delays, and occasional data losses degrade signal quality, while filtering delays, model simplifications, and sensitivity to outliers further impair real-time performance. As a result, existing schemes struggle to deliver reliable disturbance-related information within the short time window required for fast frequency control.

The evolution of frequency nadir prediction reveals a hierarchical progression of limitations across existing methodologies. Traditional analytical approaches are constrained by structural simplifications that fail to capture the rapid, non-linear dynamics of IBR-dominated systems. While extended analytical methods incorporate more complex structures, they remain mathematically rigid and highly sensitive to pre-fault parameter assumptions which are often inaccurate in time-varying grids. Furthermore, data-driven approaches, despite their computational efficiency, face significant challenges related to input data accuracy. These methods typically rely on precision features that are difficult to obtain reliably in real-time operation. Consequently, uncertainties in the estimated inputs propagate through the models, reducing the overall reliability of existing prediction frameworks.

Another key parameter characterising system behaviour during frequency disturbances is the RoCoF. Traditional hardware-based and spectral techniques are primarily constrained by an inherent balance between noise suppression and reporting latency, as the extended windowing required for numerical stability obscures fast post-disturbance dynamics. Filtering-based approaches, while improving signal stability, remain limited by their reliance on fixed system models and by recursive smoothing effects that delay or attenuate the tracking of abrupt and non-uniform frequency changes during severe transients. Furthermore, existing data-driven methods often require precision input features, such as real-time system inertia and reserve levels, which are impractical to obtain accurately during the critical milliseconds following a contingency. Uncertainties in these estimated inputs propagate through the models, reducing the overall reliability of the reported RoCoF. Consequently, a clear technical gap remains for a framework capable of delivering low-latency and physically consistent RoCoF estimation by operating directly on raw, noisy measurement signals without excessive reliance on uncertain system-level parameters.

Most existing power imbalance estimation methods derive the imbalance indirectly from RoCoF measurements based on the swing equation. Consequently, the accuracy of power imbalance estimation is fundamentally constrained by the quality of RoCoF estimation. As discussed in the previous subsection, RoCoF measurements during disturbances are highly sensitive to noise, filtering, and transient effects. These limitations directly propagate into power imbalance estimation, making it difficult to obtain reliable imbalance magnitude estimates during the critical early stages of a disturbance.

These reactive, fixed-threshold schemes do not provide the real-time adaptability required for a system with rapidly evolving dynamics. Because they lack predictive capability, such schemes often face a critical balance between response speed and accuracy. They may either act too late to arrest a sharp frequency decline or alternatively shed excessive load due to a limited understanding of the disturbance severity. Furthermore, static planning fails to capture the complex interactions among fast-acting services, such as FFR and synthetic inertia, which can lead to unpredictable system behaviour when multiple resources operate simultaneously. As a result, a clear gap remains in existing schemes, which lack the capability to provide timely and informative decision support regarding disturbance severity and appropriate intervention timing in low-inertia grids.

To address these fundamental gaps, Artificial Intelligence (AI) and Machine Learning (ML) offer an effective means of addressing these limitations by enabling learning-based prediction and adaptive decision support [127]. AI refers to the general field concerned with developing systems that can support tasks such as reasoning, learning, and decision making. Within this field, ML focuses on methods that allow models to learn relationships directly from data and simulations, rather than relying on fixed rules defined in advance. In this way, ML provides the practical techniques that allow AI-based systems to adapt their behaviour based on observed system responses. In power system applications, such learning-based approaches are particularly valuable for capturing complex non-linear dynamics that are difficult to represent accurately using fixed analytical models.

2.6 Summary

This chapter has provided a comprehensive review of traditional and data-driven approaches for frequency control. The discussion covered the conventional control structure, analytical methods for frequency nadir prediction, existing UFLS schemes, and techniques for estimating RoCoF and power imbalance size.

Traditional methods are largely model-based and rely on simplified swing equation representations and fixed thresholds. Consequently, they operate in a mainly reactive manner. In such schemes, the frequency nadir is calculated using approximate formulas, UFLS is activated only when preset limits are exceeded, and the power imbalance is estimated based on steady-state assumptions. While these approaches are computationally simple and well-established, their accuracy and reliability are significantly reduced in low-inertia and highly variable operating conditions.

In addition, the chapter has also reviewed recent data-driven and ML approaches that aim to overcome these limitations. These models have been proposed for the direct prediction of the frequency nadir, for predictive UFLS based on trajectory prediction, and for the real-time estimation of power imbalance size using high resolution data. By learning non-linear mappings from operating conditions to key frequency indicators, these methods use richer system information. They provide fast, measurement-driven assessments without depending on detailed dynamic models. However, existing research has largely focused on direct control implementations or narrowly scoped prediction tasks, with comparatively limited attention given to the development of proactive, situation-aware, and practically deployable decision-support methodologies. As a result, a clear gap remains in translating recent advances in learning-based modelling into effective real-time operational support for low-inertia power systems.

In summary, it can be concluded that frequency security is experiencing a conceptual shift from reactive to proactive frequency control. Instead of simply reacting to measured frequency deviations, emerging data-driven approaches aim to predict system behaviour. This evolution enables to more resilient frequency control in low-inertia power systems, while also highlighting the need for further research on the secure integration of ML components into existing architectures.

References

[1] National Grid, "Frequency response obligations: Statutory, code and licence requirements," Tech. Rep., 2012. [Online]. Available: `https://www.nationalgrid.com/sites/default/files/documents/16889-Frequency%20Response%20Obligations.pdf`.

[2] J. Undrill, “Primary frequency response and control of power system frequency,” Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, USA, Tech. Rep. LBNL-2001105, Feb. 2018.

[3] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power Generation, Operation, and Control*, 3rd. Hoboken, NJ, USA: John Wiley & Sons, 2013.

[4] M. Mditshwa, K. A. Folly, and D. T. Oyedokun, “Enhancing automatic generation control (agc) for frequency stability in renewable-dominated power grids: Challenges, gaps, and future directions,” *Heliyon*, vol. 12, no. 1, e44305, 2026.

[5] U. Datta, A. Kalam, and J. Shi, “Battery energy storage system for aggregated inertia-droop control and a novel frequency dependent state-of-charge recovery,” *Energies*, vol. 13, no. 8, p. 2003, 2020.

[6] National Grid, “Future requirements for balancing services,” Tech. Rep., 2016. [Online]. Available: `https://www.nationalgrid.com/sites/default/files/documents/Requirements%5C%20Report%5C%202016.pdf`.

[7] European Network of Transmission System Operators for Electricity (ENTSO-E), “Policy 1: Load-frequency control and performance,” Brussels, Belgium, Tech. Rep., Mar. 2009.

[8] National Grid, “GC0048: Requirements for generators – GB banding thresholds workgroup report,” Warwick, UK, Tech. Rep., Mar. 2016. [Online]. Available: `https://www.neso.energy/document/12381/download`.

[9] National Energy System Operator (NESO). “Mandatory Frequency Response (MFR).” (), [Online]. Available: `https://www.neso.energy/document/92441/download`.

[10] M. Nedd, K. Bell, and C. Booth, “Containing loss risk in a low inertia gb power system,” in *2018 IEEE International Conference on Environment and Electrical Engineering and 2018 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, 2018, pp. 1–6.

[11] National Energy System Operator (NESO). "Clean power 2030 — annex 3: Operability and operations analysis." (2024), [Online]. Available: `https://www.neso.energy/document/346801/download`.

[12] Q. Hong, M. Nedd, S. Norris, *et al.*, "Fast frequency response for effective frequency control in power systems with low inertia," *The Journal of Engineering*, vol. 2019, no. 16, pp. 1696–1702, 2019.

[13] L. Meng, J. Zafar, S. K. Khadem, *et al.*, "Fast frequency response from energy storage systems—a review of grid standards, projects and technical issues," *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1566–1581, 2020.

[14] D. Greenwood, K. Lim, C. Patsios, P. Lyons, Y. Lim, and P. Taylor, "Frequency response services designed for energy storage," *Applied Energy*, vol. 203, pp. 115–127, 2017.

[15] J. Matevosyan, *Frequency response*, Lecture materials, Global Power System Transformation (G-PST) Consortium.

[16] National Grid ESO, "Firm frequency response (FFR): Article 18 Consultation," Warwick, United Kingdom, Tech. Rep., Sep. 2022. [Online]. Available: `https://www.nationalgrideso.com`.

[17] National Energy System Operator (NESO), "Incident report: 22nd december 2023 low frequency event," Technical Report, Feb. 2024. [Online]. Available: `https://www.neso.energy/document/321186/download`.

[18] National Grid, "Invitation to tender for enhanced frequency response (EFR)," Invitation to Tender, version 2.2, 2016. [Online]. Available: `https://www.nationalgrid.com/sites/default/files/documents/Enhanced%20Frequency%20Response%20ITT%20v2_2%20clean.pdf`.

[19] X. Luo, J. Wang, J. D. Wojcik, *et al.*, "Review of voltage and frequency grid code specifications for electrical energy storage applications," *Energies*, vol. 11, no. 5, p. 1070, 2018.

[20] National Energy System Operator (NESO), "Dynamic response services consultation," Tech. Rep., 2024. [Online]. Available: `https://www.neso.energy/document/320556/download`.

[21] National Energy System Operator (NESO). "Dynamic services (dc/dm/dr)." (2025), [Online]. Available: `https://www.neso.energy/industry-information/balancing-services/frequency-response-services/dynamic-services-dcdmdr`.

[22] S. Homan and S. Brown, "The future of frequency response in great britain," *Energy Reports*, vol. 7, pp. 56–62, 2021, ISSN: 2352-4847.

[23] National Energy System Operator (NESO), "Static firm frequency response (sffr) guidance," Warwick, United Kingdom, Tech. Rep., version 1.0, 2024. [Online]. Available: `https://www.nationalgrideso.com/document/315181/download`.

[24] National Energy System Operator (NESO), *Static firm frequency response (SFFR)*, 2025. [Online]. Available: `https://www.neso.energy/industry-information/balancing-services/frequency-response-services/static-firm-frequency-response-sffr`.

[25] National Energy System Operator (NESO), *Static FFR procurement rules*, version 2.0, Warwick, United Kingdom, Nov. 2025. [Online]. Available: `https://www.neso.energy`.

[26] R. K. Sarojini, K. Palanisamy, P. Sanjeevikumar, and J. B.-H. Nielsen, "Inertia emulation control technique based frequency control of grid-connected single-phase rooftop photovoltaic system with battery and supercapacitor," *IET Renewable Power Generation*, vol. 14, no. 7, pp. 1156–1163, 2020.

[27] M. Jafari, G. B. Gharehpetian, and A. Anvari-Moghaddam, "On the role of virtual inertia units in modern power systems: A review of control strategies, applications and recent developments," *International Journal of Electrical Power & Energy Systems*, vol. 159, p. 110 067, 2024.

[28] H. Bevrani, B. François, and T. Ise, *Microgrid dynamics and control.* John Wiley & Sons, 2017.

[29] W. Farmer and A. J. Rix, "Optimising power system frequency stability using virtual inertia from inverter-based renewable energy generation," in *2019 International Conference on Clean Electrical Power (ICCEP)*, 2019, pp. 394–404.

[30] D. Niu, J. Fang, W. Yau, and S. Goetz, "Comprehensive evaluation of energy storage systems for inertia emulation and frequency regulation improvement," *Energy Reports*, vol. 9, pp. 2566–2576, 2023.

[31] G. Song, B. Cao, and L. Chang, "Review of grid-forming inverters in support of power system operation," *Chinese Journal of Electrical Engineering*, vol. 8, no. 1, pp. 1–15, 2022.

[32] M. M. Hasan, D. Razmi, O. Babayomi, I. Davidson, V. Terzija, and Z. Zhang, "Advanced control and protection strategies for grid-forming inverters in microgrids—a review," *International Journal of Electrical Power & Energy Systems*, vol. 172, p. 111 297, 2025.

[33] H. Golpîra, A. R. Messina, and H. Bevrani, "Emulation of virtual inertia to accommodate higher penetration levels of distributed generation in power grids," *IEEE Transactions on Power Systems*, vol. 34, no. 5, pp. 3384–3394, 2019.

[34] H. Bevrani, T. Ise, and Y. Miura, "Virtual synchronous generators: A survey and new perspectives," *International Journal of Electrical Power & Energy Systems*, vol. 54, pp. 244–254, 2014.

[35] J. Fang, H. Li, Y. Tang, and F. Blaabjerg, "Distributed power system virtual inertia implemented by grid-connected power converters," *IEEE Transactions on Power Electronics*, vol. 33, no. 10, pp. 8488–8499, 2017.

[36] M. F. M. Arani and E. F. El-Saadany, "Implementing virtual inertia in dfig-based wind power generation," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1373–1384, 2012.

[37] J. Ma, Y. Qiu, Y. Li, W. Zhang, Z. Song, and J. S. Thorp, "Research on the impact of dfig virtual inertia control on power system small-signal stability considering the phase-locked loop," *IEEE Transactions on Power Systems*, vol. 32, no. 3, pp. 2094–2105, 2016.

[38] A. Fernández-Guillamón, E. Gómez-Lázaro, E. Muljadi, and Á. Molina-García, "Power systems with high renewable energy sources: A review of inertia and frequency control strategies over time," *Renewable and Sustainable Energy Reviews*, vol. 115, p. 109 369, 2019.

[39] A. A. Milani, M. T. A. Khan, A. Chakrabortty, and I. Husain, "Equilibrium point analysis and power sharing methods for distribution systems driven by solid-state transformers," *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 1473–1483, 2018.

[40] H. Bevrani and T. Hiyama, *Intelligent automatic generation control.* CRC press New York, 2011.

[41] J. Zhao, X. Lyu, Y. Fu, X. Hu, and F. Li, "Coordinated microgrid frequency regulation based on dfig variable coefficient using virtual inertia and primary frequency control," *IEEE Transactions on Energy Conversion*, vol. 31, no. 3, pp. 833–845, 2016.

[42] P. Babahajiani, Q. Shafiee, and H. Bevrani, "Intelligent demand response contribution in frequency control of multi-area power systems," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1282–1291, 2016.

[43] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.

[44] A. Molina-García, F. Bouffard, and D. S. Kirschen, "Decentralized demand-side contribution to primary frequency control," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 411–419, 2011.

[45] Y. Mu, J. Wu, J. Ekanayake, N. Jenkins, and H. Jia, "Primary frequency response from electric vehicles in the great britain power system," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, 2013.

[46] J. Zhang, C. Lu, and Y. Han, "Mimo identification of power system with low level probing tests: Applicability comparison of subspace methods," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 2907–2917, 2013.

[47] X. Wang, J. W. Bialek, and K. Turitsyn, “Pmu-based estimation of dynamic state jacobian matrix and dynamic system state matrix in ambient conditions,” *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 681–690, 2017.

[48] K. Tuttelberg, J. Kilter, D. Wilson, and K. Uhlen, “Estimation of power system inertia from ambient wide area measurements,” *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 7249–7257, 2018.

[49] F. Zeng, J. Zhang, Y. Zhou, and S. Qu, “Online identification of inertia distribution in normal operating power system,” *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 3301–3304, 2020.

[50] W. Xing, H. Wang, L. Lu, X. Han, K. Sun, and M. Ouyang, “An adaptive virtual inertia control strategy for distributed battery energy storage system in microgrids,” *Energy*, vol. 233, p. 121 155, 2021.

[51] H. Golpira, H. Bevrani, A. Roman Messina, and B. Francois, “A data-driven under frequency load shedding scheme in power systems,” *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1138–1150, 2023.

[52] G. Cai, C. Jiang, D. Yang, *et al.*, “Data-driven predictive based load frequency robust control of power system with renewables,” *International Journal of Electrical Power & Energy Systems*, vol. 154, p. 109 429, 2023.

[53] V. Terzija, “Improved recursive newton-type algorithm for frequency and spectra estimation in power systems,” *IEEE Transactions on Instrumentation and Measurement*, vol. 52, no. 5, pp. 1654–1659, 2003.

[54] V. Terzija, M. Djuric, and B. Kovacevic, “A new self-tuning algorithm for the frequency estimation of distorted signals,” *IEEE Transactions on Power Delivery*, vol. 10, no. 4, pp. 1779–1785, 1995.

[55] P. M. Anderson and M. Mirheydar, “A low-order system frequency response model,” *IEEE Transactions on Power Systems*, vol. 5, no. 3, pp. 720–729, Aug. 1990.

[56] P. Kundur, *Power System Stability and Control.* McGraw-Hill, 1994.

[57] M. L. Chan, R. D. Dunlop, and F. Schweppe, “Dynamic equivalents for average system frequency behavior following major distribances,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-91, no. 4, pp. 1637–1642, 1972.

[58] J. O’Sullivan and M. O’Malley, “Economic dispatch of a small utility with a frequency based reserve policy,” *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1648–1653, 1996.

[59] S. Dong, X. Fang, J. Tan, N. Gao, X. Cui, and A. Hoke, “A unified analytical method to quantify three types of fast frequency response from inverter-based resources,” in *22nd Wind and Solar Integration Workshop (WIW 2023)*, vol. 2023, 2023, pp. 669–674.

[60] D. Al Kez, A. M. Foley, F. Ahmed, and D. J. Morrow, “Overview of frequency control techniques in power systems with high inverter-based resources: Challenges and mitigation measures,” *IET Smart Grid*, vol. 6, no. 5, pp. 447–469, 2023.

[61] L. Liu, W. Li, Y. Ba, J. Shen, C. Jin, and K. Wen, “An analytical model for frequency nadir prediction following a major disturbance,” *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2527–2536, 2020.

[62] S. Dong, X. Fang, J. Tan, N. Gao, X. Cui, and A. Hoke, “A unified analytical method to quantify three types of fast frequency response from inverter-based resources,” in *22nd Wind and Solar Integration Workshop (WIW 2023), IET Conference Proceedings*, The Institution of Engineering and Technology (IET), Nov. 2023.

[63] J. Lee, S. Jeong, H. Kim, *et al.*, “Analytical approach for fast frequency response control of vsc hvdc,” *IEEE Access*, vol. 9, pp. 91 303–91 313, 2021.

[64] P. Ju, Y. Zheng, Y. Jin, C. Qin, Y. Jiang, and L. Cao, “Analytic assessment of the power system frequency security,” *IET Generation, Transmission & Distribution*, Mar. 2021.

[65] J. Jung, J. Park, H. Lee, and B. Lee, “Evaluation of inertia resource for securing nadir frequency in hvdc interconnected system with high penetration of res,” *Energy Reports*, vol. 9, pp. 1374–1383, 2023.

[66] Z. Li, L. Guo, S. S. Yu, *et al.*, “An efficient full-response analytical model for probabilistic production simulation in fast frequency response reserve planning,” *Energy*, vol. 273, p. 127 268, 2023.

[67] J. Zhou, X. Tong, J. Zhou, R. Liu, and S. Bai, “Feature selection of power system data for frequency prediction with bo-lightgbm-boruta,” *Energy and AI*, vol. 21, p. 100 581, 2025.

[68] National Renewable Energy Laboratory, “Comparison of machine learning methods for frequency nadir prediction,” Golden, CO, USA, Tech. Rep., 2022. [Online]. Available: `https://www.nrel.gov`.

[69] M. J. Jiménez-Navarro, J. M. Riquelme-Dominguez, M. Carranza-García, and F. M. González-Longatt, “A real-time machine learning-based methodology for short-term frequency nadir prediction in low-inertia power systems,” *Neurocomputing*, vol. 626, p. 129 583, 2025.

[70] D. Zografos, T. Rabuzin, M. Ghandhari, and R. Eriksson, “Prediction of frequency nadir by employing a neural network approach,” in *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, 2018, pp. 1–6.

[71] P. Tielens and D. Van Hertem, “The relevance of inertia in power systems,” *Renewable and Sustainable Energy Reviews*, vol. 55, pp. 999–1009, 2016.

[72] Q. Hong, M. Karimi, M. Sun, *et al.*, “Design and validation of a wide area monitoring and control system for fast frequency response,” *IEEE Transactions on Smart Grid*, vol. 11, no. 4, pp. 3394–3404, 2020.

[73] “Frequency risk and control policy,” National Grid Electricity System Operator, London, UK, Tech. Rep., Dec. 2020. [Online]. Available: `https://www.neso.energy/document/183426/download`.

[74] S.-H. Kim, *Electric Motor Control: DC, AC, and BLDC Motors.* Amsterdam, Netherlands: Elsevier, May 2017.

[75] Y. Yu, Y. Liu, C. Qin, and T. Yang, "Theory and method of power system integrated security region irrelevant to operation states: An introduction," *Engineering*, vol. 6, no. 7, pp. 754–777, 2020.

[76] M. Dreidy, H. Mokhlis, and S. Mekhilef, "Inertia response and frequency control techniques for renewable energy sources: A review," *Renewable and Sustainable Energy Reviews*, vol. 69, pp. 144–155, 2017.

[77] "Frequency measurement requirements and usage," ENTSO-E RG-CE System Protection and Dynamics Sub Group, Brussels, Belgium, Tech. Rep. Final Version 7, Jan. 2018. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf`.

[78] A. Gupta, R. Thakur, and S. Murarka, "A review on zero crossing detector," *International Journal of Electronics & Communication Technology*, vol. 3, no. 4, pp. 366–368, 2012, IJECT Vol. 3, Issue 4, Oct–Dec 2012. [Online]. Available: `https://www.iject.org/vol34/3/ankit.pdf`.

[79] R. Wall, "Simple methods for detecting zero crossing," in *IECON'03. 29th Annual Conference of the IEEE Industrial Electronics Society (IEEE Cat. No.03CH37468)*, vol. 3, 2003, 2477–2481 Vol.3.

[80] G. Minas, J. Martins, and C. Couto, "A microcontroller based voltage space vector modulator suitable for induction motor drives," in *ISIE '99. Proceedings of the IEEE International Symposium on Industrial Electronics (Cat. No.99TH8465)*, vol. 2, 1999, 469–473 vol.2.

[81] I.-M. Gordan, "Design and microcontroller implementation of a three phase SCR power converter," *Acta Electrotehnica*, vol. 45, no. 6, pp. 659–664, 2004. [Online]. Available: `https://ie.utcluj.ro/files/acta/2004/Number%206/Paper18_Gordan.pdf`.

[82] R. E. Best, *Phase-Locked Loops: Design, Simulation, and Applications*, 5th ed. New York: McGraw-Hill, 2003.

[83] V. Kaura and V. Blasko, "Operation of a phase locked loop system under distorted utility conditions," *IEEE Transactions on Industry Applications*, vol. 33, no. 1, pp. 58–63, 1997.

[84] European Network of Transmission System Operators for Electricity (ENTSO-E), "Inertia and rate of change of frequency (rocof)," Brussels, Belgium, version 17, Dec. 16, 2020, SPD Inertia Task Force. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Inertia%20and%20RoCoF_v17_clean.pdf`.

[85] H. Yin, Y. Wu, W. Qiu, *et al.*, "Precise rocof estimation algorithm for low inertia power grids," *Electric Power Systems Research*, vol. 209, p. 107 968, 2022.

[86] IEEE Standards Association, *IEEE Standard for Synchrophasor Measurements for Power Systems*, IEEE Std C37.118.1-2011 (Revision of IEEE Std C37.118-2005), 2011.

[87] IEEE Standards Association, *IEEE Standard for Synchrophasor Measurements for Power Systems – Amendment 1: Modification of Selected Performance Requirements*, IEEE Std C37.118.1a-2014 (Amendment to IEEE Std C37.118.1-2011), 2014.

[88] G. Rietveld, P. S. Wright, and A. J. Roscoe, "Reliable rate-of-change-of-frequency measurements: Use cases and test conditions," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 9, pp. 6657–6666, 2020.

[89] A. J. Roscoe, S. M. Blair, B. Dickerson, and G. Rietveld, "Dealing with front-end white noise on differentiated measurements such as frequency and rocof in power systems," *IEEE Transactions on Instrumentation and Measurement*, vol. 67, no. 11, pp. 2579–2591, 2018.

[90] A. G. Phadke and J. S. Thorp, *Synchronized Phasor Measurements and Their Applications.* Springer, 2008.

[91] F. Ding, C. D. Booth, and A. J. Roscoe, "Peak-ratio analysis method for enhancement of lom protection using m-class pmus," *IEEE Transactions on Smart Grid*, vol. 7, no. 1, pp. 291–299, 2016.

[92] P. Romano and M. Paolone, “Enhanced interpolated-dft for synchrophasor estimation in fpgas: Theory, implementation, and validation of a pmu prototype,” *IEEE Transactions on Instrumentation and Measurement*, vol. 63, no. 12, pp. 2824–2836, 2014.

[93] L. Chen, W. Zhao, F. Wang, Q. Wang, and S. Huang, “Enhanced interpolated dynamic dft synchrophasor estimator considering second harmonic interferences,” *Sensors*, vol. 18, no. 9, 2018.

[94] A. Derviškadić, P. Romano, and M. Paolone, “Iterative-interpolated dft for synchrophasor estimation: A single algorithm for p- and m-class compliant pmus,” *IEEE Transactions on Instrumentation and Measurement*, vol. 67, no. 3, pp. 547–558, 2018.

[95] M. Bertocco, G. Frigo, C. Narduzzi, C. Muscas, and P. A. Pegoraro, “Compressive sensing of a taylor-fourier multifrequency model for synchrophasor estimation,” *IEEE Transactions on Instrumentation and Measurement*, vol. 64, no. 12, pp. 3274–3283, 2015.

[96] A. K. Singh and B. C. Pal, “Rate of change of frequency estimation for power systems using interpolated dft and kalman filter,” *IEEE Transactions on Power Systems*, vol. 34, no. 4, pp. 2509–2517, 2019.

[97] R. Ferrero, P. A. Pegoraro, and S. Toscani, “Dynamic synchrophasor estimation by extended kalman filter,” *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 7, pp. 4818–4826, 2020.

[98] R. Ferrero, P. A. Pegoraro, and S. Toscani, “Synchrophasor estimation for three-phase systems based on taylor extended kalman filtering,” *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 9, pp. 6723–6730, 2020.

[99] W. Strielkowski, A. Vlasov, K. Selivanov, K. Muraviev, and V. Shakhnov, “Prospects and challenges of the machine learning and data-driven methods for the predictive analysis of power systems: A review,” *Energies*, vol. 16, no. 10, 2023.

[100] A. Panwar, Z. H. Rather, S. Doolla, A. Liebman, and R. Dargaville, "A machine learning based approach for frequency response prediction in low inertia power system," in *2022 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia)*, 2022, pp. 175–179.

[101] A. B. Kilembe, P. N. Papadopoulos, and R. I. Hamilton, "Data-driven approach to capturing wide-area frequency response dynamics," in *2024 59th International Universities Power Engineering Conference (UPEC)*, IEEE, 2024, pp. 1–6.

[102] R. Azizipanah-Abarghooee, M. Malekpour, M. Paolone, and V. Terzija, "A new approach to the online estimation of the loss of generation size in power systems," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 2103–2113, 2019.

[103] G. Frigo, P. A. Pegoraro, and S. Toscani, "Tracking power system events with accuracy-based pmu adaptive reporting rate," *International Journal of Electrical Power & Energy Systems*, vol. 153, p. 109 384, 2023.

[104] M. N. H. Shazon, N.-A. Masood, and A. Jawad, "Frequency control challenges and potential countermeasures in future low-inertia power systems: A review," *Energy Reports*, vol. 8, pp. 6191–6219, 2022.

[105] J. Sánchez Cortés, M. Rezaei Jegarluei, P. Aristidou, K. Li, and S. Azizi, "Size/location estimation for loss of generation events in power systems with high penetration of renewables," *Electric Power Systems Research*, vol. 219, p. 109 242, 2023.

[106] 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.

[107] B. Delfino, S. Massucco, A. Morini, P. Scalera, and F. Silvestro, "Implementation and comparison of different under frequency load-shedding schemes," in *2001 Power Engineering Society Summer Meeting. Conference Proceedings (Cat. No.01CH37262)*, vol. 1, 2001, 307–312 vol.1.

[108] U. Rudez and R. Mihalic, "Wams-based underfrequency load shedding with short-term frequency prediction," *IEEE Transactions on Power Delivery*, vol. 31, no. 4, pp. 1912–1920, 2016.

[109] U. Rudez and R. Mihalic, "Analysis of underfrequency load shedding using a frequency gradient," *IEEE Transactions on Power Delivery*, vol. 26, no. 2, pp. 565–575, 2011.

[110] V. Chuvychin, N. Gurov, S. Venkata, and R. Brown, "An adaptive approach to load shedding and spinning reserve control during underfrequency conditions," *IEEE Transactions on Power Systems*, vol. 11, no. 4, pp. 1805–1810, 1996.

[111] Y. Tofis, Y. Yiasemi, and E. Kyriakides, "A plug and play, approximation-based, selective load shedding mechanism for the future electrical grid," in *International Workshop on Critical Information Infrastructures Security*, Springer, 2013, pp. 74–83.

[112] UCTE Operation Handbook Team (OpHB-Team), "A1 – appendix 1: Load-frequency control and performance," Tech. Rep. Version 1.8, Level E (final draft), Mar. 2004. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/pre2015/publications/ce/oh/appendix1_v18.pdf`.

[113] "Ieee guide for the application of protective relays used for abnormal frequency load shedding and restoration," *IEEE Std C37.117-2007*, pp. 1–55, 2007.

[114] "Prc-006: Automatic underfrequency load shedding," North American Electric Reliability Corporation (NERC), Tech. Rep., 2018.

[115] V. Terzija and H.-J. Koglin, "Adaptive underfrequency load shedding integrated with a frequency estimation numerical algorithm," *IEE Proceedings - Generation, Transmission and Distribution*, vol. 149, pp. 713–718, 6 2002.

[116] "Rate of change of frequency (rocof) withstand capability and inertia considerations," ENTSO-E, Tech. Rep., 2018.

[117] P. Anderson and M. Mirheydar, "An adaptive method for setting underfrequency load shedding relays," *IEEE Transactions on Power Systems*, vol. 7, no. 2, pp. 647–655, 1992.

[118] S. S. Banijamali and T. Amraee, "Semi-adaptive setting of under frequency load shedding relays considering credible generation outage scenarios," *IEEE Transactions on Power Delivery*, vol. 34, no. 3, pp. 1098–1108, 2019.

[119] V. Terzija, "Adaptive underfrequency load shedding based on the magnitude of the disturbance estimation," *IEEE Transactions on Power Systems*, vol. 21, no. 3, pp. 1260–1266, 2006.

[120] M. Sun, G. Liu, M. Popov, V. Terzija, and S. Azizi, "Underfrequency load shedding using locally estimated rocof of the center of inertia," *IEEE Transactions on Power Systems*, vol. 36, no. 5, pp. 4212–4222, 2021.

[121] U. Rudez and R. Mihalic, "Monitoring the first frequency derivative to improve adaptive underfrequency load-shedding schemes," *IEEE Transactions on Power Systems*, vol. 26, no. 2, pp. 839–846, 2011.

[122] D. Kottick and O. Or, "Neural-networks for predicting the operation of an under-frequency load shedding system," *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1350–1358, 1996.

[123] M. Rajabdorri, M. C. Troffaes, B. Kazemtabrizi, M. Sarvarizadeh, L. Sigrist, and E. Lobato, "Data-driven estimation of the amount of under frequency load shedding in small power systems," *Engineering Applications of Artificial Intelligence*, vol. 139, p. 109 617, 2025.

[124] R. Hooshmand and M. Moazzami, "Optimal design of adaptive under frequency load shedding using artificial neural networks in isolated power system," *International Journal of Electrical Power & Energy Systems*, vol. 42, no. 1, pp. 220–228, 2012.

[125] R. Horri and H. M. Roudsari, "Reinforcement-learning-based load shedding and intentional voltage manipulation approach in a microgrid considering load dy-

namics," *IET Generation, Transmission & Distribution*, vol. 16, no. 17, pp. 3384–3401, 2022.

[126] J. Xie and W. Sun, "Distributional deep reinforcement learning-based emergency frequency control," *IEEE Transactions on Power Systems*, vol. 37, no. 4, pp. 2720–2730, 2022.

[127] N. Kühl, M. Schemmer, M. Goutier, and G. Satzger, "Artificial intelligence and machine learning," *Electronic Markets*, vol. 32, no. 4, pp. 2235–2244, 2022.

Chapter 3

Probabilistic Nadir Prediction and Response Scheduling

Frequency nadir is a critical security metric for system operators when considering the scheduling of frequency response in order to prevent unnecessary UFLS. Traditional approaches often rely on deterministic scheduling of frequency response to ensure that the nadir remains above a prescribed threshold, assuming fixed system conditions while neglecting uncertainties in system inertia, disturbance magnitude, and frequency response delivery. This is typically achieved either by mathematically characterising frequency behaviour or by using offline representative system models.

However, in a highly complex IBR-dominated system, where frequency response may be delivered by different resources and triggered through different mechanisms, mathematical formulations are unlikely to estimate the post-disturbance frequency nadir accurately, and therefore may not dispatch the appropriate volume of frequency response. Offline simulation models are often regarded as sufficiently accurate, provided the system is represented appropriately. However, they are computationally intensive and, in practice, are usually limited to a small number of specific worst-case scenarios that then form the basis for frequency service scheduling. This will not be adequate for future systems with dramatically higher volatility.

This chapter therefore leverage learning-based techniques, where an RF and uncertainty-informed approach is proposed for prediction of frequency nadir. The proposed approach is highly lightweight, enabling fast execution to meet the requirements of volatile system conditions. Furthermore, the training process is carried out using simulation data from credible models that are already available to system operators and currently used for frequency ancillary service scheduling, thus ensuring realism. By incorporating quantified prediction errors from the proposed frequency nadir prediction model into the subsequent optimisation-based scheduling process, the overall approach balances procurement costs with system security.

This chapter is organised as follows: First, frequency nadir prediction based on the RF method is presented, including the analysis of physical influencing factors as well as feature and learning model selection. The RF-based prediction method is then introduced. Then, a sensitivity-based uncertainty quantification method is developed, and the optimal scheduling of frequency response services based on nadir prediction is discussed. Finally, case studies are conducted to evaluate the predictive performance, validate the uncertainty bounds, and demonstrate the effectiveness of the proposed scheduling approach.

3.1 Frequency Nadir Prediction Based on RF

3.1.1 Physical Factors Influencing the Frequency Nadir

The frequency nadir represents the maximum frequency drop following a sudden active power imbalance, and its ultimate depth is fundamentally governed by the dynamic response of the power system to this event. System frequency behaviour after a disturbance can be described by the swing equation, as expressed in (2.3.15), which relates frequency acceleration to the imbalance between mechanical and electrical power and the aggregate system inertia.

Accordingly, the initial RoCoF is inversely proportional to the post-disturbance system inertia (H) and directly proportional to the magnitude of the active power deficit (ΔP). During the initial moments following a disturbance, the depth of the frequency excursion is therefore primarily influenced by the total system inertia and the size of the initiating event. Lower inertia results in a higher RoCoF and a more rapid frequency decline, increasing the sensitivity of the frequency nadir to the size of the power imbalance.

As the frequency continues to deviate following the initial inertial response, the system transitions into the primary frequency response phase, during which the ultimate depth of the frequency nadir is no longer determined by inertia alone, but is strongly influenced by subsequent control actions. During this phase, a multitude of system parameters interact to arrest the frequency decline. These include the inherent load damping factor, which provides a passive stabilising effect, as well as the active frequency response from generating units, which is governed by controller characteristics (e.g. droop), governor deadbands, and response time constants. The combined influence of these mechanisms determines both the timing and effectiveness of frequency arrest.

In modern power systems with a high penetration of IBRs, this process becomes significantly more complex. In addition to conventional primary frequency response, FFR delivered by power electronic converters plays an increasingly important role. The diverse control strategies adopted by IBRs, such as synthetic inertia and FFR, introduce highly non-linear and heterogeneous dynamics. The magnitude and timing of these responses depend on control parameters, activation thresholds, and communication or measurement delays, which may vary significantly between resources.

In conventional power systems dominated by SGs, the parameters governing inertia, droop response, and damping are relatively stable and well characterised, allowing analytical models to provide reasonable approximations of frequency behaviour. However, in low-inertia systems, these parameters can change with operating conditions and over time. Consequently, the relationship between system conditions and the resulting frequency nadir is no longer governed by a fixed set of parameters, but instead emerges from the interaction of multiple uncertain and dynamically coupled factors. This complexity limits the applicability of analytical formulations and motivates approaches that can capture the underlying non-linear mapping between system operating conditions and frequency nadir behaviour.

3.1.2 Feature and Learning Model Selection for Frequency Nadir Prediction

Feature Selection

To accurately estimate the frequency nadir, it is necessary to identify a set of physically meaningful variables that capture the dominant mechanisms driving frequency dynamics. Based on the physical analysis presented in the previous section, the frequency nadir can be modelled as a non-linear function of system inertia, the magnitude of the active power imbalance, and the subsequent active power injections.

To apply this for the proposed model, these physical concepts are mapped to specific system parameters. Let $x = [H, \Delta P, P_{dc}, P_{static}, P_d]^\top \in R^5$ be the selected feature vector and let $y \in R$ represent the post-event frequency nadir target. The total system demand (P_d) is included as an indicator of operating conditions and associated load damping effects, while the DC (P_{dc}) and static response (P_{static}) capacities quantify the available frequency response resources. The specific input features and target variable considered in this study are summarised in Table 3.1.1.

Table 3.1.1: Input features and target variable for the frequency nadir estimation model

Name	**Symbol**	**Description / Unit**
Inertia constant	H	Equivalent system inertia (s)
Disturbance size	ΔP	Largest credible power imbalance size (MW)
DC	P_{dc}	Pre-event DC provision (MW)
Static response (Static)	P_{static}	Pre-event static response (MW)
Demand	P_d	Total system demand (MW)
Target	y	Post-event frequency nadir (Hz)

While modern power systems incorporate an evolving array of frequency response products, such as DM or various sub-types of FFR, the features P_{dc} and P_{static} are selected as representative categories for fast-acting and slow-acting frequency response services, respectively. Regarding the equivalent system inertia H is extracted directly from the Real-Time Digital Simulator (RTDS) environment. In this study, H is aggregated from the individual parameters of all operational synchronous units

$$H = \frac{\sum_{i=1}^{n} S_i H_i}{\sum_{i=1}^{n} S_i} \tag{3.1.1}$$

where H_i and S_i represent the inertia constant and rated capacity of each individual generator i.

While these fundamental parameters dictate the energy balance, disturbances in large transmission networks may also induce regional frequency variations. However, system security and reserve scheduling typically focus on the CoI frequency, which represents the global average frequency trajectory by filtering out local oscillations. Consequently, the global CoI nadir is fundamentally determined by the aggregated system parameters defined in Table 3.1.1. This simplifies the model input without compromising prediction accuracy. This aggregated modelling structure is subsequently validated using a microgrid as a case study, reflecting the CoI assumption.

Learning Model Selection

Mapping these selected features to the frequency nadir introduces complexities that limit the applicability of traditional offline simulations and analytical models. First, the relationship between the input variables and the frequency nadir is inherently non-linear, as the combined effects of inertia, power imbalance, and frequency response mechanisms vary significantly across different operating conditions. Second, complex interactions occur between these features, including disturbance size, available response, and system demand, which further complicates the predictive mapping.

The ensemble-based RF method is effective for this task because it can capture these non-linear mappings and feature interactions while maintaining strong generalisation performance on simulation-based datasets of moderate size [1]. Specifically, the RF regressor enhances predictive performance by aggregating the outputs of multiple decision trees into an ensemble predictor. While an individual decision tree may suffer from high variance and limited generalisation capability, the ensemble structure of RF produces a significantly more stable and robust model by combining multiple diverse learners [1], [2]. Therefore, this chapter adopts RF as the core predictive model for frequency nadir estimation.

For practical deployment, resilience beyond standard operating conditions is essential. Accordingly, the RF architecture is selected for its robustness under structural system transitions and extreme operational uncertainties. It also accommodates structural changes not explicitly represented in the primary feature set. For instance, in the event of a transition to the islanding mode or disconnection of an IBR, the system dynamics change significantly as the connection to the bulk power system is lost. While such a transition alters the magnitude of available inertia and the characteristic frequency response, the RF model can maintain predictive reliability if the training dataset encompasses a sufficiently diverse range of inertia and power imbalance levels that reflect these low inertia states.

As illustrated in Figure 3.1.1, the previously defined feature vector x is provided as a common input to a collection of independently trained decision trees. Diversity across the ensemble is introduced through randomisation during the training process, ensuring that different trees focus on different aspects of the input feature space. This diversity leads to heterogeneous internal structures and decision boundaries across trees, enables the ensemble to capture complex non-linear relationships while mitigating over-fitting.

Each tree in the RF is trained on a sample of the original dataset, and at each node only a random subspace of the input features is considered for splitting. For a given input vector x, the forest prediction is defined as the arithmetic average of the outputs from all individual trees. This ensemble averaging reduces the prediction variance relative to a single decision tree, thereby improving stability and generalisation performance.

From a power system modelling perspective, this ensemble-based structure allows RF to represent the highly non-linear mapping between pre-event operating conditions and the resulting frequency nadir. Moreover, because individual trees are trained on different subsets of the data and features, the influence of noise or extreme samples is typically confined to a subset of trees. Aggregating predictions across the ensemble therefore reduces the impact of outliers, resulting in a more stable predictor. The

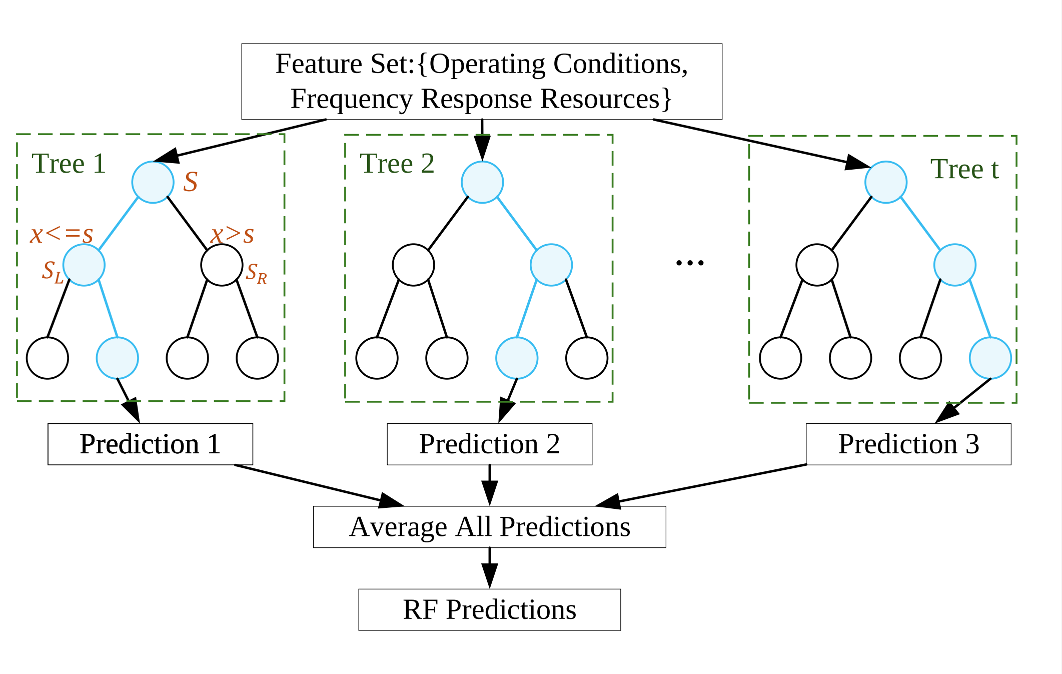

Fig. 3.1.1: Structure of the RF model [3]

detailed construction of individual trees and the associated randomisation mechanisms are described in the following section.

3.1.3 Applying RF for Frequency Nadir Prediction

The objective of the proposed approach is to approximate the non-linear mapping from the pre-event operating state to the post-event frequency nadir. The training dataset is defined as $D = \{(x_i, y_i)\}_{i=1}^{n}$, where $y_i \in R$ represents the observed nadir and x_i is the corresponding feature vector. Consistent with the feature selection analysis in Table 3.1.1, the input vector x_i incorporates key physical determinants of the frequency response

$$x_i = [H, \Delta P, P_{dc}, P_{static}, P_d]^{\top} \tag{3.1.2}$$

To establish this predictive mapping, the RF regressor constructs an ensemble comprising T regression trees, denoted as $\{f_t\}_{t=1}^{T}$. Within this algorithmic architecture, randomisation is applied across two different levels. The tree-level determines the initial setup of each independent learner, dictating the specific subset of training data assigned to a tree before its construction begins. Conversely, the node-level focuses on the individual splitting points within each tree during its training process, dictating the specific subset of features evaluated at each branch. As illustrated in Figure 3.1.2, the detailed mathematical construction of these T independent trees relies on the exact implementation of tree-level bootstrap sampling and node-level feature subspace selection, which collectively enable the ensemble to map complex operating conditions.

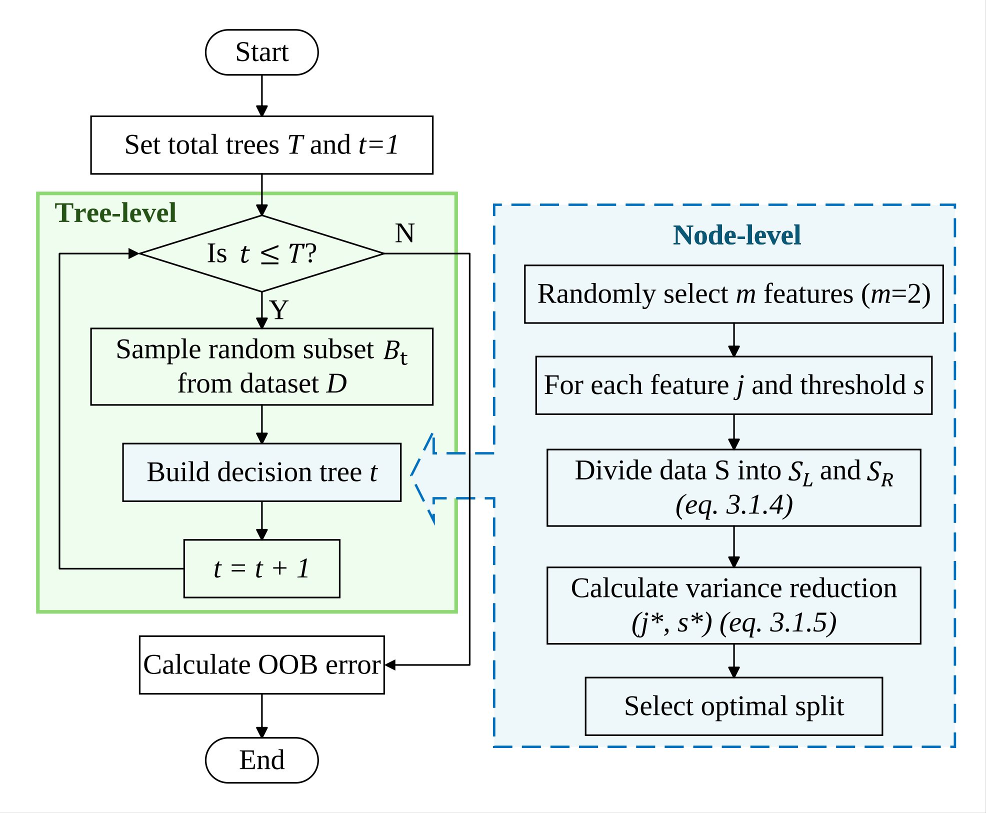

Fig. 3.1.2: Frequency nadir prediction workflow with an RF algorithm

Tree-Level Randomisation: Bagging

As illustrated in the tree-level block of Fig. 3.1.2, tree-level diversity is introduced by training each decision tree on a different randomly resampled subset of the original training data. For each tree t, a specific training subset B_t is generated by sampling n observations from the original dataset D with replacement. This means each sample is drawn and subsequently returned to the dataset, allowing it to be selected multiple times.

Under this random resampling scheme, the probability that a specific observation remains unselected in a single draw is $1 - \frac{1}{n}$. Consequently, for n independent draws, the probability that an observation is excluded from B_t converges as follows

$$\lim_{n\to\infty} \left(1 - \frac{1}{n}\right)^n \approx e^{-1} \approx 0.368 \tag{3.1.3}$$

Approximately 36.8% of the training data are remain unselected for each tree. These unselected samples form the Out-of-Bag (OOB) set, denoted as $O_t = D \setminus B_t$. Because O_t contains observations that were never seen by tree t during its training process, these samples act as unseen test data for that specific tree. This provide an unbiased source for validation without requiring a separate hold-out set.

Node-Level Recursive Partitioning and Feature Subspace Selection

Once the training subset B_t is established, each regression tree is constructed by repeated binary splits. As illustrated in the node-level block of Fig. 3.1.2, each node N applies two connected steps:

1. Random Feature Selection. The input feature vector x_i contains $p = 5$ variables (as defined in Eq. 3.1.2). At each individual node, instead of evaluating all p features, the algorithm randomly selects a smaller subset of size m (where $m < p$). In this implementation, $m = 2$ is adopted. This configuration aligns with standard ensemble learning methodologies for regression, where a restricted subset (e.g. $m \approx \sqrt{p}$ to $p/3$) balances the predictive accuracy of individual trees and the structural diversity across the forest. Restricting the candidate features at a node prevents a single dominant predictor from appearing in every split. For instance, even if H is the strongest predictor overall, its exclusion from the random subset at certain nodes forces the model to use alternative predictors such as P_{dc} or ΔP. This forced decorrelation ensures that the model explores the influence of secondary and critical factors including reserve availability and load demand, thereby enhancing the reliability of the ensemble and improving its adaptability to variations in any single input.

2. Split Optimisation. Following feature selection, the algorithm searches for the optimal split point using the m randomly selected features. For each feature j in this small subset, and for every possible threshold s, the current data S is divided into left (S_L) and right (S_R) groups

$$S_L = \{i \in S : x_{ij} \leq s\}, \quad S_R = \{i \in S : x_{ij} > s\} \tag{3.1.4}$$

The objective is to find the specific feature and threshold (j^*, s^*) that maximises the reduction in sample variance

$$(j^*, s^*) = \arg\max_{j,s} \left[\mathrm{Var}(S) - \frac{|S_L|}{|S|}\mathrm{Var}(S_L) - \frac{|S_R|}{|S|}\mathrm{Var}(S_R) \right] \tag{3.1.5}$$

where $\mathrm{Var}(\cdot)$ is the sample variance. A terminal node stores the mean of the training targets assigned to that region, and following the sequence of splits determines the tree output for a new input. Averaging the outputs of all trees reduces the influence of fluctuations from individual simulations and ensures that the predictor reflects the actual relationships rather than incidental noise.

For frequency nadir prediction, this variance reduction criterion has a direct connection to system physics. Minimising the variance within a node groups together operating conditions that lead to similar post-event frequency trajectories. These groups correspond to comparable stability conditions, such as similar levels of inertia. When the response service approaches its delivery limits, the nadir becomes highly sensitive to small changes in the imbalance size. This creates a clear boundary between stable regulation and degradation, and splits near this boundary produce substantial reductions in node variance. As a result, the optimisation process naturally identifies these transition points.

3.2 Sensitivity-Based Uncertainty Quantification Method

While the RF model provides accurate point predictions, the inherent measurement uncertainty in system inertia H necessitates a formal quantification of the prediction risk. This section develops a sensitivity-based approach to construct uncertainty bands, specifically isolating the impact of H on the frequency nadir. The proposed methodology begins by establishing an empirical justification for focusing on inertia through feature importance analysis, and mathematical formulation of sensitivity bands.

3.2.1 Uncertainty in System Inertia

Before formulating the uncertainty quantification framework, it is essential to evaluate the relative contribution of the input features to the prediction model. The swing equation relation between inertial response and the nadir suggests that H should be a dominant factor, and the model-based analysis confirms this.

The feature importance is quantified using the permutation method on the test set, as illustrated in Fig. 3.2.1. The bar height represents the mean drop in R^2 across repeated permutations, while the horizontal error bars indicate the corresponding standard deviation. The inertia H has the largest average importance (approximately 0.76), significantly exceeding the contributions of disturbance size ΔP and system demand P_d. Although H shows a wider variation in importance across different permutations, its ranking remains consistently the highest, confirming it as the most influential input for the predictive model.

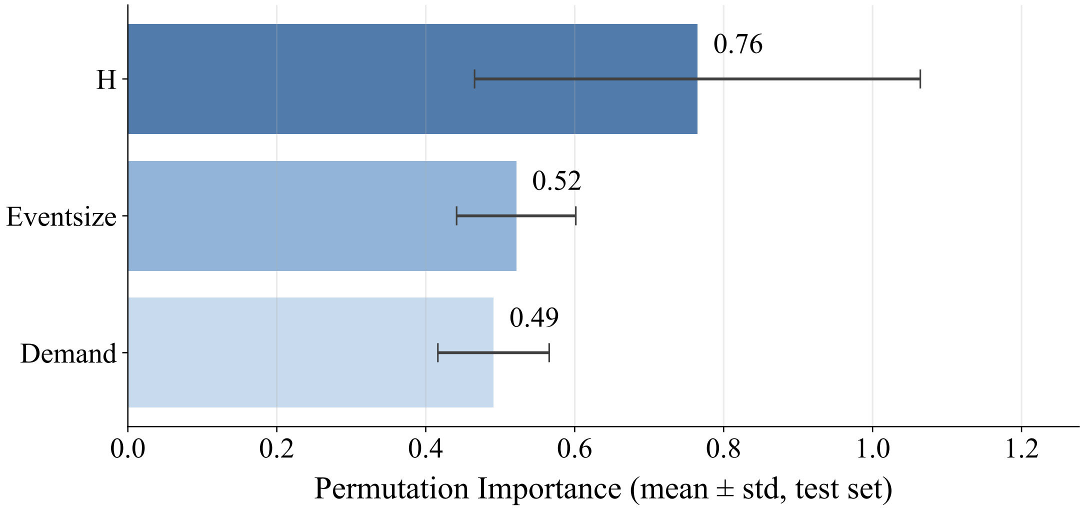

Fig. 3.2.1: Feature permutation importance on the test set

At the system level, the equivalent inertia H aggregates the inertial contributions of all online resources. For SGs, H is a stable parameter derived from physical machine ratings. However, for IBRs-dominated system the inertial contribution is dictated by control strategies rather than rotating physical mass, as power electronic converters decouple the resource speed from the grid frequency. Parameters such as virtual inertia constants and droop gains can vary significantly across different time scales and operating modes. Consequently, in systems with high IBR penetration, the aggregate H becomes increasingly uncertain. This variability motivates the subsequent sensitivity analysis to quantify how measurement and estimation uncertainties in H affect the reliability of the frequency nadir predictions.

3.2.2 Mathematical Formulation of Sensitivity Bands

With the critical role of inertia in prediction reliability established, this subsection develops the sensitivity-based approach to construct uncertainty bands. Given the non-negligible uncertainty in H, a local sensitivity band is constructed by varying H within a realistic range while holding other inputs fixed at their observed values. The measurement of H is modelled using an additive error term

$$H_{meas} = H_{true} + \varepsilon_H, \qquad \varepsilon_H \sim \mathcal{N}(0, \sigma_H^2) \tag{3.2.1}$$

where the error ε_H is assumed to be zero-mean (free from systematic bias) and follows a Gaussian distribution with standard deviation σ_H. For a two-sided 95% confidence level, the standard normal quantile is $z_{0.975} = \Phi^{-1}(0.975) \approx 1.96$. Thus, a variation of $\pm 1.96\sigma_H$ captures approximately 95% of the expected measurement errors.

For each test sample x_i, the local sensitivity is assessed by perturbing only the inertia component h_i (the normalised value of H in x_i. Two boundary values for H are defined based on the 95% confidence interval

$$\tilde{h}_i^- = h_i - z_{0.975}\sigma_H, \qquad \tilde{h}_i^+ = h_i + z_{0.975}\sigma_H \tag{3.2.2}$$

Correspondingly, two modified input vectors, $\tilde{x}_i^-$ and $\tilde{x}_i^+$, are generated by substituting the original H value in x_i with the calculated lower and upper bounds

$$\tilde{x}i^{-} : xi, j_H \leftarrow \tilde{h}_i^{-}, \qquad \tilde{x}i^{+} : xi, j_H \leftarrow \tilde{h}_i^{+} \tag{3.2.3}$$

By inputting these modified vectors $\tilde{x}_i^{-}$ and $\tilde{x}_i^{+}$ into the trained model f, the corresponding predicted responses are obtained

$$\hat{y}_i^{-} = f(\tilde{x}_i^{-}), \qquad \hat{y}_i^{+} = f(\tilde{x}_i^{+}) \tag{3.2.4}$$

The sensitivity band for sample i is constructed as the interval $[\min(\hat{y}_i^{-}, \hat{y}_i^{+}), \max(\hat{y}_i^{-}, \hat{y}_i^{+})]$. This range represents the potential variation in the prediction caused specifically by the measurement error in H (at the 95% confidence level). It is important to note that this band isolates the influence of H and excludes other sources of uncertainty.

To evaluate the performance of the generated prediction intervals $[\hat{y}_i^{L}, \hat{y}_i^{U}]$, two complementary metrics are employed: the Prediction Interval Coverage Probability (PICP) and the Mean Prediction Interval Width (MPIW). These are defined as:

$$PICP = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\{\hat{y}_i^{\mathrm{L}} \leq y_i \leq \hat{y}_i^{\mathrm{U}}\} \tag{3.2.5}$$

$$MPIW = \frac{1}{n} \sum_{i=1}^{n} \left(\hat{y}_i^{\mathrm{U}} - \hat{y}_i^{\mathrm{L}}\right) \tag{3.2.6}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function and n is the total number of test samples.

PICP estimates the probability that the true target y_i falls within the predicted bounds, while MPIW measures the average width of the intervals. The empirical coverage measured by PICP is expected to match the nominal confidence level (e.g. PICP $\approx$ 95% for $\alpha = 0.05$). Values significantly below $1 - \alpha$ indicate under-coverage (intervals are too narrow or miscalibrated), whereas values significantly above suggest the intervals are excessively wide. For a given coverage level, a smaller MPIW is preferred as it implies tighter bounds. Consequently, PICP and MPIW must be evaluated in tandem since coverage reflects reliability while width reflects precision.

3.3 Optimal Scheduling of Frequency Response Services Based on Nadir Prediction

This section demonstrates the operational utility of the proposed nadir predictor in supporting the frequency response scheduling process. To ensure the approach remains computationally efficient for real-time scheduling, P_{dc} and P_{static} are treated as final delivered capacities provided by the respective resources. By focusing on the net frequency gain (β_{dc} and β_{static}) evaluated at the current operating point, the non-linear predictive mapping is simplified into a linear sensitivity-based formulation

$$\hat{y} \approx \beta_0 + \beta_{dc} P_{dc} + \beta_{static} P_{static} + \beta_H H + \beta_d P_d \tag{3.3.1}$$

here, the coefficients β represent the local sensitivities obtained by finite-difference analysis of the trained RF model.

Given a security threshold y_{min} (e.g. 49.5 Hz), the potential security shortfall r is first quantified. This metric represents the frequency gap between the required threshold and the predicted nadir in the absence of additional frequency services

$$r \;=\; y_{\min} - \big(\beta_0 + \beta_H H + \beta_d P_d\big) \tag{3.3.2}$$

If $r \leq 0$, the system is considered secure, and no mitigation is required ($P_{dc} = P_{static} = 0$). Conversely, if $r > 0$, frequency response services must be scheduled to address this deficit. The scheduling problem is structured as a linear optimisation to determine the optimal combination of P_{dc} and P_{static} that satisfies the security requirement $\hat{y} \geq y_{min}$ while minimising operational costs

$$\begin{aligned} \min_{P_{dc}, P_{static} \geq 0} \quad & c_{dc} P_{dc} + c_{static} P_{static} \\ \text{subject to} \quad & \beta_{dc} P_{dc} + \beta_{static} P_{static} \geq r \end{aligned} \tag{3.3.3}$$

Here, the parameters c_{dc} and c_{static} denote unit costs or preference weights, which are both set to 1 in this study to minimise the total procured capacity. To robustly incorporate inertia uncertainty, the security constraint is enforced based on the lower bound H_{lb} of the 95% confidence interval. This introduces a security margin δ, defined as

$$\delta = \hat{y}(H) - \hat{y}(H_{lb}) \tag{3.3.4}$$

This formulation ensures that the scheduled response services are adequate to maintain system stability even under the worst-case credible inertia scenario.

3.4 Case Studies

3.4.1 Power System Modelling and Scenario Generation

Establishment of the Simulation and Microgrid Structure

To implement and validate the proposed frequency nadir prediction scheme, a microgrid system is implemented in RTDS and used for scenario generation and testing. It is based on a modified IEEE 9-bus benchmark as illustrated in Fig. 3.4.1. The Point of Common Coupling (PCC) with the main grid is at Bus 1; distributed generation, storage and loads are located in Buses 2-7. The key elements are:

- **Bus 1 (PCC with the main grid)** Bus 1 serves as the interface between the microgrid and the bulk power system. It accommodates Load 0 and Load Extra, which represent demand directly supplied from the main grid side, and it functions as the entry point through which external disturbances may affect the microgrid. In addition, Bus 1 enables the import and export of both active and reactive power, making it the critical node that links the microgrid to the wider grid environment.

- **Bus 2 (PV and Feeder Bus)** Bus 2 connects to the main entry point and acts as a feeder bus supplying Load 1 and Load 2. It hosts the PV unit (1 MVA), representing converter-interfaced solar generation.

- **Bus 3 (Local Generation Hub)** Bus 3 integrates a CHP SG (G2), representing a dispatchable thermal unit, and supplies Load 3.

- **Bus 4 (Storage Node)** Bus 4 connects the BESS, which provides fast frequency support, and supplies Load 4.
- **Bus 5 (Virtual Synchronous Machine (VSM) Node)** Bus 5 hosts the VSM, an IBR unit configured to emulate the inertia and damping of a SG. It also supplies Load 5.
- **Bus 6 (Wind and Load Center)** Bus 6 is a major load center hosting Load 6, Load 7 and Load 8. It also connects a wind turbine generator.
- **Bus 7 (Hydro Generation)** Bus 7 connects the hydro unit (G1), a SG with governor and automatic voltage regulator control. It also supplies Load 10, modelling distributed demand in this section of the grid.

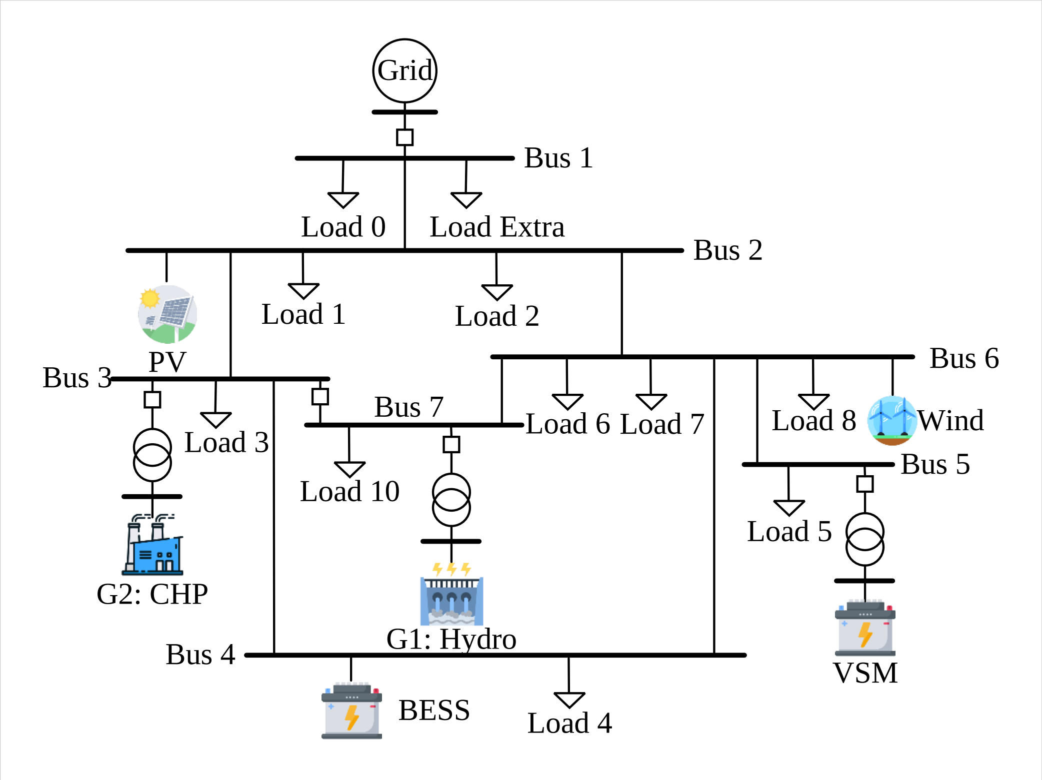

Fig. 3.4.1: Configuration of the microgrid model implemented in RTDS simulator

Feeder lines are represented by π-equivalent models with positive-sequence parameters (R_p, X_p, X_{cp}) and zero-sequence parameters (R_z, X_z, X_{cz}). Table 3.4.1 lists the values used in dynamic simulations.

Table 3.4.1: Line π-model parameters of the microgrid network

Line	R_p **(ohm)**	X_p **(ohm)**	X_{cp} **(MΩ)**	R_z **(ohm)**	X_z **(ohm)**	X_{cz} **(MΩ)**
Line 12	0.0865	0.0865	0.5221	0.1753	0.8994	0.8994
Line 23	0.1038	0.2590	0.4351	0.2104	1.0792	1.0616
Line 26	0.2249	0.5612	0.2008	0.2008	0.2008	0.4900
Line 34	0.3979	0.9929	0.1135	0.8064	4.1370	0.2769
Line 65	0.2941	0.7334	0.1536	0.5960	3.0578	0.3747
Line 64	0.2422	0.6044	0.1865	0.4908	2.5182	0.4550
Line 67	0.0865	0.2159	0.5221	0.1753	0.8994	1.2739

SG parameters are given in Table 3.4.2. Rated apparent power, voltage levels and inertia constants determine the available rotational inertia, a key factor shaping the post-event frequency response.

Table 3.4.2: SG parameters used in the RTDS microgrid model

Generator	**Rated power** S_{rated} **(MVA)**	**Rated Voltage (kV)**	**Inertia constant** H **(s)**
G1: Hydro (Bus 7)	1	0.4	2
G2: CHP (Bus 3)	1	0.4	2

IBRs include PV, wind, BESS and VSM units with distinct control modes. The PV operates with Maximum Power Point Tracking (MPPT) and grid-following PQ control; the wind and BESS use grid-following PQ control; the VSM is grid-forming to provide virtual inertia and frequency support. Table 3.4.3 summarises rated capacities and key control parameters. In this configuration, the BESS is the primary provider of FFR to enhance the dynamic response by arresting the frequency fall. In addition, the static response and the long-term frequency stability under severe contingencies are provided by supplementary control measures, such as the proposed adaptive load shedding scheme, which replaces or complements traditional generation reserves.

Table 3.4.3: Parameters of converter-interfaced units in the RTDS microgrid model

Unit	S_{rated} **(MVA)**	V_{rated} **(kV)**	**Control type**	**Ref/Init Values**	**Key Parameters**
PV (Bus 2)	1	0.4	MPPT + PQ	$P_{ref} = 0.6$ pu	$\tau = 1$ ms
BESS (Bus 4)	1	0.4	PQ	Droop=5%	$\tau = 1$ ms
Wind (Bus 6)	1	0.4	PQ	$P_{ref} = 0.2$ pu	Limit=1.2 pu
VSM (Bus 5)	1	0.4	Virtual Sync	$H = 2$ s	$D = 50$

Data Collection and Scenario Generation

The data generation process, illustrated in Fig. 3.4.2, aims to produce a diverse and valid dataset for training and validating the ML models.

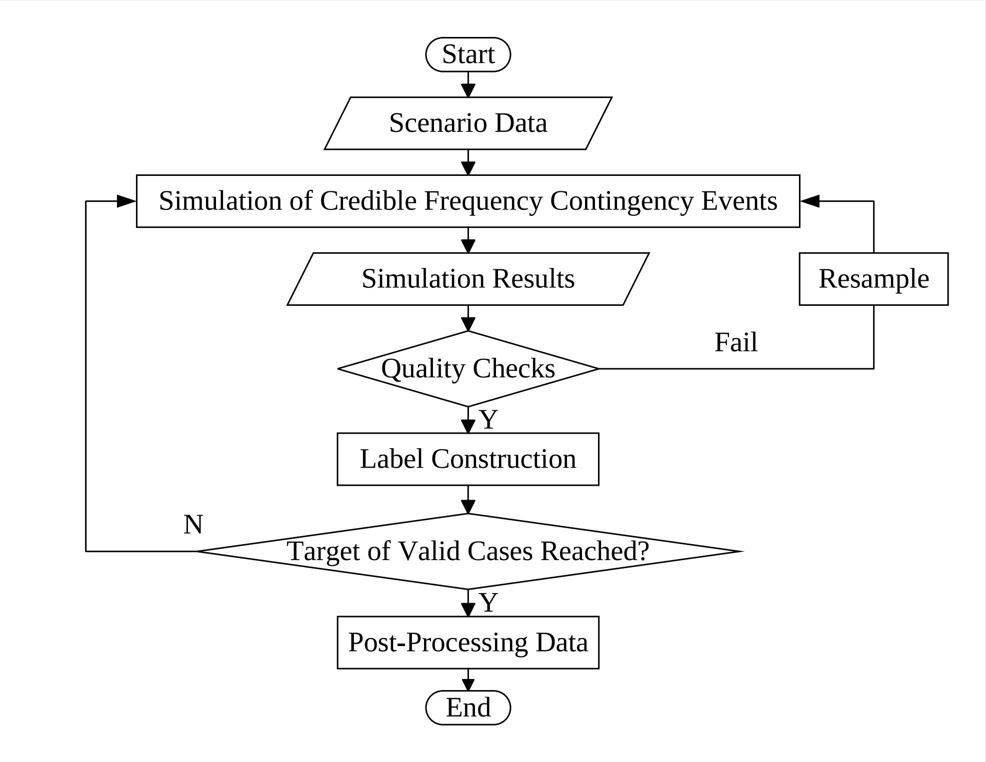

Fig. 3.4.2: Flowchart of data generation process

Each simulation case is defined by a scenario vector that specifies the system operating conditions and the associated disturbance event. To ensure the dataset covers the full operational range, key variables such as inertia constants ($H \in [2, 8]\,\text{GVA·s}$), renewable penetration, and power imbalances ($\Delta P \in [0.1, 1.8]$ GW) are sampled. Each scenario further specifies the system parameters, including demand level, generation mix, and controller settings for frequency response services (e.g. static response and DC), as well as the simulation configuration, comprising the simulation horizon T, integration step Δt, and the disturbance schedule defined by its timestamp and magnitude.

The microgrid dynamics are simulated, with all output signals automatically aligned to the event timestamp to ensure feature-label consistency. A quality control process then filters the results, discarding cases with missing data or physical inconsistencies (e.g. negative demand). Failed scenarios are resampled until the target number of valid cases is reached. Finally, pre-event features are paired with the post-event frequency nadir to form the labels. The resulting dataset is partitioned into training, validation, and test sets to evaluate the proposed RF approach.

3.4.2 Performance Evaluation and Sensitivity Analysis

To validate the proposed approach, this section evaluates the RF model from three perspectives: overall predictive accuracy, sensitivity-based uncertainty quantification, and scheduling of frequency response services.

Determination of Optimal Ensemble Configuration and Assessment of Predictive Performance

To determine the optimal ensemble configuration, the OOB error is evaluated during the model training. As established in Fig. 3.1.2, OOB samples provide an internal estimate of generalisation performance. For a specific training sample i, the OOB prediction is calculated by averaging the outputs only from the subset of trees that did not observe i during training. The aggregation is simply calculated as

$$\hat{y}_i^{oob} = \frac{1}{N_i} \sum f_t(x_i) \tag{3.4.1}$$

where N_i is the number of trees for which i is OOB. The resulting MSE_{OOB}

$$RMSE_{OOB} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i^{OOB})^2} \tag{3.4.2}$$

serves as an unbiased approximation of the test error.

The optimal ensemble size T is determined by monitoring this OOB error. All other hyperparameters (e.g. feature subset size m, maximum tree depth, and minimum leaf size) are held constant while the forest is expanded in steps of ΔT (e.g. 10 to 25 trees). To ensure statistical stability, the error metrics are evaluated only after the forest reaches a minimal size T_0 (e.g. $T_0 = 50$), guaranteeing that a sufficient proportion of training samples possess valid OOB predictions.

Figure 3.4.3 illustrates the evolution of $RMSE_{OOB}$ (left axis) and the OOB R^2 score (right axis) as functions of the ensemble size T. The error decreases rapidly for small forests, attains its minimum value of 0.398 at $T = 175$ with $R^2 = 0.797$, and exhibits no further meaningful improvement thereafter. Hence, $T^* = 175$ is adopted for the final model configuration.

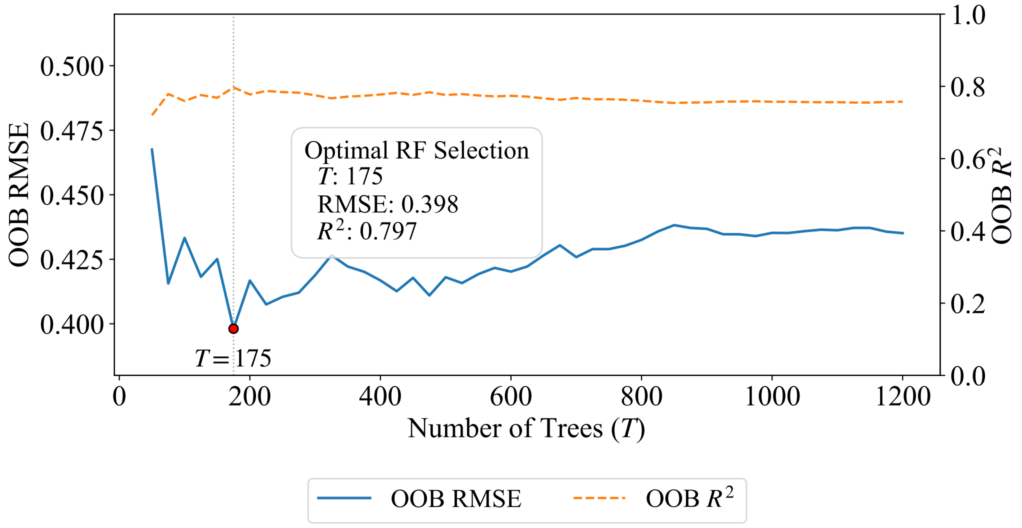

Fig. 3.4.3: OOB error and R^2 as functions of the ensemble size

The final RF model is then retrained on the entire training set using $T = T^*$, and its generalisation performance is formally evaluated on the isolated test set. The prediction accuracy is quantified using the Mean Absolute Error (MAE), Root Mean Square Error

(RMSE), and the Coefficient of Determination (R^2).

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{3.4.3}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{3.4.4}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{3.4.5}$$

where $\bar{y}$ is the mean of the observed targets.

Fig. 3.4.4 compares the predicted frequency nadir and the actual values, with a diagonal reference line $y = x$. Point colours represent the local density of samples, where warmer colours indicate areas of higher concentration. As shown, most observations cluster tightly along the diagonal. This alignment indicates that the model has no significant systematic bias. Furthermore, the histograms along the axes demonstrate that the distribution of predicted values closely matches the width and shape of the actual values.

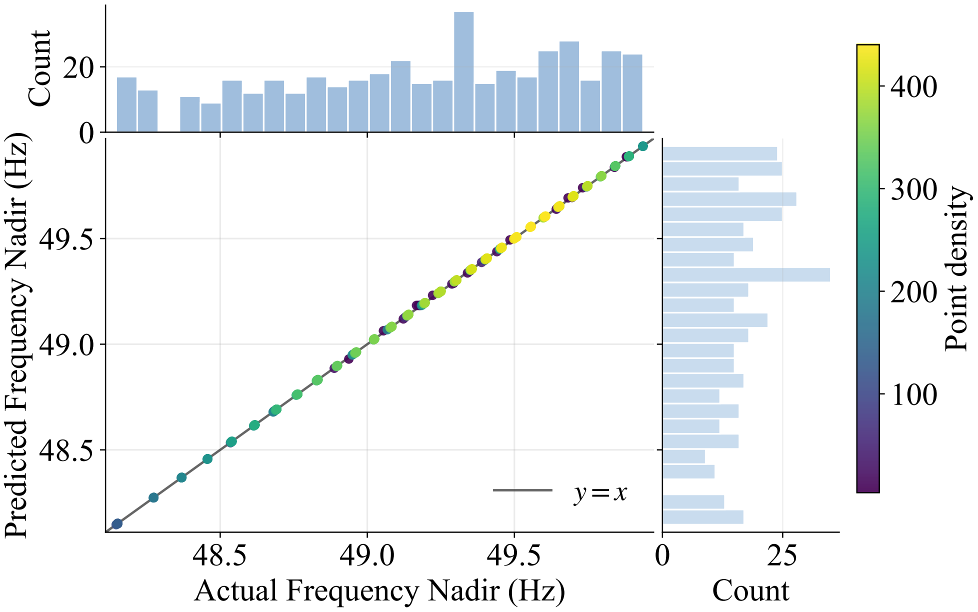

Fig. 3.4.4: Overall model performance for frequency nadir prediction

Validation of Sensitivity-Based Prediction Intervals

To complement the overall accuracy analysis, 10 representative test cases are analysed in detail. Fig. 3.4.5 displays these selected samples, showing the actual frequency nadir (y_{true}, blue dot) alongside the predicted value ($\hat{y}$, orange cross). A vertical error bar represents the sensitivity band derived from the measurement uncertainty in H (at the 95% confidence level). Due to the non-linear structure of the RF model, the intervals are not necessarily symmetric around the point prediction. Their width varies across different operating points, reflect the fluctuating local sensitivity of the nadir to deviations in system inertia.

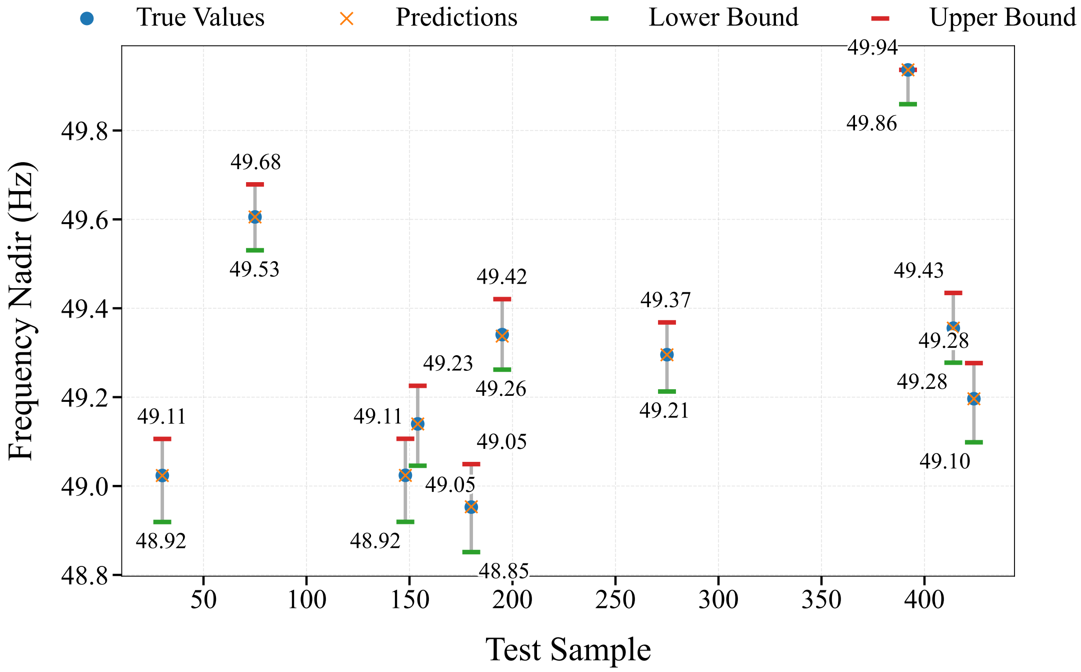

Fig. 3.4.5: Nadir prediction and sensitivity intervals for selected test cases

Table 3.4.4 summarises the results for the ten representative cases. In all cases, the observed nadir is successfully captured within the interval generated by perturbing only H. The interval width ranges from 0.0772 Hz to 0.198 Hz, showing that the influence of inertia is highly dependent on the specific operating condition. Larger interval widths represent high-sensitivity regions where the frequency nadir is unstable, meaning slight variations in the available inertia H can trigger drastically different dynamic outcomes.

Table 3.4.4: Detailed results for the ten selected test samples

$y_{\mathbf{true}}$	$\hat{y}$	**Lower Bound**	**Upper Bound**	**Width**	**Covered**
49.1965	49.1963	49.0984	49.2766	0.1782	Y
49.6055	49.6054	49.5304	49.6787	0.1482	Y
48.9526	48.9533	48.8140	49.0493	0.1979	Y
49.0237	49.0236	48.9191	49.1060	0.1869	Y
49.9363	49.9360	49.8592	49.9364	0.0772	Y
49.2956	49.2953	49.2127	49.3682	0.1555	Y
49.3555	49.3550	49.2776	49.4345	0.1570	Y
49.1398	49.1397	49.0457	49.2256	0.1799	Y
49.3407	49.3373	49.2619	49.4205	0.1586	Y
49.0242	49.0239	48.9194	49.1064	0.1870	Y

Given that system inertia is the primary physical determinant of the frequency response and is subject to significant measurement uncertainty, the constructed sensitivity bands are evaluated here as simplified prediction intervals. Table 3.4.5 presents the statistical performance across the full test set ($n = 432$) and reports an empirical coverage probability of PICP is 0.965. This value slightly exceeds the nominal 0.95 level and remains well within its associated confidence interval, suggesting a robust and slightly conservative estimation of uncertainty.

Table 3.4.5: Quality of the simplified prediction intervals on the test set

Metric	**Value**
Test sample size (n)	432
Target confidence level ($1 - \alpha$)	0.950
Actual coverage (PICP)	0.965
PICP 95% CI	[0.948, 0.982]
Average Width (MPIW)	0.1934 Hz
Underflow rate ($y <$ Lower Bound)	0.019
Overflow rate ($y >$ Upper Bound)	0.016
Expected Uncovered Count	21.6
Actual Uncovered Count	15

The fact that considering uncertainty exclusively in H achieves such high coverage indicates that measurement errors in inertia account for the dominant share of total predictive uncertainty. This implies that other potential error sources, such as model approximation bias or noise in secondary features, are relatively negligible compared to the impact of H. Additionally, the MPIW is 0.1934 Hz, ensuring that the intervals remain narrow enough to be operationally informative. The balanced underflow and overflow rates further confirm that the sensitivity bands are statistically centred and free from systematic bias. Therefore, despite isolating only a single source of uncertainty, these H-based bands demonstrate high efficacy as reliable prediction intervals for real-time power system applications.

Validation of Scheduling of Frequency Response Services

The proposed scheduling scheme is tested across various dynamic conditions. Table 3.4.6 summarises the results for the 10 operating scenarios, comparing the observed frequency nadirs with and without the optimised services. In all scenarios, without mitigation, the nadir falls below the critical threshold of 49.5 Hz. However, by dispatching the optimised P_{dc} and P_{static}, the system reliably maintains the nadir above 49.5 Hz. These results demonstrate that the proposed method effectively secures the system frequency while minimising the procurement quantity, validating both the predictive accuracy and the operational viability of the scheduling scheme.

Two representative contingencies (case 1 and case 5) are detailed in Fig. 3.4.6 and Fig. 3.4.7. To evaluate the proposed approach, the frequency response without additional ancillary services is compared with the optimised response. In both cases, the optimised strategy substantially improves frequency security by mitigating the frequency nadir and accelerating system recovery. Specifically, the frequency nadir is raised above the 49.5 Hz threshold, while the time to nadir is noticeably advanced, indicating a faster arrest of the frequency decline. The baseline response represents the inherent system behaviour from existing synchronous inertia and standard primary frequency control, whereas the optimised response incorporates the pre-event frequency ancillary services developed in this work. Consistent improvements are also observed across the full test set, with a median nadir enhancement of 0.32 Hz. It should be noted that the

Table 3.4.6: Ten operating scenarios: predicted schedules and observed nadir improvements

Case	H	Demand (GW)	Event (GW)	Nadir (Baseline, Hz)	P_{dc} (GW)	P_{static} (GW)	Nadir (With Svc, Hz)
1	**2.00**	**4.65**	**0.90**	**49.114**	**0.72**	**0.63**	**49.517**
2	4.00	4.35	0.60	49.464	0.33	0.29	49.575
3	6.00	4.46	0.71	49.366	0.46	0.40	49.548
4	3.00	4.57	0.82	49.230	0.61	0.54	49.527
5	**2.00**	**4.53**	**0.78**	**49.243**	**0.57**	**0.50**	**49.530**
6	6.00	4.64	0.89	49.159	0.69	0.60	49.508
7	2.00	4.71	0.96	49.041	0.80	0.70	49.509
8	4.00	4.80	1.05	48.971	0.90	0.79	49.511
9	7.00	4.86	1.11	48.903	0.96	0.85	49.518
10	3.00	4.39	0.64	49.436	0.38	0.34	49.563

minor transient fluctuation observed shortly after the event is a measurement artifact. This effect arises from the numerical differentiation of the phase angle ($d\theta/dt$) within the estimation window, where rapid phase shifts at disturbance inception introduce non-fundamental components, causing a momentary spike in the estimated frequency [4]. It's a common phenomenon also observed in real world events [5].

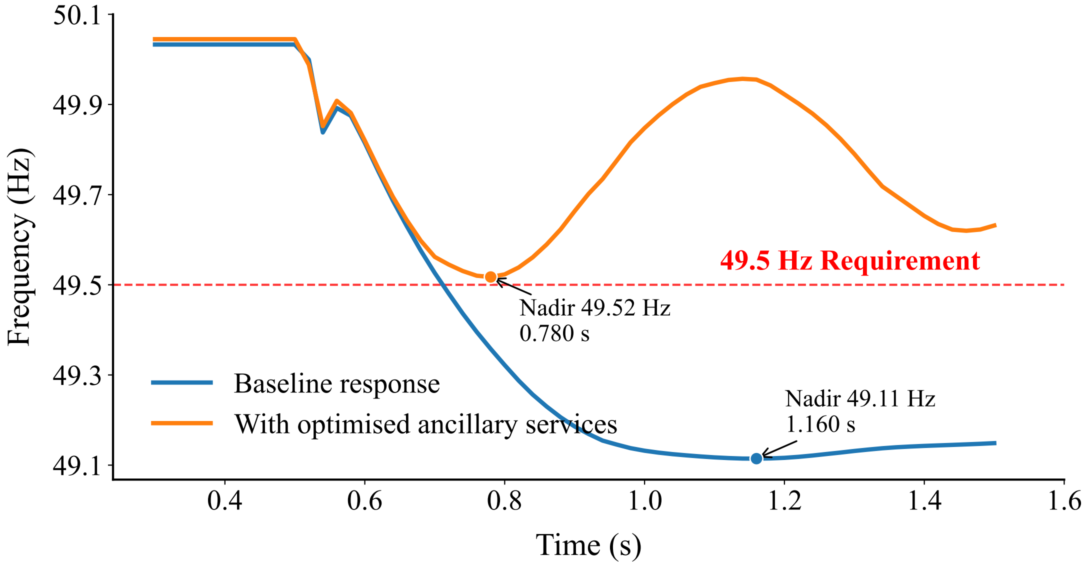

Fig. 3.4.6: Frequency response profile for case 1

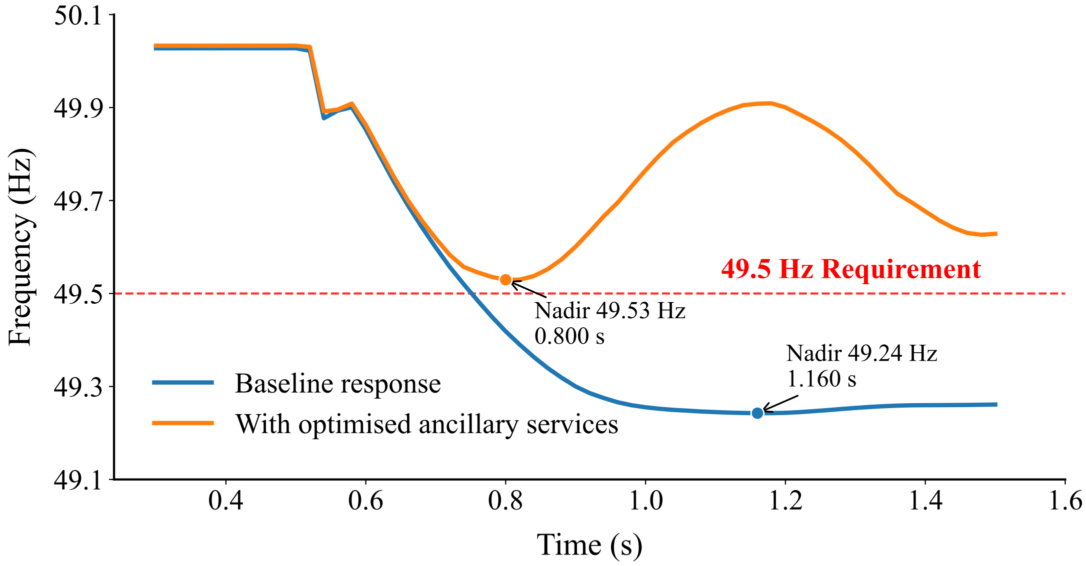

Fig. 3.4.7: Frequency response profile for case 5

3.5 Summary

This chapter has developed and validated an RF-based learning model for predicting the post-event frequency nadir, quantifying associated predictive uncertainty, and optimising the procurement of frequency response services. The approach leverages measurable system features to provide real-time situational awareness, ensuring that the scheduling process is both computationally efficient and security-informed. The RF-based frequency nadir predictor has been evaluated across diverse operating scenarios, which demonstrates high predictive accuracy with a strong alignment between predicted and actual values. This is evidenced by a high overall R^2 and tight sample clustering on unseen data. Feature importance analysis confirms that system equivalent inertia H is the dominant physical driver of the frequency nadir, followed by power imbalance size and demand, which is highly consistent with the theoretical expectations of the swing equation.

Sensitivity-based uncertainty bounds, constructed by propagating a 95% confidence level measurement error in system inertia, achieved a PICP of 0.965 with a MPIW of 0.1934 Hz. By embedding this uncertainty-aware predictor into a linear optimisation algorithm, the approach enables the optimal procurement of DC and static response. Case studies demonstrate that this method successfully maintains the frequency nadir above the 49.5 Hz security threshold in all test scenarios while avoiding the over-procurement typical of conservative, fixed-amount reserves. The algorithm achieved a median nadir improvement of 0.32 Hz, effectively eliminating frequency security excursions and ensuring stability even under the most credible worst-case inertia conditions.

Overall, the proposed security-informed scheduling approach enables system operators to dynamically adapt procurement quantities to the specific requirements of the prevailing operating state. By directly incorporating nadir predictions and their associated uncertainty bounds into the optimisation process, the algorithm provides a robust and efficient solution to safeguard the security of low-inertia power systems.

References

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

[2] L. Breiman, "Bagging predictors," *Machine Learning*, vol. 24, no. 2, pp. 123–140, 1996.

[3] M. Belgiu and L. Drăguţ, "Random forest in remote sensing: A review of applications and future directions," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 114, pp. 24–31, 2016.

[4] European Network of Transmission System Operators for Electricity (ENTSO-E), "Frequency measurement in synchronous areas and its impact on LFDD settings," Technical Report, Sep. 2018. [Online]. Available: `https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf`.

[5] S. Engelken, C. Strafiel, and E. Quitmann, "Frequency measurement for inverter-based frequency control," in *Proc. 15th Int. Wind Integr. Workshop*, 2016, pp. 1–18.

Chapter 4

Real-Time Estimation of RoCoF and Power Imbalance Size

While the frequency nadir prediction approach and the resulting frequency response scheduling scheme presented in Chapter 3 ensure that an appropriate level and mix of ancillary services are procured to contain frequency deviations, the effective delivery of frequency response ultimately depends on the real-time measurement and monitoring of key power system parameters. During contingency events, e.g. a sudden loss of generation or demand, the RoCoF and power imbalance size serve as primary indicators of event severity, which play a central role in frequency control decision making. However, RoCoF is not a directly measured quantity but is obtained as the time derivative of the system frequency. As a result, any measurement error, noise, or latency present in frequency estimation is inherently amplified in the computation of $\mathrm{d}f/\mathrm{d}t$. Consequently, existing approaches for RoCoF measurement and power imbalance estimation are either highly susceptible to measurement noise or suffer from significant delays introduced by filtering, limiting their effectiveness in the critical early post-event time window.

To address this challenge, this chapter presents learning-based approach for accurate real-time estimation of RoCoF and power imbalance. Instead of relying purely on direct numerical differentiation and conventional filtering procedures, ML techniques are capable of calibrating the measured RoCoF to more accurate measurements with appropriate training. Following this principle, a multi-resolution RoCoF estimation approach is developed, which integrates DWT-SVR to achieve reliable early time frequency gradients from noisy measurements. Furthermore, an XGBoost model is employed to directly estimate the power imbalance size from frequency measurement and selected operating conditions, thereby bypassing the need for explicit inertia and RoCoF calculations. These approaches demonstrate how learning-based methods can enhance the observability of fast frequency dynamics and support timely frequency control actions in low-inertia power systems.

4.1 DWT-SVR Approach for RoCoF Estimation

4.1.1 Overview of the RoCoF Estimation Methodology

Traditional RoCoF measurement methods introduce errors and time delays under transient conditions due to the adverse effects of measurement noise and filtering processes. This results in deviations between the measured RoCoF and the actual system RoCoF, especially in the early stages following an event. However, it is difficult to achieve fast response while ensuring high accuracy by simply improving the numerical differentiation method or independently adjusting the filter parameters [1]. To address this technical gap, this section presents a method that employs DWT to extract transient features from measured signals, and utilises SVR algorithm as an alternative approach to learn the relationship between the measured signal and the actual RoCoF from the data to enable a fast and accurate estimation.

Rather than requiring fundamental modifications to measurement algorithms/devices, which could entail hardware changes and may not be practical, the proposed RoCoF estimation approach operates directly on the raw, noisy outputs of existing RoCoF measurement devices, such as PMUs. The methodology focuses on calibrating the output measurement signals to restore their physical accuracy while significantly reducing the time delay associated with traditional low-pass filtering. The overall structure of the proposed approach is illustrated in Fig. 4.1.1. As shown in the figure, the methodology involves four main stages:

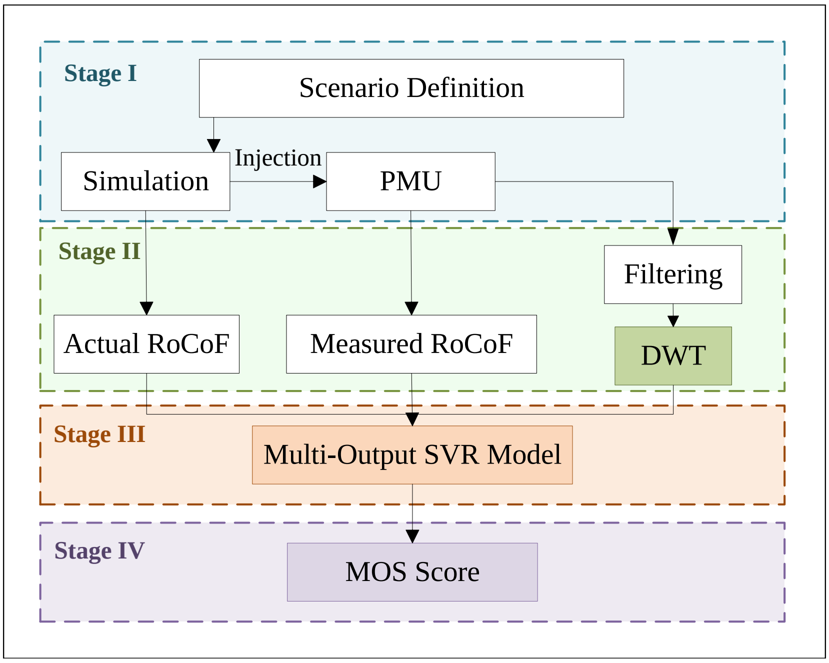

Fig. 4.1.1: Overview flowchart of RoCoF estimation

- **Stage I: Scenarios Generation and PMU Response Characterisation.** In this initial stage, a wide range and representative frequency events are firstly defined, which are used for enabling dynamic simulations to create a diverse set of power system disturbance scenarios. The simulated waveforms are injected to PMUs (or the targeted RoCoF measurement devices), which generates time-aligned pairs of actual RoCoF (the ground truth reference) and measured RoCoF from the PMU.

- **Stage II: Feature Construction.** A preliminary filtering step is applied to the measurements to simulate typical industrial noise-suppression processes. Following, DWT is applied to decompose the filtered RoCoF into multi-resolution components, which captures both fast transients and slower frequency trends. These features with the corresponding measurement data, are consolidated into the training data set.

- **Stage III: Model Training.** A multi-output SVR model is trained using the constructed features. SVR is selected for its robust generalisation capabilities.

The model is specifically optimised to learn the mapping from the multi-resolution features to the actual RoCoF trajectory, effectively compensating for measurement errors introduced in PMUs measurements.

- **Stage IV: Performance Evaluation.** The final stage assesses the accuracy and reliability of the trained models. A composite Model Output Statistics (MOS) score (incorporating bias and variance metrics) is used to evaluate the candidate models, ensuring that the selected estimator provides high precision and minimal latency for real-time applications.

4.1.2 Design and Implementation of DWT-SVR-Based RoCoF Estimation

Stage I: Scenario Generation and PMU Response Characterisation

The main objective of the initial stage is to characterise the measurement response of a RoCoF measurement device, i.e. in this case a PMU, for a diverse range of synthesised signals [1]. Constructing a reliable dataset that contains both measured RoCoF signals and their corresponding true system values is essential for supervised model training. However, extensive and representative RoCoF records are rarely available from field measurements, and more importantly, the actual system RoCoF is generally not directly observable in real world events. The lack of reliable ground-truth reference data makes it difficult to use real measurement data alone for model development and validation. To address these limitations, a synthetic data generation methodology based on offline dynamic simulations is employed to generate a diverse set of frequency disturbance scenarios with known reference RoCoF trajectories.

The detailed process of Stage I is outlined in Fig. 4.1.2, which shows that the stage is structured to construct paired datasets of actual and measured RoCoF. The scenario definition part consists of two categories of training events. First, the training events are formulated by explicitly defining the desired frequency variations. These

---

[1] Please be noted that the proposed method is equally applicable for calibrating RoCoF measurement from other similar devices for enhanced accuracy, but PMU is used as an example in this work.

variations encompass both controlled mathematical signal injections (such as frequency step changes).

For SFR-based scenarios, the defined disturbance parameters are applied as inputs to the analytical SFR model [2]. In the SFR model, mechanical power P_m represents the aggregated power supplied by synchronous generation and frequency control services, while electrical power P_e corresponds to the net system demand following a disturbance. The resulting power imbalance $\Delta P = P_m - P_e$ drives the frequency and RoCoF dynamics through the swing equation, with the equivalent inertia H_s governing the initial frequency acceleration. Primary frequency control acts by adjusting P_m to restore the power balance and arrest the frequency decline.

Despite its reduced complexity, the SFR model preserves the fundamental physics governing system frequency dynamics, as it is explicitly grounded in the swing equation. It captures the essential relationship between aggregated system inertia (H), net power imbalance (ΔP), and frequency response services, providing a physically consistent representation of system-level frequency and RoCoF behaviour. At the same time, the aggregated formulation significantly reduces computational cost, enabling the rapid generation of large and diverse disturbance datasets required for learning-based model training and validation.

Driven by these specific input events, the analytical model simulates the corresponding physical frequency evolution of the system. By varying the disturbance magnitudes and characteristics across different scenarios, a diverse and comprehensive set of frequency disturbances is generated. This process ultimately outputs a precise, analytical ground-truth trajectory, denoted as the actual RoCoF, which serves as the noise-free reference $r_{act}(k)$ required for the training of the SVR-based estimation model.

To bridge the gap between these analytical trajectories and realistic system conditions, the resulting frequency profiles are applied to a three-phase controlled voltage source, a representative dynamic network model, to synthesise the corresponding voltage waveforms (V_a, V_b, V_c). As the three-phase voltage waveforms pass through the controlled network model (comprising transmission lines, transformers, and renewable generation sources), it naturally acquires broadband noise and high-frequency transient components that typically obscure physical behaviour in real world grids. Finally, these network signals are injected into the PMU (e.g. via an analogue amplifier) to obtain the measured RoCoF $r_{pmu}(k)$. Through this complete flow, the approach effectively captures the complex measurement dynamics and creates the paired sets of measured RoCoF and actual RoCoF required for training the subsequent RoCoF estimation model.

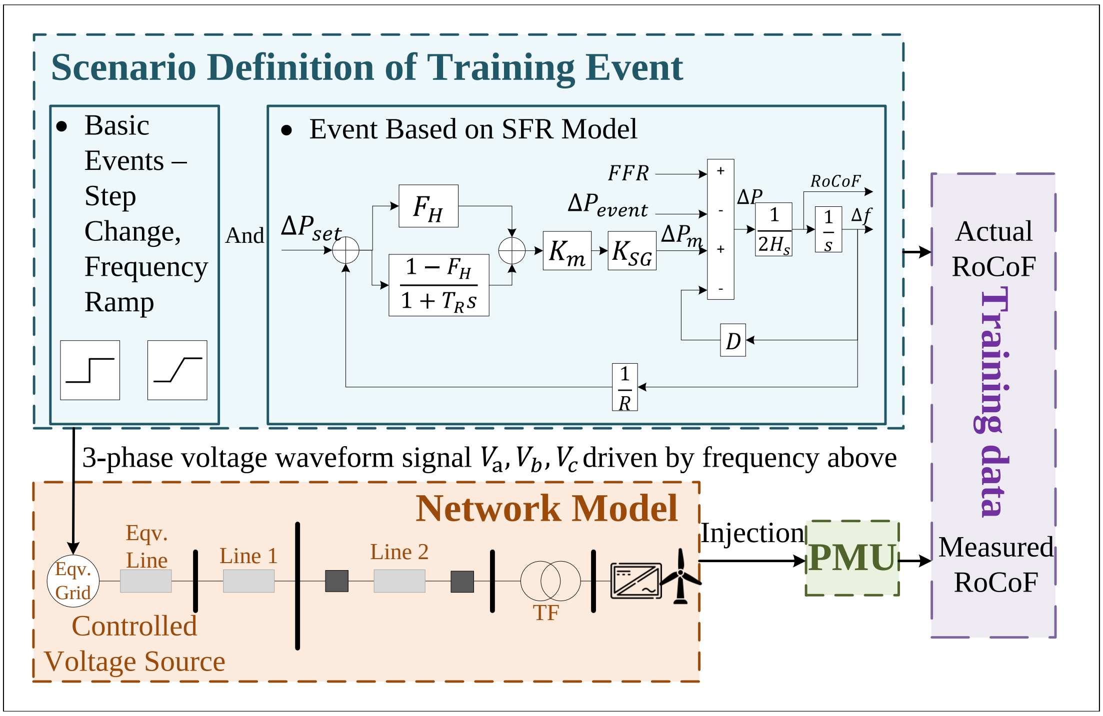

Fig. 4.1.2: Scenario generation and PMU response characterisation for RoCoF estimation (stage I)

Stage II: Feature Construction

PMU-based RoCoF measurements inherently contain substantial noise and transient fluctuations that obscure the underlying system dynamics. To mitigate these effects and isolate critical disturbance characteristics, a multi-stage feature construction is employed, comprising low-pass filtering, DWT denoising, and a progressive windowing scheme.

First, to attenuate broadband noise and provide baseline smoothing, a low-pass filter with a time constant τ ($\tau = 0.1\,\mathrm{s}$ in this work) is applied. In continuous time, the filter satisfies

$$r_{pmu}(t) = \tau \frac{dy(t)}{dt} + y(t) \tag{4.1.1}$$

where $y(t)$ represents the filtered RoCoF signal. With sampling period T_s, its discrete-time realisation is

$$r_{filt}(k) = \alpha r_{filt}(k-1) + \beta r_{pmu}(k) + \beta r_{pmu}(k-1) \tag{4.1.2}$$

with coefficients

$$\alpha = \frac{\tau - 0.5\,T_s}{\tau + 0.5\,T_s} \qquad \beta = \frac{0.5\,T_s}{\tau + 0.5\,T_s} \tag{4.1.3}$$

The filtered sequence $r_{filt}(k)$ retains the dominant transient behaviour while introducing only limited latency, making its causal recursion highly suitable for real-time applications.

For each simulated event, the filtered sequence $r_{filt}(k)$ and the corresponding reference RoCoF $r_{act}(k)$ are segmented. To capture the complete transient dynamics and provide sufficient signal length for the subsequent DWT decomposition, the sequences are restricted to the first $T = 100$ samples (approximately 2 s after disturbance initiation).

While the low-pass filter handles baseline smoothing, further processing is required to extract non-stationary transient disturbances. Traditional techniques (e.g. the Fourier transform) often fail to adequately capture these localised dynamics. To address this, DWT is employed, which provides a multi-resolution time and frequency analysis that has been widely applied in power system protection and transient event analysis [3]–[5]. Due to its multi-scale decomposition capability, DWT is particularly effective at handling RoCoF signals, which inherently contain transient, non-stationary fluctuations.

Specifically, DWT-based decomposition enables the suppression of residual high-frequency noise while retaining transient characteristics that are critical for accurate RoCoF estimation. It achieves this by decomposing the signal into approximation coefficients, which reflect the slow-varying fundamental trend, and multiple detail coefficients, which isolate progressively small-scale transient features [6]. Mathematically, the DWT expands the signal using an orthonormal basis generated by the dyadic scaling and shifting of a mother wavelet $\psi(t)$:

$$\psi_{j,k}(t) = 2^{j/2}\psi(2^j t - k), \quad j, k \in Z \tag{4.1.4}$$

where j and k denote the scale and translation indices, respectively. This dyadic structure eliminates information redundancy and ensures high computational efficiency for real-time analysis.

For the discrete sequence $x[n]$, approximation coefficients $A_j[n]$ and detailed coefficients $D_j[n]$ are obtained by the analysis of low-pass and high-pass filters $h[m]$ and $g[m]$

$$A_j[n] = \sum_m h[m]\ A_{j-1}[2n-m] \qquad A_0[n] = x[n] \tag{4.1.5}$$

$$D_j[n] = \sum_m g[m]\ A_{j-1}[2n-m] \tag{4.1.6}$$

where n represents the discrete time index at each decomposition level.

Decomposition proceeds to level J, which is chosen to capture transient behaviour across multiple resolutions. The approximation vector A_J carries the trend, whereas the detail coefficients $\{D_j\}_{j=1}^{J}$ capture progressively finer disturbance-related structures. To suppress residual noise without discarding disturbance related components, soft threshold is applied to each detail coefficient vector. Coefficients below a magnitude threshold are reduced toward zero, whereas larger coefficients associated with transient dynamics are retained. The full processing procedure is summarised in Algorithm 4.1.1, which produces a multi-resolution feature sequence, denoted as F_{dwt}.

Since the lengths of disturbance windows vary under the progressive windowing scheme, the resulting DWT feature sequences have different sizes. However, the SVR requires input vectors of fixed dimensionality. Therefore, the DWT approximation and detailed components are resized to predefined lengths by truncating longer sequences or zero-padding shorter sequences. The resulting dimension of the final input feature vector z_i is defined as

$$M \ = \ L^*_{A_J} + \sum_{j=1}^{J} L^*_{D_j} \tag{4.1.7}$$

Algorithm 4.1.1 DWT-based denoising and multi-resolution feature extraction for RoCoF estimation

```
Input: Segmented RoCoF sequence $x[n] \in \mathrm{R}^T$; mother wavelet $\psi(t)$; decomposition level $J \leq \lfloor \log_2(T) \rfloor$.
Decomposition: Perform multi-level DWT on $x[n]$ to obtain the approximation coefficients $A_J$ and the set of detail coefficients $\{D_1, \dots, D_J\}$ via Eqs. (4.1.5)–(4.1.6).
for $j = 1$ to $J$ do
    Determine the noise threshold $\lambda_j$ (e.g. using the universal threshold rule).
    Denoising: Apply soft threshold to the detail coefficients $D_j$:
        $\tilde{D}_j \leftarrow \mathrm{sign}(D_j) \cdot \max\big(|D_j| - \lambda_j,\, 0\big)$
end for
Feature Construction: Combine the approximation coefficients and the denoised detail coefficients to form the final feature sequence:
    $F_{dwt} = \big[A_J^\top, \tilde{D}_J^\top, \tilde{D}_{J-1}^\top, \dots, \tilde{D}_1^\top\big]^\top$
Output: Multi-resolution feature sequence $F_{dwt}$.
```

which remains fixed for all windows. Stacking all feature vectors generates the final feature matrix $Z = [z_1, z_2, \dots, z_N]^T \in R^{N \times M}$, where N denotes the total number of training samples. The corresponding target matrix is $Y \in R^{N \times w}$, where w is the length of the same observation window used for the input features.

Stage III: Multi-Output SVR Model Training

Support vector machines implement structural risk minimisation by solving an optimisation problem that simultaneously controls model smoothness and prediction accuracy [7]. They are widely used in classification, regression, and time series analysis [8]. When applied to continuous-valued targets, the formulation becomes SVR, which estimates a regression function of the form

$$f(z) = \theta^T \phi(z) + b \tag{4.1.8}$$

where z is the input feature vector, $\phi(\cdot)$ is the kernel-induced feature mapping, and θ and b are the parameters to be determined during training. SVR introduces an ε-insensitive loss, where estimation errors within a tolerance tube do not influence the learning objective. Deviations outside the tube are captured by slack variables, thereby reducing the impact of measurement noise in PMU data. For a dataset $\{(z_i, y_i)\}_{i=1}^{N}$, the optimisation problem is expressed as

$$\begin{aligned} \min_{\theta,b,\xi,\xi^*} \quad & \frac{1}{2}||\theta||^2 + C\sum_{i=1}^{n}(\xi_i + \xi_i^*) \\ \text{subject to} \quad & \begin{cases} y_i - f(z_i) \le \varepsilon + \xi_i \\ f(z_i) - y_i \le \varepsilon + \xi_i^* \\ \xi_i \ge 0, \xi_i^* \ge 0, \quad i = 1, \dots, n \end{cases} \end{aligned} \tag{4.1.9}$$

where z_i and y_i denote the input RoCoF window and the corresponding actual output value, respectively. The hyperparameter $C > 0$ controls the balance between the model regularisation term and the penalty assigned to errors exceeding the margin ε.

The concept of the ε-tube is illustrated in Fig. 4.1.3. The red curve denotes the learned regression function, while the region enclosed by the dashed boundaries corresponds to the tolerance zone. Samples located within this region do not contribute to the loss function, whereas samples outside the zone influence the cost through the slack terms ξ_i and ξ_i^*, which quantify their deviation from the allowable error margin ε. This mechanism enables the estimator to ignore small noise within the tolerance zone while robustly penalising significant transient outliers.

To meet the requirements of real-time applications and avoid the error propagation associated with recursive methods, a direct multi-step strategy is adopted. In this approach, an independent SVR model is trained to estimate each discrete time step within the observation window. If the current window contains w samples, the estimator becomes an ensemble of w parallel regressors. For an input feature vector $z_i \in R^M$ extracted from the noisy disturbance window, the corresponding noise-free RoCoF sequence is estimated simultaneously as

$$\hat{y}_i = [f^{(1)}(z_i), f^{(2)}(z_i), \dots, f^{(w)}(z_i)]^T \in R^w \tag{4.1.10}$$

Each SVR estimator $f^{(s)}$ corresponds to a specific step s within the current observa-

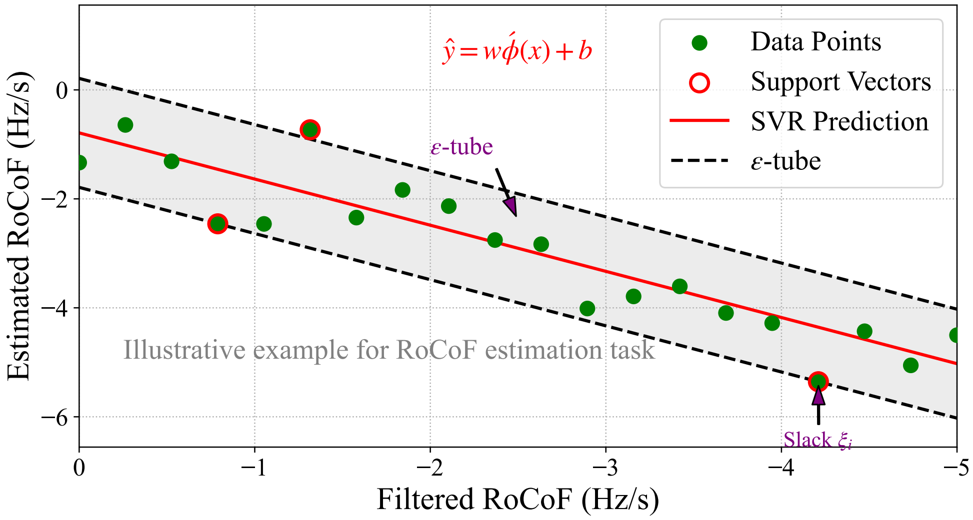

Fig. 4.1.3: ε-insensitive tube and slack variables in SVR

tion window. Consequently, all w regressors function independently and are evaluated in parallel.

Before training the SVR ensemble, the consolidated feature matrix $Z \in R^{N\times M}$ (generated in Stage II) is standardised using statistics derived exclusively from the training set. Standardisation prevents coefficients with large magnitude from dominating the kernel computation and improves the numerical stability of the Radial Basis Function (RBF) kernel. After preprocessing, the set of w SVR regressors is trained, where the s^{th} regressor is fitted using the s^{th} column of the target matrix Y. The hyperparameters $\{C, \varepsilon\}$ are optimised through a logarithmic grid search combined with time-series cross-validation, utilising RMSE as the evaluation metric. The overall training process is summarised in Algorithm 4.1.2.

During online operation, once a disturbance is detected, DWT features are continuously extracted from the current data window of length w, zero-padded to a fixed length M, standardised using the stored (μ, σ) parameters, and evaluated by all w regressors in parallel. The integration of DWT-based multi-resolution feature extraction with an independent multi-output SVR approach ensures both accuracy and low-latency estimation, making the proposed DWT-SVR approach highly suitable for fast frequency response and power system protection applications.

Algorithm 4.1.2 Direct multi-output SVR training procedure

```
Input: Feature matrix $Z \in R^{N\times M}$, target matrix $Y \in R^{N\times w}$.
Preprocessing: Standardise $Z$ using training set statistics, store the mean vector $\mu$ and standard deviation vector $\sigma$.
Model Selection: Select optimal $\{C, \varepsilon\}$ via logarithmic grid search with cross-validation (Objective: minimise RMSE).
for $s = 1$ to $w$ do
    Extract the $s^{th}$ target column $y^{(s)} = Y_{:,s}$.
    Fit SVR estimator $f^{(s)}$ on the pair $\left(Z,\, y^{(s)}\right)$ using the selected hyperparameters.
end for
Output: The ensemble of estimators $\{f^{(s)}\}_{s=1}^{w}$ and scaling parameters $(\mu, \sigma)$.
```

Stage IV: Performance Evaluation

This stage focuses on evaluating the estimation accuracy of the DWT-SVR estimator using standard performance metrics and identifying an optimal observation window size w through a statistical window-selection scheme. To achieve this, a statistical approach adapted from MOS is developed. Unlike classical MOS, which corrects model outputs directly, the proposed MOS-based method adjusts the model structure by selecting the most appropriate window size, thereby mitigating the inherent conflict between estimation noise and detection latency.

MOS Variant for Window Selection: Numerical models fundamentally display systematic biases derived from their structural parameters. In numerical weather prediction, MOS is widely used to correct such biases by building statistical relationships between model outputs and historical observations [9], [10]. For RoCoF estimation, the main structural bias arises from the selected window length w. Short windows amplify noise despite offering a fast response, while long windows smooth the estimates but introduce significant delay. This conflict between variance and latency motivates a learning-based approach to selecting an optimal window size.

Inspired by the MOS bias-correction principle, a modified approach is developed to choose the window length w^* that offers the best balance between stability and accuracy. This is achieved by analysing a set of statistical descriptors extracted from the estimated trajectories. For each candidate window size w, the following metrics are computed to quantify estimation quality:

1. Trajectory Smoothness ($\mathcal{S}_{sm}$): This descriptor measures the average stability of the predicted RoCoF trajectory by computing the mean absolute difference between successive estimates. Lower values indicate a smoother trajectory with reduced high-frequency fluctuations.

$$\mathcal{S}_{sm}(w) = \frac{1}{n-1} \sum_{k=2}^{n} |\hat{y}_k(w) - \hat{y}_{k-1}(w)| \tag{4.1.11}$$

2. Fluctuation Volatility ($\mathcal{S}_{vol}$): This descriptor penalises irregular variations typically arising from noise amplification, especially in short windows. It is defined as the standard deviation of successive differences.

$$\mathcal{S}_{vol}(w) = \mathrm{std}\left(\hat{y}_k(w) - \hat{y}_{k-1}(w)\right) \tag{4.1.12}$$

3. Estimation Accuracy ($\mathcal{S}_{acc}$): While stability is crucial, the estimate must remain accurate relative to the actual system dynamics. This metric calculates the local RMSE using a set of known true values to ensure the model closely tracks the disturbance.

$$\mathcal{S}_{acc}(w) = \sqrt{\frac{1}{n} \sum_{k=1}^{n} \big(\hat{y}_k(w) - y_k\big)^2} \tag{4.1.13}$$

Composite Score and Weight Optimisation: To objectively combine these diverse descriptors into a single selection criterion, a composite scoring function is formulated. To ensure that different metrics contribute equitably regardless of their absolute magnitudes, each statistic $X(w)$ is normalised to the range $[0, 1]$ using the min-max scaling

$$X_{norm}(w) = \frac{X(w) - \min_w X(w)}{\max_w X(w) - \min_w X(w)} \tag{4.1.14}$$

The composite score for a window w is then defined as the weighted sum of these normalised metrics

$$Score(w, \lambda) = \lambda_1 \mathcal{S}_{sm}^{norm}(w) + \lambda_2 \mathcal{S}_{vol}^{norm}(w) + \lambda_3 \mathcal{S}_{acc}^{norm}(w) \tag{4.1.15}$$

subject to the constraint $\sum_{i=1}^{3} \lambda_i = 1$ and $\lambda_i \geq 0$.

Determining the weight vector $\lambda = [\lambda_1, \lambda_2, \lambda_3]$ manually can introduce subjective bias. Moreover, directly minimising the composite score is unreliable. Because the results may be dominated by performance indicators with different numerical ranges, rather than reflecting truly better estimation behaviour. For a given λ, the window size selected by the score is assessed based on the actual estimation error achieved on an independent validation dataset. The optimal configuration λ^* is therefore defined as the one that selects a window w^* with the lowest validation error.

The optimisation procedure performs a grid search over the feasible range of λ. For each weight combination:

1. **Window Selection (Inner Level):** The observation window is selected by minimising the composite score under the current weight configuration

$$w^*(\lambda) = \arg\min_{w} \mathrm{Score}(w; \lambda) \tag{4.1.16}$$

2. **Performance Validation (Outer Level):** The estimation accuracy of this selected window $w^*(\lambda)$ is evaluated by comparing its estimation error with the RMSE calculated using the validation dataset.

Among all tested weight configurations, the weight vector λ^* that produces the minimum validation RMSE is selected, and the corresponding window size is adopted as the final optimal window w^*. To quantify the overall predictive performance, standard error metrics including MAE, R^2, and RMSE are utilised following the formulations established in Section 3.4.2.

4.1.3 Case Studies

This section evaluates the performance proposed DWT-SVR RoCoF estimation approach and demonstrate its further application through three case studies, as illustrated in the test setup in Fig. 4.1.4 and described as follows.

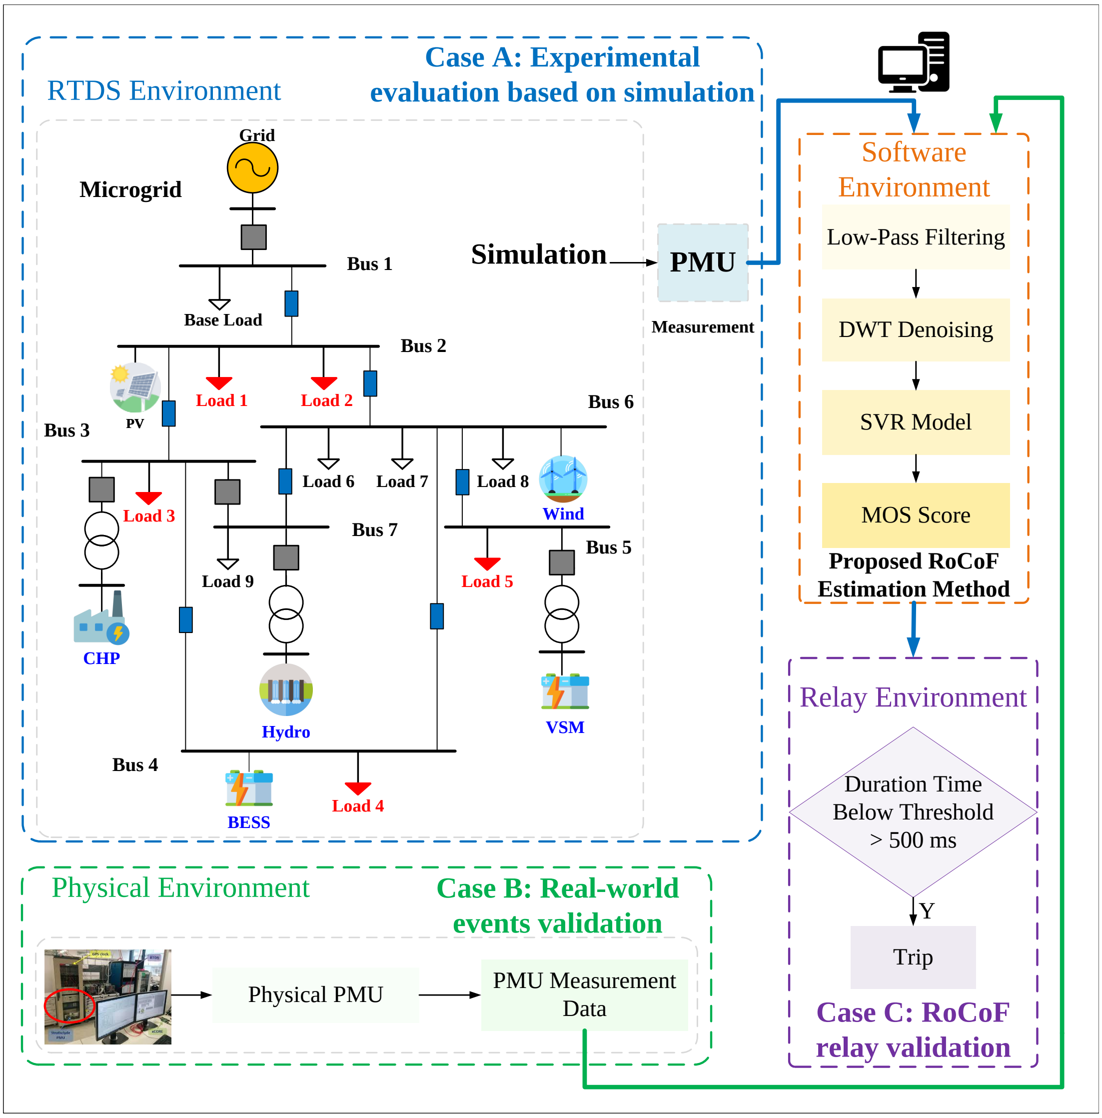

Fig. 4.1.4: Test setup for RoCoF estimation using DWT-SVR

- **Case A: Experimental Evaluation Based on Simulation:** This stage utilises simulated data from microgrid and PMU models in RTDS to benchmark the proposed method with conventional low-pass filtering. These simulations are designed to demonstrate the theoretical feasibility of the approach, allowing for a detailed sensitivity analysis of the window size w.

- **Case B: Real World Event Validation:** This stage utilises measurements acquired from a physical PMU connected to the actual GB power system from real world frequency events. By comparing the estimates with official ground-truth data provided by NESO, this estimation validates the practical feasibility of the proposed method. This confirms its effectiveness in addressing hardware noise and achieving accurate measurements in real world scenarios.
- **Case C: Application of the proposed RoCoF Estimation in a LoM Relay:** This case demonstrates how the proposed RoCoF estimator can not only enable enhanced monitoring, but also protection actions. The output RoCoF from the proposed method is integrated with a LoM relay logic to enable a more robust RoCoF protection for DERs. This showcases its ability to satisfy the requirement for critical operational applications.

The proposed method first applies a low-pass filter to the input data to suppress noises and measurement spikes. Subsequently, a DWT is utilised to extract multi-scale features, based on which the SVR model maps to the target RoCoF values. The sequence employs the MOS mechanism to select the optimal window w^*, generating the final RoCoF estimation.

Case A: Experimental Evaluation of the Proposed RoCoF Estimation Method Based on Simulation

In this case study, a wide rang of disturbances were simulated using the RTDS environment as illustrated in Fig. 4.1.4. As an initial step of the evaluation, a sensitivity study was conducted to verify the stability of the estimator across different temporal scales. Fig. 4.1.5 illustrates the predictive performance across the overall test dataset as the window size w varies.

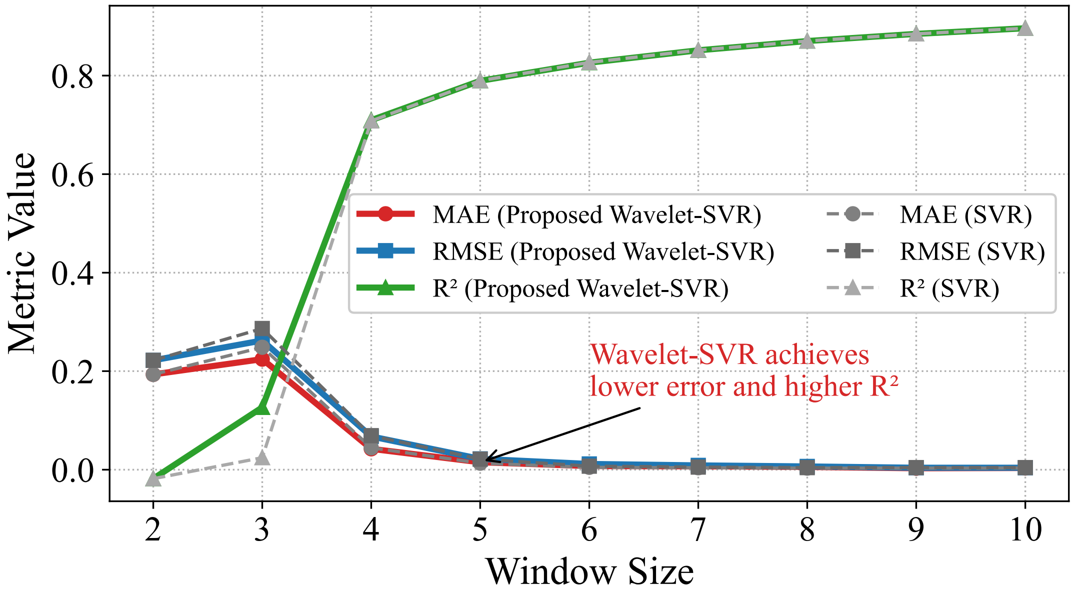

Fig. 4.1.5: Overall test set performance of SVR and DWT-SVR methods across window sizes

As the window size increases from 2 to 4, the estimation error reduces markedly, while R^2 rises above 0.85 and stabilises for larger window lengths. This suggests that the model is able to capture the underlying dynamics more effectively when provided with extended temporal information. These results indicate the importance of selecting a window that offers a suitable balance between noise suppression and dynamic responsiveness, thereby ensuring reliable RoCoF prediction in real-time applications.

To provide a more comprehensive comparison between the various estimation methods across different window configurations, Fig. 4.1.6 illustrates the influence of the input window size w on the estimated RoCoF trajectories. While the estimated results are time-aligned with the ideal RoCoF to facilitate a direct comparison of numerical accuracy, an inherent physical latency exists in real-time implementation. Specifically, since the estimator requires a sequence of w samples, the output is only generated after the full window is captured, introducing a delay of approximately $(w - 1) \times 20$ ms at a 50 Hz reporting rate.

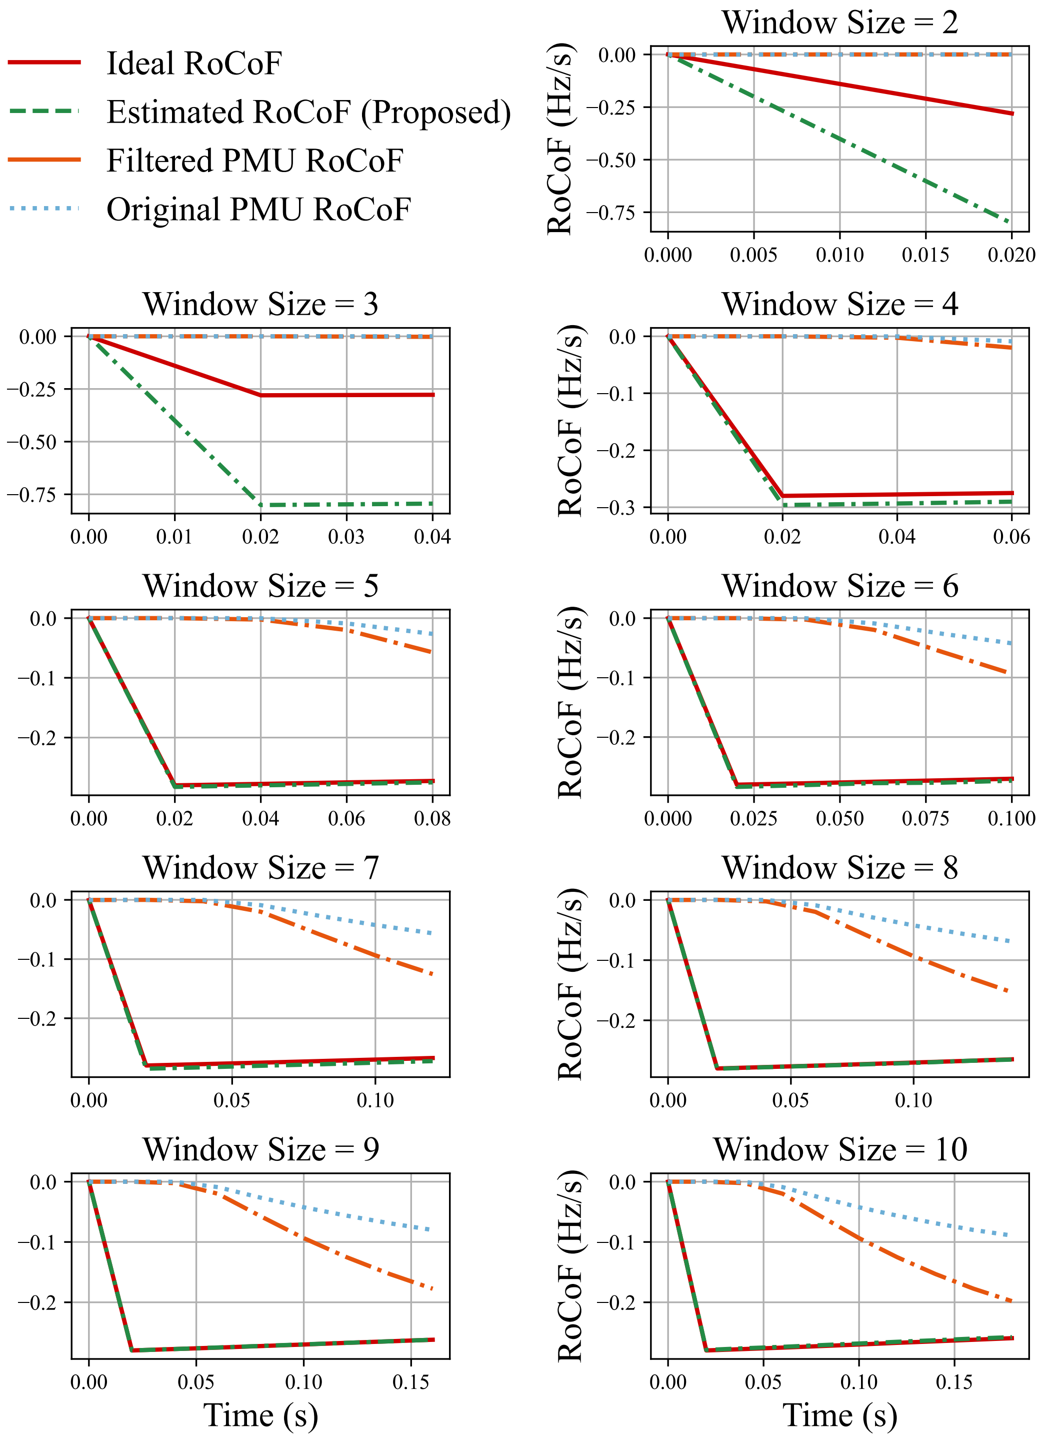

Fig. 4.1.6: Comparison of RoCoF estimates for different window sizes

The results highlight a balance between stability and latency: small window sizes (e.g. $w = 3$) are highly susceptible to measurement noise due to insufficient temporal context, leading to the instability shown in the plots. Conversely, increasing w enhances accuracy and smoothness but introduces greater causal delay, as the estimated value reflects the trend over the preceding window rather than the instantaneous state. As depicted, the proposed method achieves an optimal balance at $w = 4 \sim 6$, providing a rapid response that significantly mitigates the lag typically found in traditional low-pass filtering methods.

Table 4.1.1 provides a detailed comparison of the proposed DWT-SVR method with the filtered PMU and a standard SVR model, using the window selected by the MOS. The filtered PMU method reduces noise but still incurs substantial error and relatively low R^2 due to its inability to capture rapid dynamics.

Table 4.1.1: Accuracy evaluation of filtered, SVR and DWT-SVR RoCoF estimation methods

No.	**Actual RoCoF**	**Filtered Method**		**SVR Method**		**Proposed Method**	
	(Hz/s)	**MAE**	**RMSE**	**MAE**	**RMSE**	**MAE**	**RMSE**
1	-1.165	0.908	1.015	0.187	0.216	0.031	0.035
2	-1.500	1.167	1.305	0.058	0.067	0.001	0.001
3	-1.040	0.809	0.374	0.043	0.048	0.004	0.037
4	-0.435	0.339	0.379	0.017	0.019	0.004	0.005
5	-0.930	0.632	0.641	0.015	0.024	0.004	0.006
6	-0.555	0.432	0.483	0.023	0.027	0.019	0.021
7	-0.640	0.499	0.558	0.078	0.090	0.012	0.014
8	0.127	0.196	0.224	0.016	0.016	0.006	0.008
9	-0.460	0.358	0.401	0.017	0.019	0.011	0.012
10	-0.350	0.272	0.305	0.004	0.004	0.007	0.007

The standard SVR model improves estimation accuracy by learning non-linear relationships from the data, it remains limited in capturing transient behaviour under non-stationary conditions. In contrast, the proposed DWT-SVR model consistently outperforms both baselines across all metrics. By leveraging multi-scale time-frequency features extracted via DWT, the proposed method achieves the lowest MAE and RMSE (below 0.04 Hz/s and 0.1 Hz/s, respectively) and the highest R^2. These improvements demonstrate that integrating DWT-based features with SVR learning enables accurate representation of both steady-state and transient RoCoF characteristics.

In summary, the results confirm that appropriate window selection is essential for stable RoCoF estimation, and the proposed DWT-SVR approach provides consistently reliable predictions across the test scenarios.

Case B: Performance Assessment using Real World Events

To comprehensively assess the effectiveness of the proposed DWT-SVR RoCoF estimation approach, three real world power system disturbance events in the GB were analysed. The experimental data was acquired using a Schweitzer Engineering Laboratories (SEL) 451 PMU [11], which is connected to into the low voltage distribution network and installed at the Dynamic Power System Laboratory at the University of Strathclyde, Glasgow, as illustrated in Fig. 4.1.7. This device fully supports the IEEE C37.118 standard for synchrophasor data [12]. By using this physical hardware, the capability of the proposed method for RoCoF estimation in real world events is evaluated.

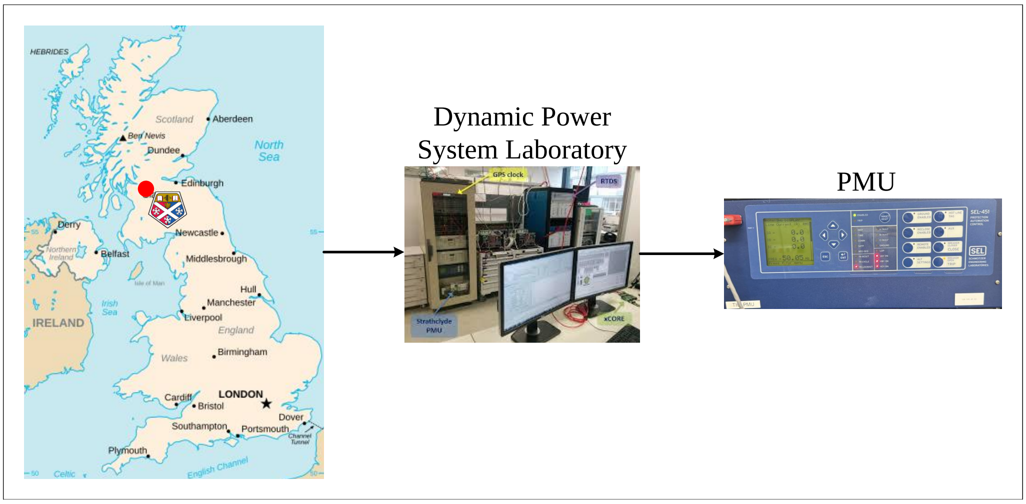

Fig. 4.1.7: PMU in dynamic power system laboratory

For each event, a consistent assessment procedure was followed: (1) DWT-based feature extraction, (2) MOS-driven window selection, (3) RoCoF estimation using the proposed method, and (4) comparison with reference values. This subsection focuses on the first of the three events, as it is representative of a moderate generation loss scenario. The remaining events are summarised later to highlight broader trends and validate generalisability.

Event 1 on 22 December 2023

The first event occurred on 22 December 2023 and involved the sudden disconnection of a large generation unit, resulting in a rapid system frequency decline [13]. Using inertia data provided by the ESO, the theoretical RoCoF was derived through the swing equation. The event corresponds to an interconnector trip with a frequency deviation of -0.145 Hz/s, arising from a 1000 MW import with a system inertia of 172 GVA $\cdot$ s.

Fig. 4.1.8 illustrates the feature extraction stage applied to the recorded PMU measurements. Fig. 4.1.8 (a) presents the original PMU-measured RoCoF signal, which is dominated by broadband noise. Fig. 4.1.8 (b) shows the corresponding low-pass filtered signal, where high-frequency noise are significantly attenuated. The signal is decomposed using a three-level DWT ($J = 3$) with the Daubechies-4 (db4) mother wavelet, generating one approximation vector A_3 and three detail vectors D_3, D_2, D_1, where the resulting approximation coefficients A_3 are displayed in Fig. 4.1.8 (c). These coefficients represent the dominant low-frequency behaviour and serve as the actual SVR model input features.

In this study, only the A_3 coefficients are selected as input features for the SVR model. This selection effectively captures dominant underlying dynamics of the system while suppressing the high-frequency noise and transient disturbances. The detailed coefficients (D_1, D_2, and D_3) are deliberately excluded, as these typically represent noise and short transient disturbances, which would otherwise reduce model accuracy. Specifically, the detail coefficients carry limited useful information regarding the actual RoCoF trend but increase data variance. Thus, the selection of only approximation coefficients simplifies the learning process and improves SVR estimation accuracy.

If the DWT were applied directly to the unfiltered PMU (shown previously in Fig. 4.1.8 (a)), broadband noise and impulsive spikes would contaminate the resulting

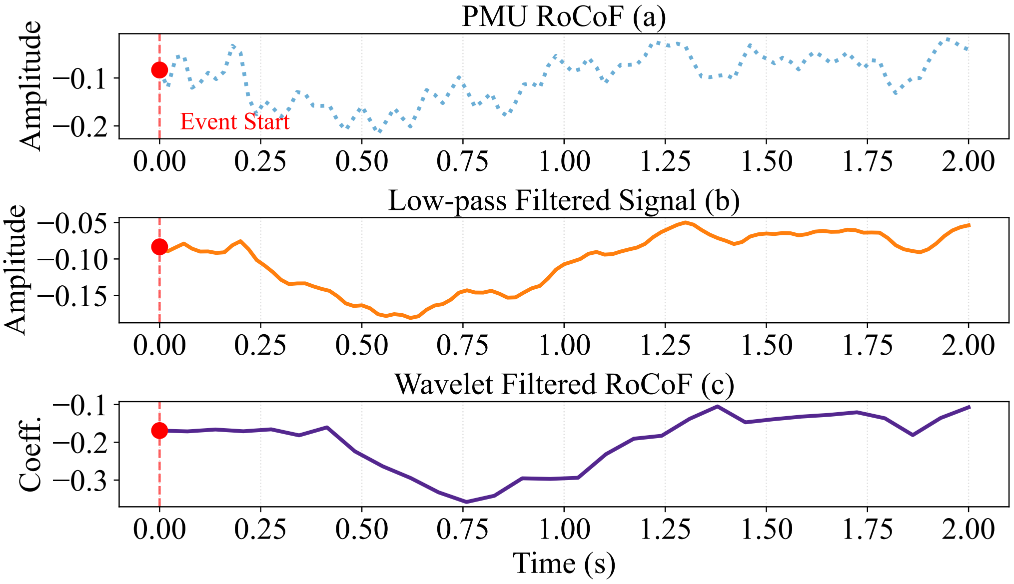

Fig. 4.1.8: DWT decomposition results used for SVR feature generation (case 1)

approximation level. This contamination would increase the complexity of the regression problem and degrade the SVR's predictive performance. The preliminary low-pass filtering step removes slow drifts and stabilises the input, ensuring that the retained A_3 coefficients capture only the relevant low-frequency structure of the RoCoF response.

With the reporting rate of 20 ms ($f_s = 50\,\mathrm{Hz}$) and three dyadic decomposition levels. This configuration produces a passband restricted to approximately

$$\frac{f_s}{2^{3+1}} = \frac{50}{16} \approx 3.1\,\mathrm{Hz}$$

This low-frequency range aligns with the dominant dynamics typically examined in power system stability analysis. Because each DWT decomposition level halves the data length, the coefficients A_3 represent a dimensionally reduced data set of only 1/8 of the original length. This reduction lead to approximately 87.5% fewer data points, which considerably simplifies SVR training without compromising the representation of slow RoCoF dynamics behaviour.

The feature extraction procedure therefore generates a clean and compact representation of the RoCoF trajectory, capturing the dominant low-frequency behaviour while discarding components that do not contribute meaningfully to the underlying dynamics. With these refined inputs, the next step is to determine an appropriate window size that balances noise suppression with responsiveness. All candidate window lengths were evaluated using the composite score defined in Eq. (4.1.15), and the three best performance configurations are summarised in Table 4.1.2.

Table 4.1.2: Optimal window selection for final RoCoF estimation (case 1)

No.	**Model**	**MAE (Hz/s)**	**RMSE (Hz/s)**	**Std (Hz/s)**	**Std Diff (Hz/s)**	**Min (Hz/s)**
1	SVR	0.0021	0.0027	0.0357	0.0053	-0.1226
2		0.0082	0.0086	0.0367	0.0030	-0.1321
3		0.0037	0.0039	0.0373	0.0041	-0.1461
1	DWT-SVR	0.0124	0.0133	0.0434	0.0097	-0.1593
2		0.0133	0.0142	0.0410	0.0106	-0.1428
3		0.0370	0.0394	0.0383	0.0086	-0.1239

Table 4.1.3 compares the RoCoF estimates produced by the filtered PMU method, the standard SVR model, and the proposed DWT-SVR estimator with the calculated reference value. The proposed method achieves the smallest deviation and consistently yields the most accurate RoCoF estimates.

Table 4.1.3: RoCoF estimates compared with the calculated reference values (case 1)

Calculated (Hz/s)	**PMU (Hz/s)**	**Filt. PMU (Hz/s)**	**SVR est. (Hz/s)**	**Proposed est. (Hz/s)**
-0.145	-0.216	-0.181	-0.134	-0.142
Rel. error vs. Ref. (%)	52.75	27.32	5.63	2.51

Fig. 4.1.9 presents the RoCoF trajectories obtained from the four processing methods considered in this study. The PMU-derived RoCoF exhibits substantial high-frequency fluctuations that mask the underlying system behaviour. Applying a conventional low-pass filter reduces these fluctuations, although part of the dynamic shape is simultaneously smoothed out. The standard SVR estimator further suppresses noise but tends to respond more slowly during periods of rapid change. In contrast, the proposed

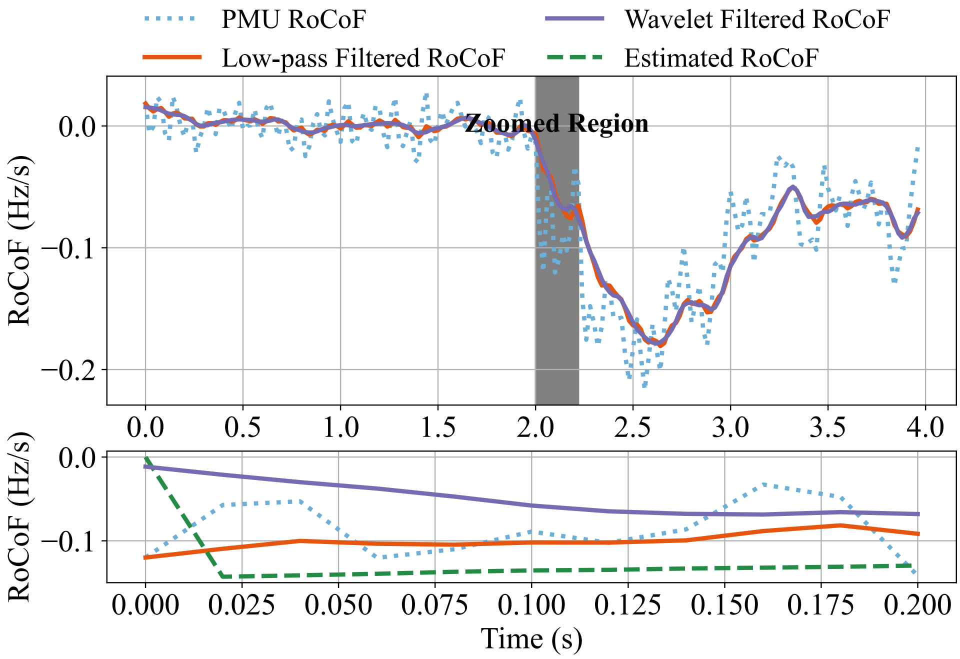

Fig. 4.1.9: RoCoF estimation obtained from different processing methods (case 1)

DWT-SVR approach appeared to be able to estimate a more realistic evaluation of RoCoF values while maintaining robustness to measurement noise. The estimated trajectory closely follows the calculated reference from the onset of the disturbance, capturing the initial frequency decline without inducing oscillations or noticeable delay. The zoomed-in portion of the figure highlights these differences, showing that the proposed method is better able to represent fast dynamics while retaining stability in the presence of noise.

Event 2 on 22 December 2023

A second validation scenario evaluates another generation loss event that occurred on 22 December 2023 and resulted in a measurable frequency deviation [13]. In contrast with the previous case, this event involved a smaller active power imbalance (approximately 440 MW), while the system inertia remained the same (172 GVA · s). This configuration provides a complementary test scenario in which the disturbance magnitude is reduced but the underlying system conditions are unchanged, enabling an assessment of how the proposed method performs under a different disturbance magnitude.

According to the swing equation, the theoretical maximum RoCoF during this event is approximately -0.064 Hz/s. The PMU-derived RoCoF signal was processed using the same feature extraction process as in the previous case. Fig. 4.1.10 shows the corresponding three-level db4 DWT results, including the approximation coefficients used to construct the SVR model input features.

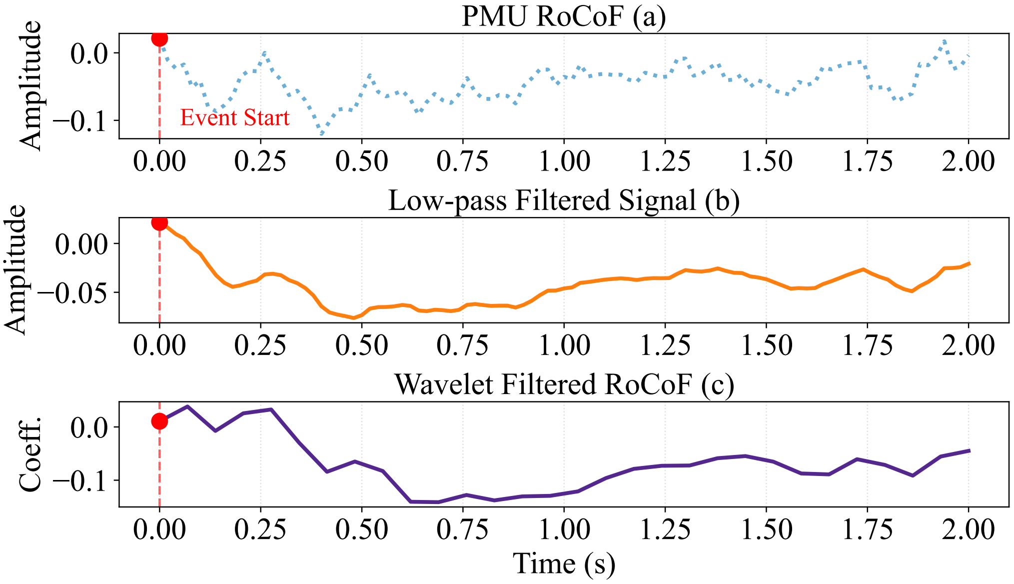

Fig. 4.1.10: DWT decomposition results for SVR feature generation (case 2)

Following feature extraction, all candidate windows were evaluated using a composite score as defined in Eq. (4.1.15). Table 4.1.4 summarises the three best performance window sizes and their accuracy metrics.

Quantitative comparisons in Table 4.1.5 indicate that the proposed DWT-SVR method achieves superior accuracy, closely matching the calculated reference value. The results demonstrate significantly lower relative errors compared to other methods, presenting the accuracy of the proposed algorithm under conditions of increased imbalance and reduced inertia.

Table 4.1.4: Optimal window selection for final RoCoF estimation (case 2)

No.	**Model**	**MAE (Hz/s)**	**RMSE (Hz/s)**	**Std (Hz/s)**	**Std Diff (Hz/s)**	**Min (Hz/s)**
1	SVR	0.0022	0.0024	0.0184	0.0069	-0.0695
2		0.0004	0.0005	0.0151	0.0102	-0.0647
3		0.0004	0.0005	0.0156	0.0099	-0.0652
1	DWT–SVR	0.0032	0.0033	0.0177	0.0076	-0.0659
2		0.0015	0.0016	0.0152	0.0100	-0.0657
3		0.0015	0.0015	0.0153	0.0102	-0.0639

Table 4.1.5: RoCoF estimates compared with the calculated reference values (case 2)

Calculated (Hz/s)	**PMU (Hz/s)**	**Filt. PMU (Hz/s)**	**SVR est. (Hz/s)**	**Proposed est. (Hz/s)**
-0.06395	-0.120	-0.0765	-0.067	-0.065
Rel. error vs. Ref. (%)	84.62	17.69	4.77	1.64

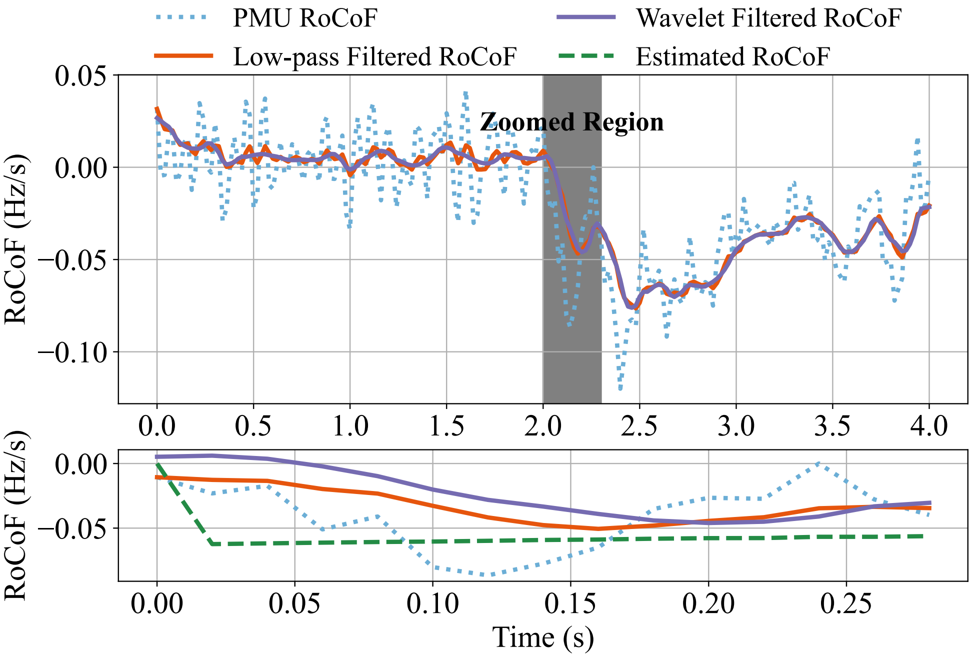

Fig. 4.1.11: RoCoF estimates obtained using different methods (case 2)

Fig. 4.1.11 shows the RoCoF trajectories produced by the different estimation methods. As in the previous case, the PMU-derived RoCoF exhibits substantial measurement noise, and the low-pass filtered approach smooths these fluctuations at the expense of attenuating the dynamic response and introducing time delay. The DWT-filtered approach reduces noise further but responds more gradually to sudden deviations. The proposed DWT-SVR method provides a closer match to the calculated reference across the entire time window, including the initial stage of the disturbance where the RoCoF changes most rapidly. The result indicates that the method remains effective under reduced inertia conditions and varying disturbance characteristics.

Event on 14 March 2025

The final validation scenario assess a generation loss event that occurred on 14 March 2023. In this case, a large generating unit of approximately 654.94 MW tripped, causing a distinct frequency deviation [14]. According to the swing equation and considering the ESO reported inertia of 276.4 GVA $\cdot$ s, the corresponding theoretical maximum RoCoF is estimated to be around -0.059 Hz/s.

The recorded RoCoF signal for this event was similarly decomposed using the three-level db4 DWT. The resulting DWT decomposition, presented in Fig. 4.1.12, illustrate the characteristics of this event.

Following feature extraction, candidate window lengths were assessed using the composite score defined in Eq. (4.1.15). The three best-performing configurations for both SVR and DWT-SVR models are listed in Table 4.1.6 with the associated evaluation metrics.

Quantitative results for this event are summarised in Table 4.1.7. The proposed DWT-SVR estimate shows the closest agreement with the calculated reference, with smaller deviations than those obtained using the filtered PMU signal or the standard SVR model. This outcome is consistent with the trends observed in the earlier cases.

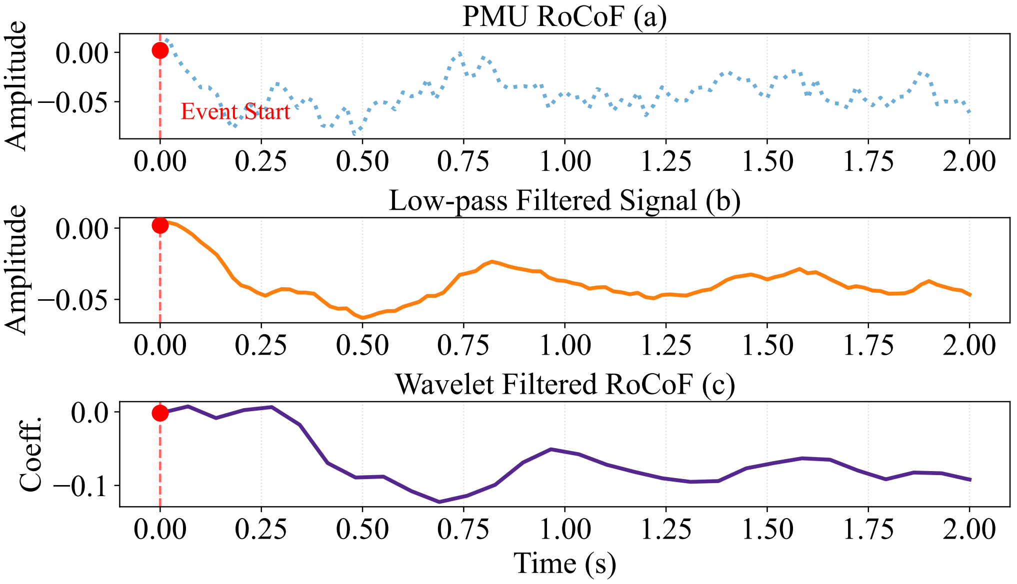

Fig. 4.1.12: DWT decomposition results for SVR feature generation (case 3)

Table 4.1.6: Optimal window selection for final RoCoF estimation (case 3)

No.	**Model**	**MAE (Hz/s)**	**RMSE (Hz/s)**	**Std (Hz/s)**	**Std Diff (Hz/s)**	**Min (Hz/s)**
1	SVR	0.0053	0.0006	0.0142	0.0093	-0.0608
2		0.0004	0.0005	0.0146	0.0096	-0.0608
3		0.0005	0.0007	0.0152	0.0098	-0.0615
1	DWT–SVR	0.0008	0.0009	0.0141	0.0094	-0.0609
2		0.0004	0.0005	0.0137	0.0094	-0.0607
3		0.0001	0.0002	0.0143	0.0099	-0.0598

Table 4.1.7: RoCoF estimates compared with the calculated reference values (case 3)

Calculated (Hz/s)	**PMU (Hz/s)**	**Filt. PMU (Hz/s)**	**SVR est. (Hz/s)**	**Proposed est. (Hz/s)**
-0.059	-0.0828	-0.105	-0.061	-0.0605
Rel. error vs. Ref. (%)	40.34	4.46	3.39	2.54

Fig. 4.1.13 presents the proposed DWT-SVR method maintains a closer alignment with the theoretical trajectory throughout the event.

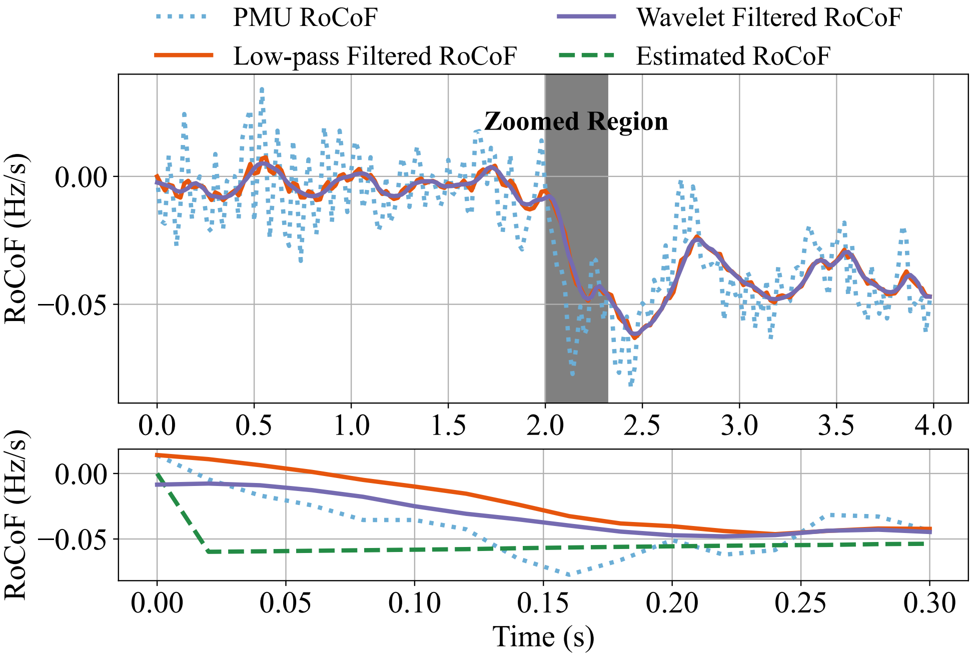

Fig. 4.1.13: RoCoF estimates obtained using different methods (case 3)

Case C: Performance of RoCoF Relay Enabled by the Proposed RoCoF Estimator

Traditional RoCoF relays are widely deployed for LoM protection [15]. Their operation relies significantly on the quality of the underlying RoCoF estimate, as measurement noise and transient fluctuations may cause premature activation or delayed response, particularly in systems with high renewable penetration. A typical relay setting uses a 500 ms measurement window with a trip threshold of −1 Hz/s, requiring the estimated RoCoF to remain below the threshold continuously for the full duration before a trip is issued. This case study will showcase how the proposed DWT-SVR could be applied in RoCoF relays to enable more dependable and secure protection against LoM events.

To evaluate the capability of the proposed method in a RoCoF relay, the integrated test setup previously illustrated in Fig. 4.1.4 is utilised. In this configuration, an event is simulated within the microgrid model in the RTDS environment. The generated data is measured by the PMU, and the raw measurements are input into the software environment. The proposed RoCoF estimation method processes this data, replacing the low-pass filtered RoCoF estimation in traditional RoCoF relays. Finally, this estimate is fed into the relay environment to determine whether a trip command should be issued based on the predefined threshold and persistence criteria.

Fig. 4.1.14 compares the RoCoF trajectories obtained from the filtered PMU method and the proposed estimation approach when evaluated against a relay threshold of -1 Hz/s. The proposed estimate enters the trip region earlier and exits it earlier as well, producing a continuous below threshold interval that satisfies the typical 500 ms persistence requirement used in RoCoF relay settings. By contrast, the filtered PMU estimate enters the trip region with a substantial delay, reflecting the smoothing and lag introduced by conventional filtering.

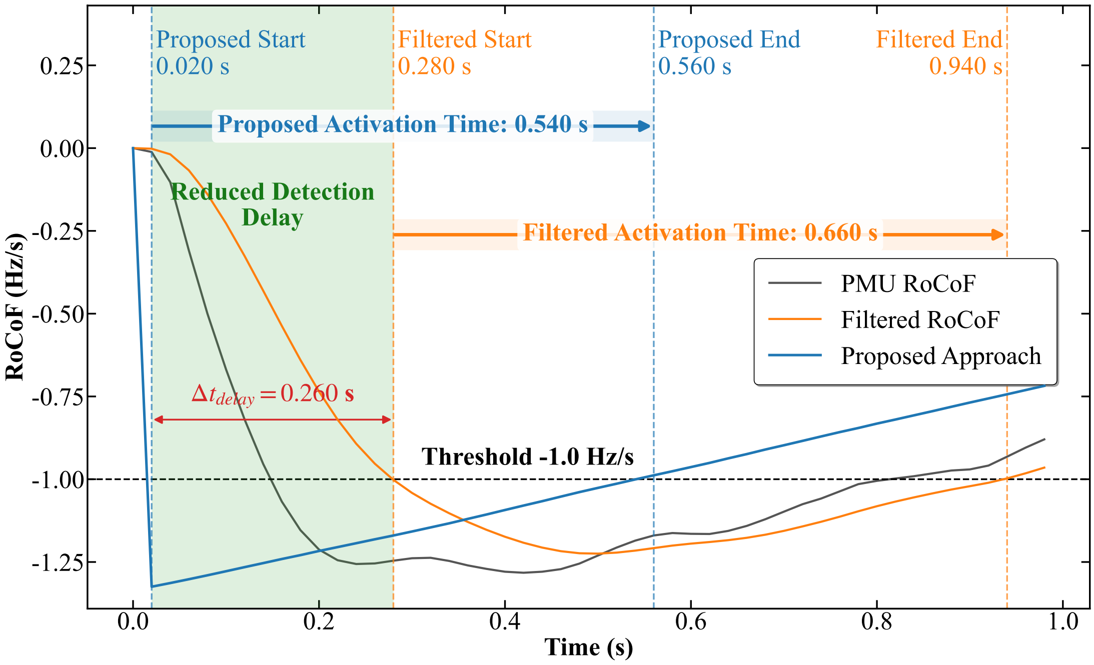

Fig. 4.1.14: RoCoF estimates used for relay assessment

Although both estimates would cause the relay to operate, the proposed approach offers a more precise indication of when the RoCoF first crosses the threshold and when it returns above it. This results in a reduced detection delay and more accurate timing information for relay decision making. The shaded regions in Fig. 4.1.14 illustrate these differences in threshold crossing behaviour.

4.2 Real-Time Estimation of Power Imbalance During Contingency Events Based on XGBoost

4.2.1 Overview of Power Imbalance Size Estimation Methodology

As noted in Chapter 2.3, real-time estimation of power imbalance currently relies on accurate measurement of RoCoF and system inertia. The RoCoF estimation method presented in the previous section contributes to a more accurate estimation of power imbalance. However, in practice, power imbalance estimation will still experience challenge associated with inertia uncertainty. This section therefore presents a new method for accurate power imbalance estimation, which leverages the learning from historically trained data and only requires frequency measurement, along with system operating conditions.

The fundamental idea of the proposed method is that the initial system frequency deviation following a contingency event will only be impacted by two main factors, i.e. the active power imbalance caused by the contingency (e.g. loss of generation) and the system operating condition, e.g. demand, generation pattern (including their controls), etc. The use of frequency measurement will largely mitigate the issues of noises and transients in RoCoF measurements. Furthermore, the precise knowledge of the system inertia becomes less critical, because the system operating conditions are directly correlated with inertia, and the learning-based model is able to implicitly learns this relationship during training. Consequently, the proposed approach is expected to overcome the limitations of the conventional method discussed in Section 2.3.

The selection of an appropriate ML algorithm is essential for ensuring computational feasibility in real-time power imbalance estimation. XGBoost is a relatively new technique in applied ML that has gained significant popularity due to its computational efficiency during training, scalability to large datasets, and consistently high predictive performance across a wide range of tasks [16]. In this work, the training data consists of time-series measurements of system frequency and system operating conditions. Although the input features are high-dimensional, the relationship between these features and the power imbalance size is not entirely linear. Under these conditions, XGBoost represents an suitable choice as it can effectively capture non-linear patterns within the data. This allows the model to maintain high predictive accuracy while remaining more computationally efficient than other algorithm. Therefore, XGBoost is adopted in this work for power imbalance size estimation.

The overall approach of the proposed XGBoost-based power imbalance size estimation method is illustrated in Fig. 4.2.1. The methodology begins with the generation of training data through detailed power system simulations that emulate a wide range of credible frequency contingency events under diverse system operating conditions. By accurately representing different load and generation patterns, the simulated data capture the dynamic frequency response of the system to varying power imbalance scenarios. Given the inherent scarcity of real world measurement data for severe contingencies, relying on such simulations is essential. Consequently, the assumption that a detailed simulation model is available for generating training data is considered realistic and practical, and it is also commonly adopted in many existing learning-based studies where real world measurement data are limited [17].

Based on the generated simulation data, raw time-series measurements are further processed to construct informative feature representations. This feature engineering process focuses on extracting key indicators that are highly sensitive to power imbalances, such as frequency. The extracted features are used to train the XGBoost regression model for power imbalance size estimation. The model training process is conducted iteratively to optimise predictive performance and ensure robustness across different operating conditions. A systematic evaluation procedure is employed to assess the accuracy and generalisation capability of the trained model, thereby verifying its effectiveness for real-time power imbalance estimation applications. Detailed discussions regarding feature engineering, as well as the training and performance validation of the XGBoost, are provided in the sections that follow.

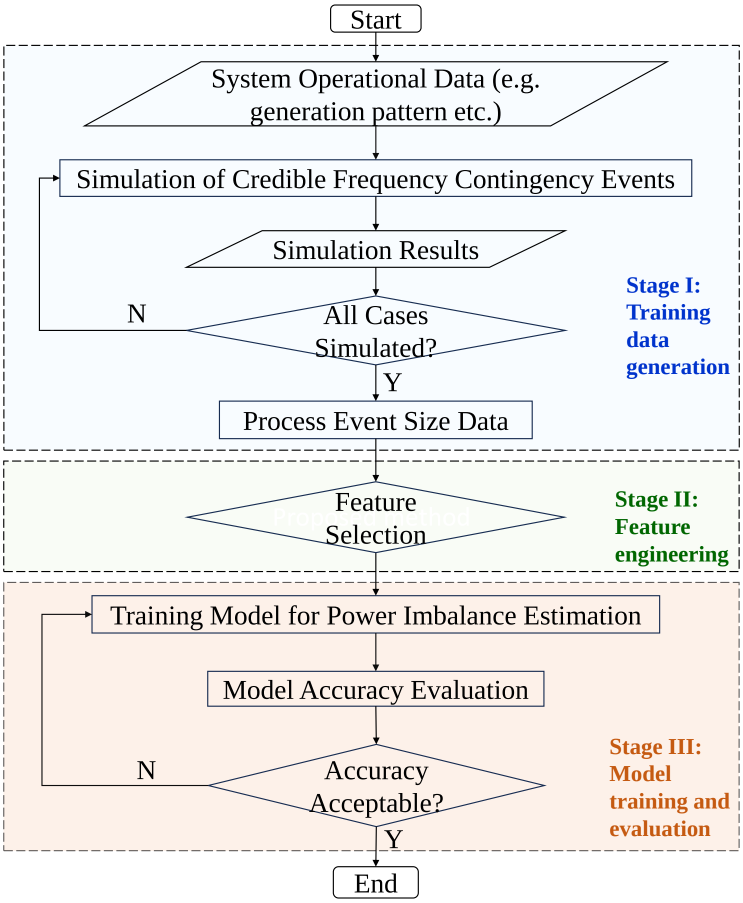

Fig. 4.2.1: The proposed power imbalance size estimation scheme

4.2.2 Design and Implementation of XGBoost-Based Power Imbalance Size Estimation

Feature Selection

The performance of XGBoost-based power imbalance size estimator highly depends on the selection and quality of its input features. Features with clear physical relevance and strong correlation to the target variable typically lead to more accurate predictions. For the estimation of power imbalance size, the initial few hundreds of milliseconds following the event (before any frequency response starts to take effect) are critical, as the frequency deviation during this short interval is only affected by the power imbalance size and system inertia. Consequently, this work employs frequency data from these initial periods alongside system operational data, i.e. generation patterns and demand levels that are directly related to system inertia, as its input features.

To quantify and interpret the contribution of each input feature to the model's output, this work employs SHapley Additive exPlanations (SHAP) analysis [18]. SHAP is an interpreter tool that treats all features as contributors to the final prediction [19]. It assigns an importance value to each feature by measuring how the model's output changes when a specific feature is included versus when it is absent. To ensure a fair distribution of importance, SHAP calculates the average marginal contribution of each feature across all possible combinations of inputs. This approach provides a reliable means to interpret the model's decision making process. Specifically, it identifies which segments of frequency variation are most instrumental in driving a particular prediction.

Fig. 4.2.2 illustrates the resulting feature importance, highlighting the relative contributions of different frequency measurement window lengths to the estimated power imbalance size. The SHAP summary plot is generated by evaluating the trained XGBoost estimator over the entire test dataset to compute the SHAP values for each sample. In this plot, each point represents an individual event, where its horizontal position denotes the feature's impact (positive or negative) on the model estimation, while its color represents the actual feature value (ranging from low to high). Overall, the analysis reveals that features derived from different temporal scales contribute unevenly to the estimation. Notably, the features are ranked by their mean absolute SHAP values, representing their global importance. The horizontal dispersion of each feature's points reflects the range of its influence on estimation. Thus, a wider distribution indicates that the feature has a more significant and varied impact on the model's output across different scenarios.

As shown in Fig. 4.2.2, the estimator relies most heavily on frequency samples taken shortly after the onset of the disturbance. Although $f_{20\text{ms}}$ maintains the highest average importance, the $f_{40\text{ms}}$ feature exhibits the largest individual SHAP values in extreme cases. This suggests that while the initial 20 ms window provides a consistent baseline for estimation, the 40 ms trajectory segment contains highly sensitive information that is particularly critical for identifying the magnitude of large power imbalances.

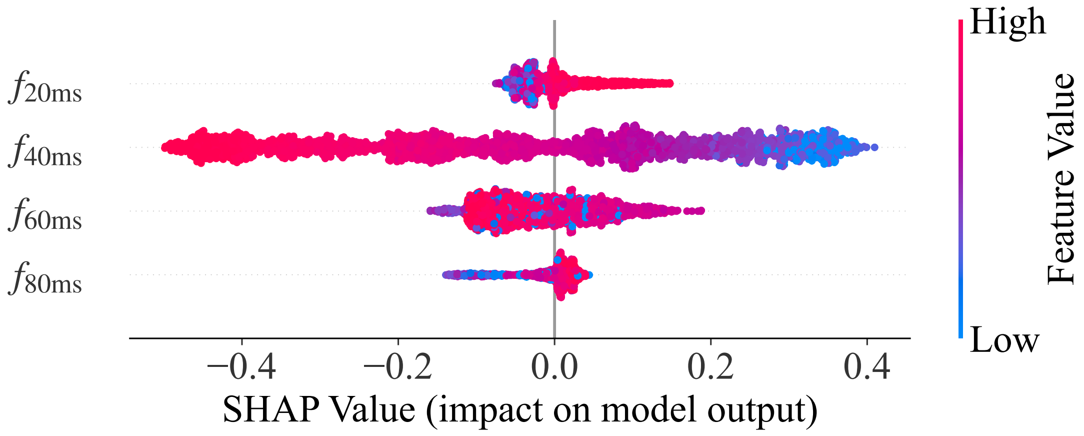

Fig. 4.2.2: Feature importance analysis using SHAP values

While SHAP quantifies the average contribution of each feature to the estimator, it does not indicate how estimation accuracy is affected when individual measurements become less reliable. To investigate this aspect, a separate sensitivity analysis was conducted on test set. For each evaluation, Gaussian noise with standard deviations $\sigma = 0.01, 0.05, 0.10, 0.20$ pu was applied into a single input feature, while all other features remained at their nominal (noise-free) values. This process was repeated independently for each of the selected features ($f_{40\text{ms}}$, $f_{60\text{ms}}$, and $f_{80\text{ms}}$) to evaluate the individual impact of each measurement's inaccuracy on the overall prediction performance. The resulting changes in prediction performance were evaluated using RMSE and Mean Absolute Percentage Error (MAPE). Whereas SHAP characterises global feature influence under nominal measurement conditions, the sensitivity analysis isolates the local dependence of the estimator on each measurement. This makes it possible to identify the features for which measurement accuracy is most critical to achieving reliable imbalance estimates.

Fig. 4.2.3 presents the sensitivity analysis results for three representative intervals, specifically, $f_{40\text{ms}}$, $f_{60\text{ms}}$ and $f_{80\text{ms}}$. Although $f_{20\text{ms}}$ exhibits high global importance, it is excluded from the sensitivity plot because 20 ms window is highly susceptible to transient measurement outliers in real world PMU data. The corresponding numerical values for each noise level are collected in Table 4.2.1. These intervals were selected due to their identified significance from the SHAP feature analysis. The results are presented in terms of RMSE and MAPE as functions of the imposed noise levels applied to each feature, ranging from 0.01 to 0.2 standard deviation, reflecting minor to substantial disturbances in the measurement data.

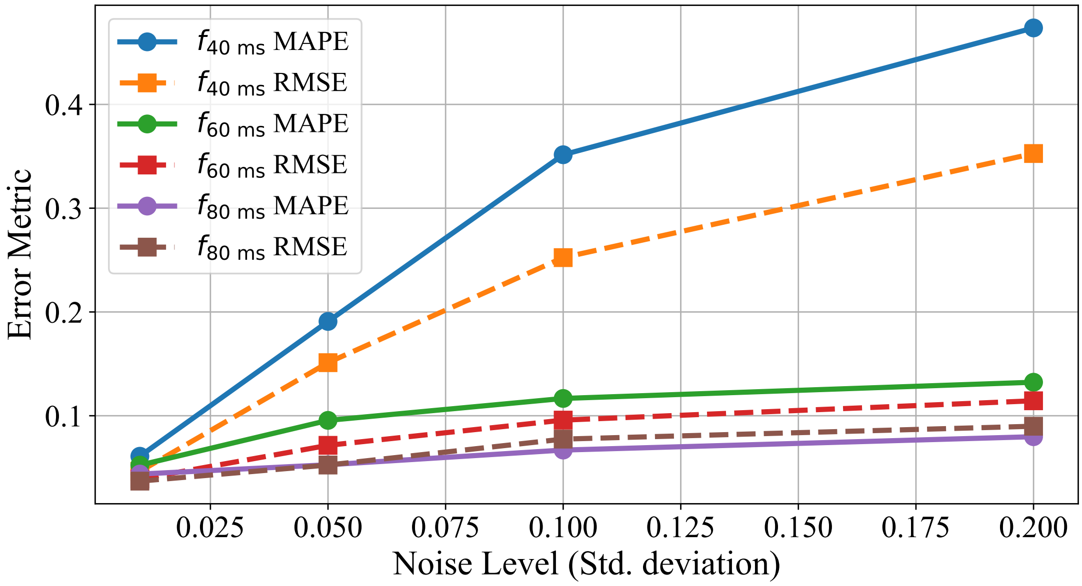

Fig. 4.2.3: Sensitivity analysis for selected frequency features

Table 4.2.1: Estimation error for representative features under increasing Gaussian noise

Noise σ (pu)	$f_{\mathbf{40ms}}$		$f_{\mathbf{60ms}}$		$f_{\mathbf{80ms}}$	
	RMSE	**MAPE**	**RMSE**	**MAPE**	**RMSE**	**MAPE**
0 (Baseline)	RMSE: 3.78 MAPE: 4.23					
0.01	4.59	5.86	4.00	5.00	3.86	4.28
0.05	15.04	18.50	7.72	9.15	5.52	5.19
0.10	26.69	34.50	10.40	11.24	6.77	6.14
0.20	35.45	49.17	13.06	14.01	8.37	7.59

Note: All values except for noise level are expressed in percentage (%).

As detailed in Table 4.2.1, the baseline errors ($\sigma = 0$) remain identical across all categories, as they represent the performance of the same trained model under ideal measurement conditions. However, a sensitivity is observed for $f_{40\mathrm{ms}}$, where the RMSE and MAPE increase significantly as the noise level rises. At a noise level of 0.10 pu, the accuracy degrades sharply, with the RMSE and MAPE increasing by nearly sevenfold compared to the baseline. It reflects the estimator's critical dependence on the initial post-disturbance sample. This is because the unique information contained in this early interval is lost to noise and cannot be reconstructed or substituted using the remaining features. In contrast, the estimator exhibits enhanced robustness toward inaccuracies in later features, such as $f_{80\mathrm{ms}}$, where error metrics remain consistently stable even under substantial noise ($\sigma = 0.20$ pu). This suggests that by 80 ms, the frequency trajectory has transitioned into a more predictable phase where information redundancy across subsequent samples allows the model to attenuate the impact of local measurement errors. The intermediate response of $f_{60\mathrm{ms}}$ confirms a gradual decrease in sensitivity as the observation window moves away from the disturbance onset, where the unique information regarding the initial power imbalance becomes less dominant.

The correspondence between the sensitivity curves and the SHAP-derived feature ranking provides mutual confirmation that the earliest frequency measurements exert the greatest influence on the estimation outcome. Another observation is that the magnitude of the sensitivity does not grow uniformly with the noise level. While increasing the noise from 0.01 pu to 0.05 pu leads to only moderate changes in the estimation error, the step from 0.05 pu to 0.10 pu for $f_{40\mathrm{ms}}$ results in a much larger deterioration. This behaviour indicates that the estimator handles small perturbations relatively well. As a consequence, precise measurement of the first $40 - 60$ ms after the disturbance is particularly important, whereas later frequency samples contribute less unique information and therefore permit more relaxed accuracy requirements.

The sensitivity analysis therefore supports the conclusions drawn from the SHAP results and further clarifies how inaccuracies in individual features influence the model's output. In addition, it highlights the necessity for precise measurement and real-time data accuracy during the earliest critical intervals following frequency disturbances, thereby informing practical recommendations for data acquisition strategies and model deployment in actual power systems.

Power Imbalance Estimation based on XGBoost

The general structure of the XGBoost model is summarised in Fig. 4.2.4. In this study, the input feature set X primarily consists of frequency measurements and system operating conditions, which are used to estimate the power imbalance size.

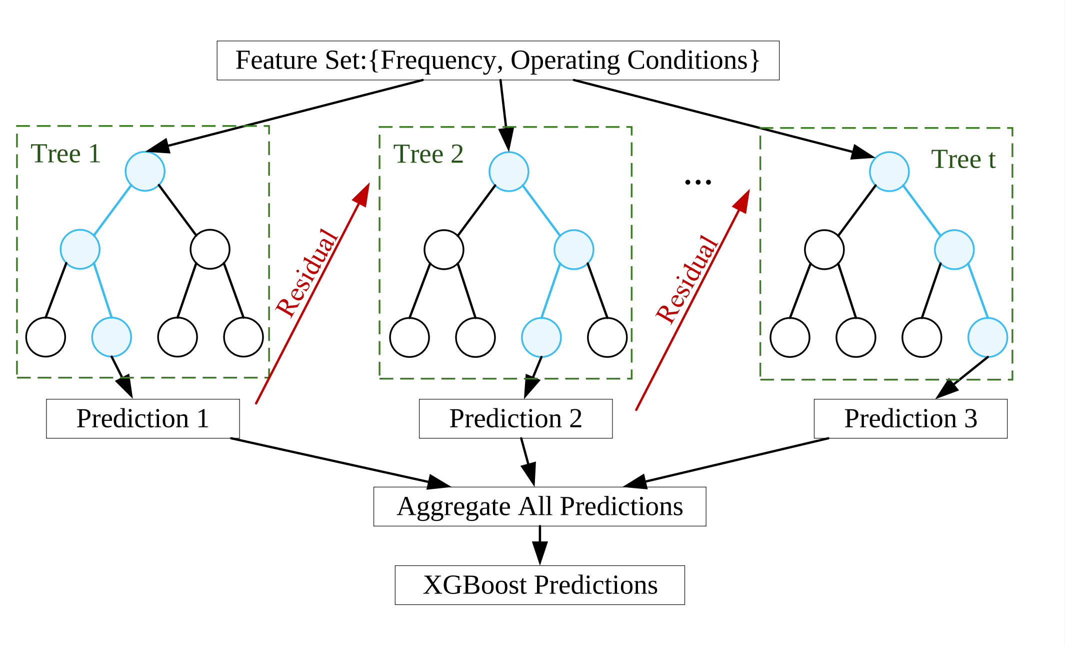

Fig. 4.2.4: Structure of the XGBoost model

Unlike RF which reduces error by averaging independent trees (as illustrated in Figure 3.1.1), XGBoost employs an additive training approach. The model begins with an initial baseline $\hat{y}_i^{(0)}$, typically representing the mean power imbalance of the training dataset, where a global initial estimation is conducted before any trees are trained. Then, an additive training approach is adopted, where each new regression tree is fitted to the current residual error (i.e. the difference between the actual power imbalance from training data and the current estimated value).

Specifically, the first tree targets the residual left by the baseline, while each subsequent tree f_t compensates for the specific estimation errors remaining from all preceding trees. After T iterations, all tree outputs are aggregated with the initial baseline to produce the final estimate. At each round t, the XGBoost model updates the estimated power imbalance size as follows

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + f_t(x_i) \tag{4.2.1}$$

where, $\hat{y}_i^{(t)}$ is the estimation value at round t, $\hat{y}_i^{(t-1)}$ is the cumulative estimation from the previous round $t-1$, and $f_t(x_i)$ is the output of the newly generated decision tree in round t.

To determine the optimal structure of each tree f_t, the objective function $\mathcal{J}^{(t)}$ is defined to balance estimation accuracy and model complexity

$$\mathcal{J}^{(t)} = \sum_{i=1}^{n} \ell\Big(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)\Big) + \Omega(f_t) \tag{4.2.2}$$

where $\ell(\cdot,\cdot)$ is the training loss function (e.g. squared error), y_i is the true power imbalance for sample i, and $\Omega(f_t)$ is the regularisation term that penalises the complexity (e.g. tree depth and leaf weights) to prevent overfitting.

Since the previous prediction $\hat{y}_i^{(t-1)}$ is fixed at round t, the loss term can be efficiently approximated using a second-order Taylor expansion. This allows the model to use the first-order (g_i) and second-order (h_i) gradients to guide the optimisation

$$\mathcal{L}^{(t)} \approx \sum_{i=1}^{n} \Big[g_i^{(t-1)}\, f_t(x_i) + \tfrac{1}{2} h_i^{(t-1)}\, f_t(x_i)^2\Big] + \Omega(f_t) \tag{4.2.3}$$

The regularisation term $\Omega(f_t)$ is defined as

$$\Omega(f_t) = \gamma J + \frac{\lambda}{2} \sum_{j=1}^{J} w_j^2 \tag{4.2.4}$$

where γ controls the complexity by penalising the number of leaves J, and λ acts as an L_2 regularisation term on leaf weights w_j to prevent excessively large values that could lead to overfitting.

To determine the optimal weight for each leaf node, the objective function is rearranged by grouping all samples assigned to the same leaf. By defining $\mathcal{I}_j$ as the set of samples falling into leaf j, the objective can be reformulated leaf-wise

$$\tilde{\mathcal{L}}^{(t)} = \sum_{j=1}^{J} \left[G_j w_j + \tfrac{1}{2} (H_j + \lambda) \, w_j^2 \right] + \gamma J \tag{4.2.5}$$

where G_j and H_j are the sums of the first-order and second-order gradients, respectively, for all samples within leaf j.

For a fixed tree structure, the optimal weight w_j on each leaf can be obtained by minimising Eq. (4.2.5). Taking the derivative of $\tilde{\mathcal{L}}^{(t)}$ with respect to w_j and setting it to zero gives

$$\frac{\partial \tilde{\mathcal{L}}^{(t)}}{\partial w_j} = G_j + (H_j + \lambda) w_j = 0 \quad \Rightarrow \quad w_j^{\star} = -\frac{G_j}{H_j + \lambda} \tag{4.2.6}$$

Substituting Eq. (4.2.6) back into Eq. (4.2.5) gives the optimal objective value

$$\tilde{\mathcal{L}}^{(t)}(w^{\star}) = -\frac{1}{2} \sum_{j=1}^{J} \frac{G_j^2}{H_j + \lambda} + \gamma J \tag{4.2.7}$$

Therefore, a tree structure that generates larger values of $\frac{G_j^2}{H_j + \lambda}$ effectively minimises the overall objective function, indicating a better fit to the residual errors and a more accurate estimation of the power imbalance.

When growing the tree, XGBoost uses Eq. (4.2.7) to evaluate each candidate split. Suppose a node with gradient sums (G, H) is split into left and right children with (G_L, H_L) and (G_R, H_R), respectively, where $G = G_L + G_R$ and $H = H_L + H_R$. The improvement in the objective function caused by this split is

$$Gain = \frac{1}{2} \left(\frac{G_L^2}{H_L + \lambda} + \frac{G_R^2}{H_R + \lambda} - \frac{G^2}{H + \lambda} \right) - \gamma \tag{4.2.8}$$

At each node, the algorithm searches over all possible features and thresholds, computes the corresponding (G_L, H_L) and (G_R, H_R), and selects the split with the largest positive gain in Eq. (4.2.8). Through this procedure, XGBoost constructs a sequence of trees where each tree is explicitly optimised to reduce the regularised loss, and the ensemble of trees in Eq. (4.2.1) provides an accurate estimate of the power imbalance size.

With the key features identified and the appropriate ML algorithm selected, this stage aims to train the model to estimate power imbalance size. Standard performance metrics, i.e. MAPE and RMSE in Section 3.4.2, are used for hyperparameter tuning and cross-validation to control overfitting and ensure reliable estimates. The MAPE is defined in Eq. (4.2.9), quantifies the average relative deviation between estimated and actual values, providing a scale-independent evaluation. $\hat{y}_i$ and y_i denote the estimated and actual power imbalance size for the i^{th} sample, respectively. n represents the total number of observations in the test set.

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{\hat{y}_i - y_i}{y_i}\right| \times 100 \tag{4.2.9}$$

The trained model is then evaluated using a validation data set to emulate the real world application, thus evaluating its performance. Iterative refinements are made until the desired accuracy is achieved.

4.2.3 Case Studies

This section present case studies that compare the performance of the proposed XGBoost-based method with the conventional method based on the swing equation. The testing environment and the processing workflow for both methods are illustrated in Fig. 4.2.5.

As shown in the figure, the test network is modelled as a microgrid connected to a main grid. The system incorporates a diverse mix of DERs, including PV, wind, hydro, CHP, BESS, and VSM. This detailed configuration enables the simulation of various system operating conditions, capturing the complex dynamic behaviours of modern grids under varying generation mixes and load profiles. During a contingency, the system frequency is measured by a PMU. Two cases are evaluated: case A (the conventional approach), which heavily relies on an accurate estimation of the total system inertia to calculate the power imbalance, and Case B, the proposed XGBoost approach.

- **Conventional Swing Equation Method:** The raw RoCoF signal from the

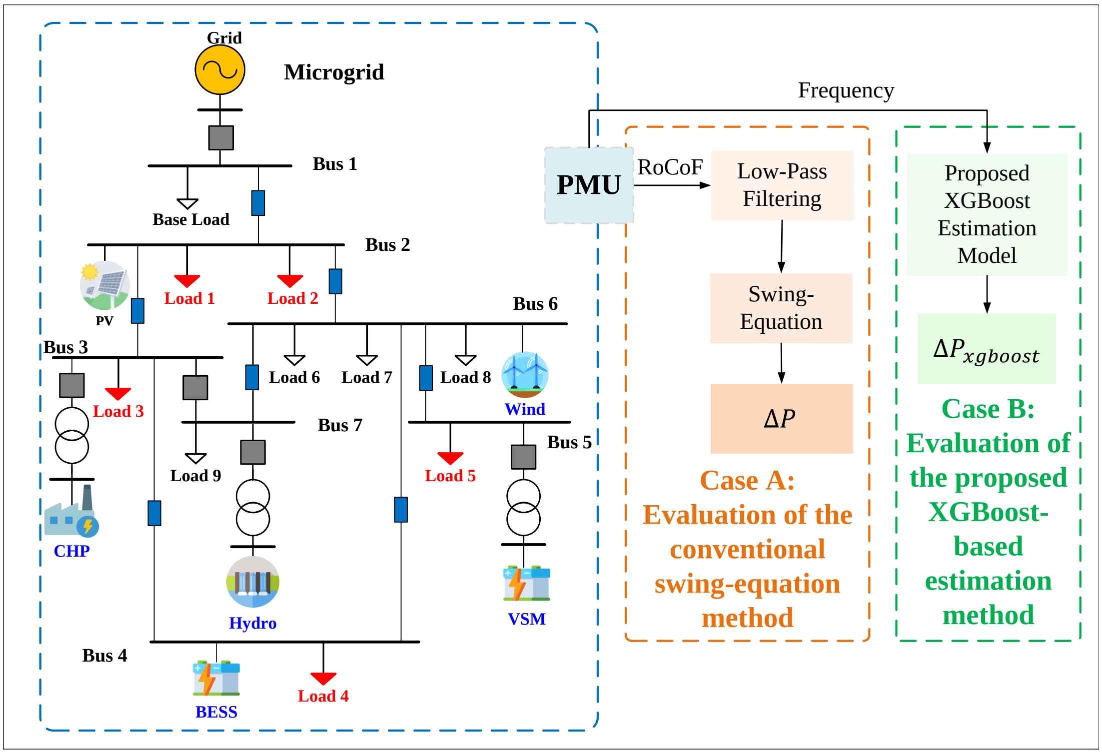

Fig. 4.2.5: Power imbalance size estimation test setup

PMU is first processed through a low-pass filter to mitigate measurement noise. The power imbalance is then calculated using the classical swing equation, which requires prior knowledge of the system inertia. This method serves as a benchmark to highlight the limitations of traditional filtering, particularly the accuracy of estimation.

- **Proposed XGBoost-based Method:** In the proposed approach, the frequency measurements are directly fed into the trained XGBoost estimation model. Unlike the conventional path, this approach bypasses the need for low-pass filtering. Instead, the model leverages learned patterns from the frequency to map the observations directly to the power imbalance size.

Evaluation of the Conventional Swing Equation Method

The performance of the conventional swing equation estimator using the measured RoCoF and system inertia is firstly evaluated. In total, there are 10 scenarios have been simulated and analysed, representing a broad range of generation loss magnitudes. As shown in Table 4.2.2, a low-pass filter with a time constant of 0.1 s is applied to reduce the noise in the measured RoCoF signal. The filtered RoCoF measurements are used for the calculation in all the cases as unfiltered measurements can be subject to severe noises, which are not commonly used in practice. The average MAPE and RMSE are 19.56% and 0.119 MW respectively, which show the estimation results are subject to a considerable level of error. As mentioned previously, this is due to the fact that filtering RoCoF, with the aim of removing measurement noise, will also lead to measured maximum RoCoF deviating from the actual value.

Table 4.2.2: Power imbalance estimation accuracy of the conventional swing equation method

No.	**Actual Generation Loss (MW)**	**MAPE (%)**	**RMSE (MW)**
1	0.11	9.09	0.01
2	0.19	20.58	0.039
3	0.24	20.50	0.049
4	0.33	21.30	0.070
5	**0.42**	**20.74**	**0.087**
6	0.53	21.09	0.112
7	0.70	20.54	0.144
8	0.77	20.52	0.158
9	**0.82**	**20.51**	**0.168**
10	0.95	20.68	0.197

Two representative cases: case 5 (0.42 MW generation loss) and case 9 (0.82 MW generation loss) are selected to illustrate the limitations of the conventional approach. The application of the low-pass filter slightly reduces the peak RoCoF in both cases, which in turn leads to a systematic underestimation of the generation loss. In case 5, the estimator reports 0.33 MW instead of the true 0.42 MW, while in Case 9 the estimate is 0.65 MW compared with the actual 0.82 MW. These deviations, each on the order of 20%, demonstrate the method's sensitivity to distortion in the RoCoF measurement.

Fig. 4.2.6 further shows that, although filtering suppresses high-frequency noise, it also attenuates the transient frequency response that the swing equation estimator relies

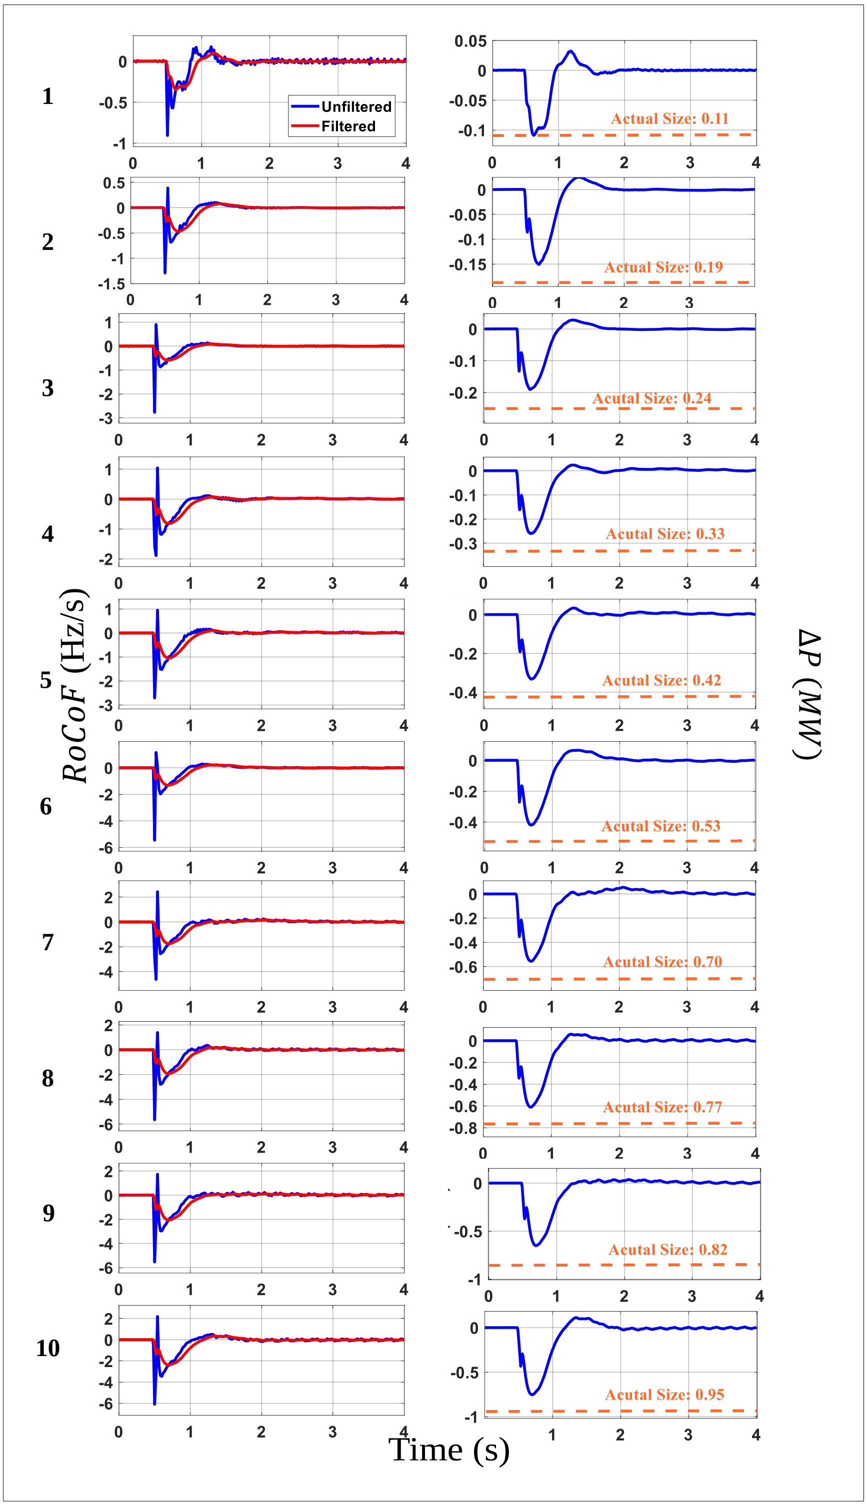

Fig. 4.2.6: Filtered and unfiltered RoCoF and estimated generation loss

on. As a result, the filtered RoCoF signal no longer preserves the information required for accurate and timely event characterisation. This illustrates that sole reliance on filtered RoCoF measurements can be misleading, and highlights the need for alternative estimation approaches capable of retaining essential transient characteristics while remaining robust to measurement noise.

Evaluation of the Proposed XGBoost-based Estimation Method

The performance of the proposed approach using XGBoost is illustrated in Fig. 4.2.7 and in Table 4.2.3. In contrast to the swing equation approach, the learning-based model exploits multiple early stage frequency samples and system-level operating conditions, enabling it to capture relationships that extend beyond the maximum RoCoF measurement. This enables the estimator to learn relationships that are not explicitly represented in analytical formulations and to extract information that would otherwise be lost due to RoCoF filtering.

Fig. 4.2.7 demonstrates a close agreement between the predicted and actual generation loss values, with the majority of points lying near the ideal reference line. The model achieves an overall MAPE of 2.67% and an RMSE of 0.01 MW, representing a substantial improvement over the swing equation method. Importantly, the cases that exhibited the largest errors under the conventional approach (case 5 and case 9) are now estimated with errors below 2%, confirming that the proposed method remains robust even when the RoCoF signal has been filtered and distorted.

A direct comparison of both methods is provided in Table 4.2.4. The proposed model reduces the average MAPE from 19.56% to 2.85% and decreases the RMSE from 0.119 MW to 0.01 MW. This improvement demonstrates that the data-driven approach effectively mitigates the limitations of RoCoF-based analytical estimation by exploiting richer temporal and operational information. Its high accuracy and robustness indicate significant potential for deployment in real-time power imbalance monitoring applications.

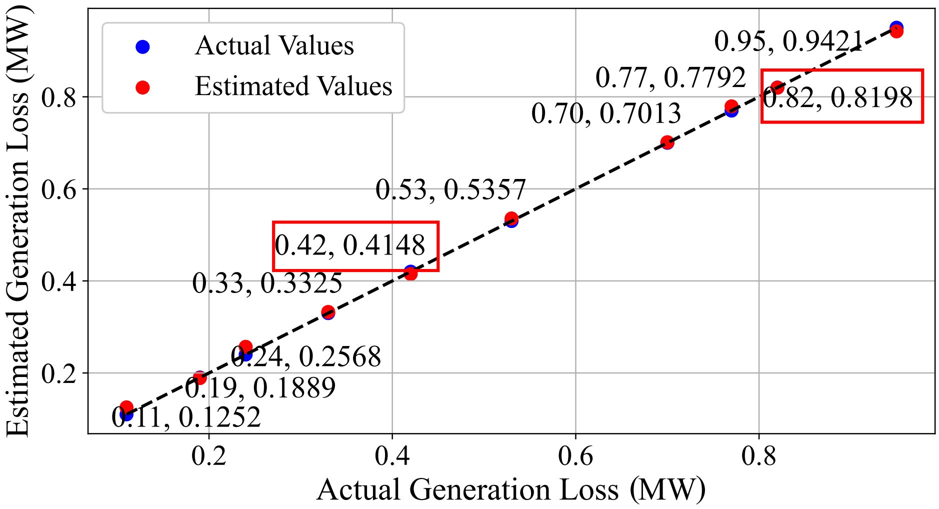

Fig. 4.2.7: Performance of the proposed power imbalance size approach using XGBoost

Table 4.2.3: Power imbalance estimation accuracy of the proposed methods

No.	**Actual Generation Loss (MW)**	**MAPE (%)**	**RMSE (MW)**
1	0.11	12.55	0.014
2	0.19	2.05	0.004
3	0.24	5.75	0.014
4	0.33	1.06	0.004
5	**0.42**	**0.81**	**0.002**
6	0.53	0.42	0.002
7	0.70	2.44	0.017
8	0.77	1.71	0.013
9	**0.82**	**1.29**	**0.011**
10	0.95	0.37	0.004

Table 4.2.4: Evaluation for different estimation methods

Method	**MAPE (%)**	**RMSE (MW)**
Conventional Approach	19.56	0.119
XGBoost Approach	2.85	0.010

4.3 Summary

This chapter has presented learning-based methods for accurate estimation of RoCoF estimation and power imbalance during contingency events, offering critical information to support effective frequency control schemes. The first part focused on accurate estimation of RoCoF using a DWT-SVR approach, while the second part introduced an estimation structure based on the XGBoost algorithm for quantifying power imbalance size in real-time.

For RoCoF estimation, the proposed method combines DWT and SVR, where the DWT extracts time-frequency information from PMU measurements, and SVR models the non-linear mapping between these features and the underlying RoCoF. This enables the estimator to capture both the fast local transients and the slower global trends embedded within the frequency signal. Extensive validation using measured PMU data demonstrated that the proposed approach consistently outperforms conventional filtered RoCoF methods across all considered window lengths. A window size analysis showed improvements in MAE and RMSE as the window size increases, with particularly large gains for short windows in which PMU data are most susceptible to noise. To support practical deployment, a MOS approach was introduced to automatically determine the optimal observation window by balancing accuracy and responsiveness. Overall, the DWT-SVR scheme provides a reliable solution for real-time RoCoF estimation under varying operating conditions.

For power imbalance estimation, a complementary XGBoost-based method has been developed to address the limitations of the swing equation approach, whose accuracy relies heavily on precise RoCoF measurements and accurate inertia values. The proposed model incorporates early stage frequency samples together with system operational variables, thereby reducing the estimator's sensitivity to RoCoF filtering errors and inertia uncertainty. Case studies covering a diverse set of operating conditions confirm that the data-driven estimator achieves consistently high accuracy, with typical errors within a ±10% margin of the true imbalance. Compared with the conventional method, the XGBoost model more effectively captures the non-linear relationships governing post-event dynamics and offers a more reliable basis for real-time power imbalance assessment.

References

[1] X. Chen, Y. Jiang, V. Terzija, and C. Lu, "Review on measurement-based frequency dynamics monitoring and analyzing in renewable energy dominated power systems," *International Journal of Electrical Power & Energy Systems*, vol. 155, p. 109 520, 2024.

[2] Q. Hong, M. Nedd, S. Norris, *et al.*, "Fast frequency response for effective frequency control in power systems with low inertia," *The Journal of Engineering*, vol. 2019, no. 16, pp. 1696–1702, 2019.

[3] D. C. Robertson, O. I. Camps, J. S. Mayer, and W. B. Gish, "Wavelets and electromagnetic power system transients," *IEEE Transactions on Power Delivery*, vol. 11, no. 2, pp. 1050–1058, 1996.

[4] S. Avdakovic, A. Nuhanovic, M. Kusljugic, and M. Music, "Wavelet transform applications in power system dynamics," *Electric Power Systems Research*, vol. 83, no. 1, pp. 237–245, 2012.

[5] A. Chakraborty and R. Mandal, "A novel technique employing dwt-based envelope analysis for detection of power system transients," in *2017 International Conference on Energy, Communication, Data Analytics and Soft Computing (ICECDS)*, 2017, pp. 346–350.

[6] O. Rioul and M. Vetterli, "Wavelets and signal processing," *IEEE Signal Processing Magazine*, vol. 8, no. 4, pp. 14–38, 1991.

[7] S. Suthaharan, "Support vector machine," in *Machine Learning Models and Algorithms for Big Data Classification: Thinking with Examples for Effective Learning.* Boston, MA: Springer US, 2016, pp. 207–235.

[8] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, pp. 273–297, 1995.

[9] H. R. Glahn and D. A. Lowry, “The use of model output statistics (mos) in objective weather forecasting,” *Journal of Applied Meteorology and Climatology*, vol. 11, no. 8, pp. 1203–1211, 1972.

[10] J. R. S. Pelland, J. Kleissl, T. Oozeki, and K. De Brabandere, “Photovoltaic and solar forecasting: State of the art,” in *Report IEA PVPS T13-01:2013*, 2013.

[11] K. Kawal, Q. Hong, S. Guo, *et al.*, “Data-driven early warning of power system oscillations in great britain,” in *2025 IEEE Power & Energy Society General Meeting (PESGM)*, IEEE, 2025, pp. 1–5.

[12] S. M. Blair, M. H. Syed, A. J. Roscoe, G. M. Burt, and J.-P. Braun, “Measurement and analysis of pmu reporting latency for smart grid protection and control applications,” *IEEE Access*, vol. 7, pp. 48 689–48 698, 2019.

[13] National Grid ESO, “Eso operational transparency forum: 17 january 2024,” Tech. Rep., Jan. 2024. [Online]. Available: `https://www.nationalgrideso.com/OTF`.

[14] National Energy System Operator (NESO), “Neso operational transparency forum: 19 march 2025,” Tech. Rep., Mar. 2025. [Online]. Available: `https://www.neso.energy/what-we-do/systems-operations/operational-transparency-forum`.

[15] A. Dyśko, D. Tzelepis, and C. Booth, “Assessment of risks resulting from the adjustment of ROCOF based loss of mains protection settings: Phase II,” University of Strathclyde, Institute for Energy and Environment, Glasgow, UK, Tech. Rep., Oct. 2015. [Online]. Available: `https://www.nationalgrid.com/sites/default/files/documents/Appendix%201%20Strathclyde%20Report%201.pdf`.

[16] Y. Wang, S. Sun, X. Chen, *et al.*, “Short-term load forecasting of industrial customers based on svmd and xgboost,” *International Journal of Electrical Power & Energy Systems*, vol. 129, p. 106 830, 2021.

[17] Defence Science and Technology Laboratory, “Machine learning with limited data,” Tech. Rep., Dec. 2020. [Online]. Available: `https://assets.publishing.`

`service.gov.uk/media/5fcdf5cd8fa8f54d61af6c9a/Dstl_Machine_Learning_with_Limited_Data_ACCESSIBLE.pdf`.

[18] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in neural information processing systems*, vol. 30, 2017.

[19] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '16, San Francisco, California, USA: Association for Computing Machinery, 2016, pp. 785–794.

Chapter 5

Real-Time Prediction of Frequency Behaviour and Emergency Frequency Control

Emergency frequency control, typically achieved via load shedding, has become increasingly challenging in modern low-inertia power systems. With faster frequency decline following major disturbances, traditional containment measures often fail to act within the narrow time window required to prevent severe nadirs. Although the previous chapters established methods for real-time estimation and prediction of frequency behaviour and critical system parameters during early stage of disturbances, the broader challenge remains to ensure the emergency actions can effectively contain frequency excursions under rapidly changing system conditions.

Conventional emergency load shedding typically relies on fixed threshold triggers and preset shedding amounts, which limits its ability to respond effectively when system inertia varies or when the disturbance size exceeds the support provided by primary frequency response [1]. Because actions occur only after thresholds are crossed, and some schemes also are also subject to hardware delays, emergency shedding actions often operate too late to prevent the frequency nadir from falling below secure limits.

This chapter introduces a learning-based approach for real-time prediction of frequency behaviour, which is subsequently used to support an emergency load shedding scheme. The proposed scheme enhances frequency control during severe frequency events by incorporating short-horizon frequency trajectory prediction into the decision process. It predicts the expected frequency evolution and the probability of critical threshold violations, thereby enabling corrective actions to be initiated with sufficient lead time. The proposed approach explicitly considers system delays and adapt the shedding decision to the estimated disturbance severity, thereby improving the responsiveness and reliability of emergency frequency control in low-inertia conditions.

5.1 Overview of Real-Time Frequency Behaviour Prediction

To support emergency load shedding decisions, three learning-based models are proposed, each addressing a distinct aspect of system frequency behaviour: (i) risk assessment of frequency threshold violations and estimation of the lead time to violation using a RF model; (ii) short-term frequency trajectory prediction in the absence of intervention based on an XGBoost model; and (iii) determination of the minimum load shedding required to arrest frequency decline using an RF-based model. An overview of the development process for these three learning models is illustrated in Fig. 5.1.1.

Stage I concerns the generation of training data, in which a range of credible frequency disturbance events are simulated to provide sufficient scenario coverage. The resulting frequency measurements and load-related parameters are then subjected to preliminary processing. In Stage II, the data are standardised prior to feature selection, with distinct feature sets identified for the different models according to their respective objectives. In parallel, three post-event observation window lengths are defined to capture different temporal characteristics of the frequency response. Stage III involves the training and evaluation of the three predictive models using the selected features and predefined window configurations.

The proposed modules are integrated into a multi-stage scheme:

- **Threshold-Crossing Risk Prediction** $\{h_{cls}^{(w)}, h_{reg}^{(w)}\}$: This module serves as the

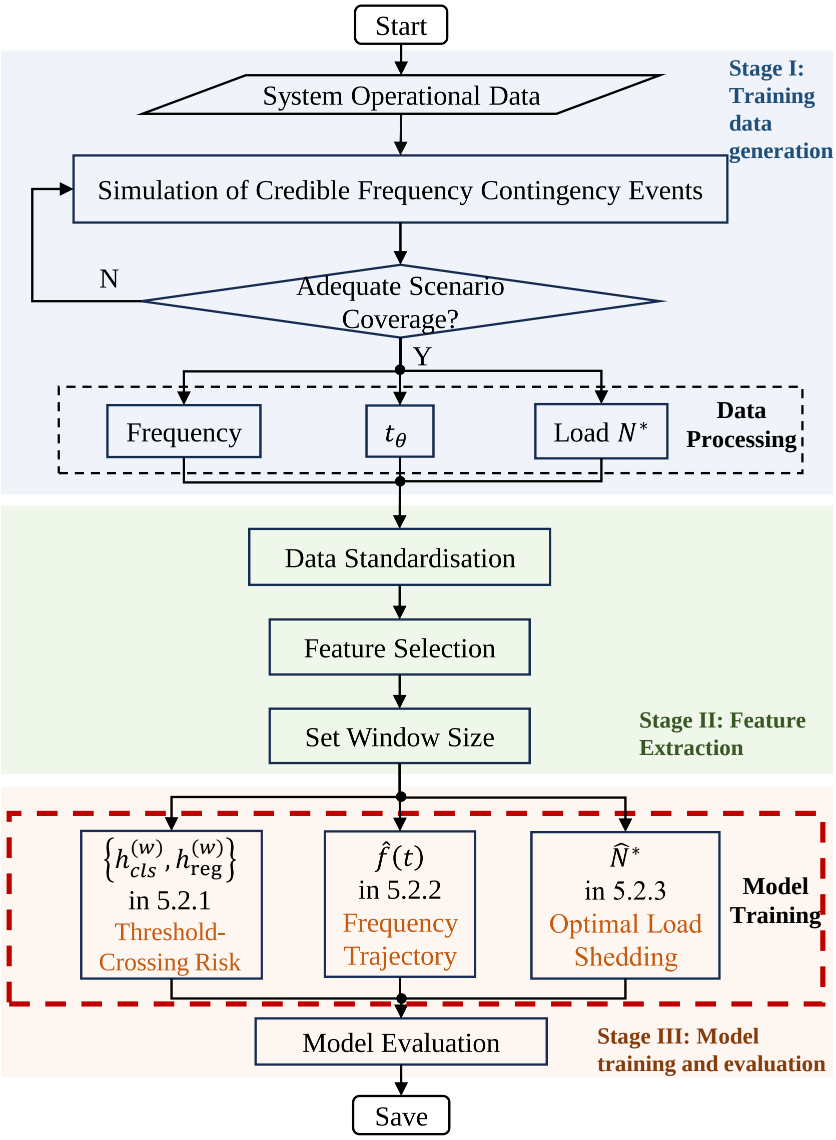

Fig. 5.1.1: Overview of the development of learning-based models for frequency behaviour prediction

detection layer by quantifying the probability of the frequency falling below the critical threshold (i.e. 47.5 Hz in this work) and providing a precise estimate of the first crossing time $\hat{t}_\theta$. The RF algorithm is selected due to its capability of handling non-linear decision boundaries. Moreover, its ensemble nature provides robustness against measurement noise and outliers commonly present in PMU data, helping to prevent over-fitting and ensuring reliable performance under transient conditions. The threshold of 47.5 Hz is selected as it represents a critical security boundary in existing literature [2], [3], beyond which the risk of generator tripping and total system collapse escalates significantly.

- **Short-Term Frequency Trajectory Predictor** $\hat{f}(t)$: This module acts as the information layer, predicting the continuous evolution of the frequency following a disturbance to provide situational awareness. The XGBoost algorithm is selected for this regression task due to its high computational efficiency and regularisation mechanisms, which prevent overfitting in high-dimensional datasets. Its well suited for capturing the dynamics of frequency trajectories following various contingencies for real-time applications.
- **Adaptive UFLS** $\hat{N}^*$: This module functions as the mitigation layer, utilising the prediction model previously to identify the minimum required load shedding amount N^* to arrest the frequency decline. By mapping specific operating conditions (H and D) and the predicted frequency response to the control action, it ensures precision in the final command, avoiding excessive intervention compared with traditional schemes. The RF algorithm is employed by predictive stability and generalisation performance. By aggregating the outputs of diverse decision trees, the model effectively mitigates the risk of prediction bias in unseen operating scenarios, ensuring that control actions are reliable across a wide range of system states.

5.2 Development of Predictive Models for Frequency Behaviour

5.2.1 Threshold Crossing Risk Prediction Based on RF

Severe disturbances cause a rapid decline in system frequency due to a sudden generation and load imbalance. Although the frequency nadir is an important indicator of disturbance severity, it occurs relatively late in the event and cannot be inferred reliably from the short post-event measurements available at detection. Relying solely on eventual nadir prediction therefore risks delaying intervention. For emergency frequency load shedding, the operational requirement is to determine, at the earliest possible stage, whether the frequency trajectory is likely to enter the emergency action band and to estimate the time at which this will occur. With a specified trigger threshold, the prediction task is posed as predicting both the occurrence of a crossing and the corresponding first crossing time. This formulation provides the information needed to intervene before unsafe conditions develop and avoids the instability associated with nadir based prediction under rapidly evolving conditions.

Fig. 5.2.1 illustrates the timing quantities used throughout this subsection, including the disturbance onset, the detection instant, the first crossing time, and the three observation windows used for prediction. This diagram serves as the temporal reference for all subsequent definitions.

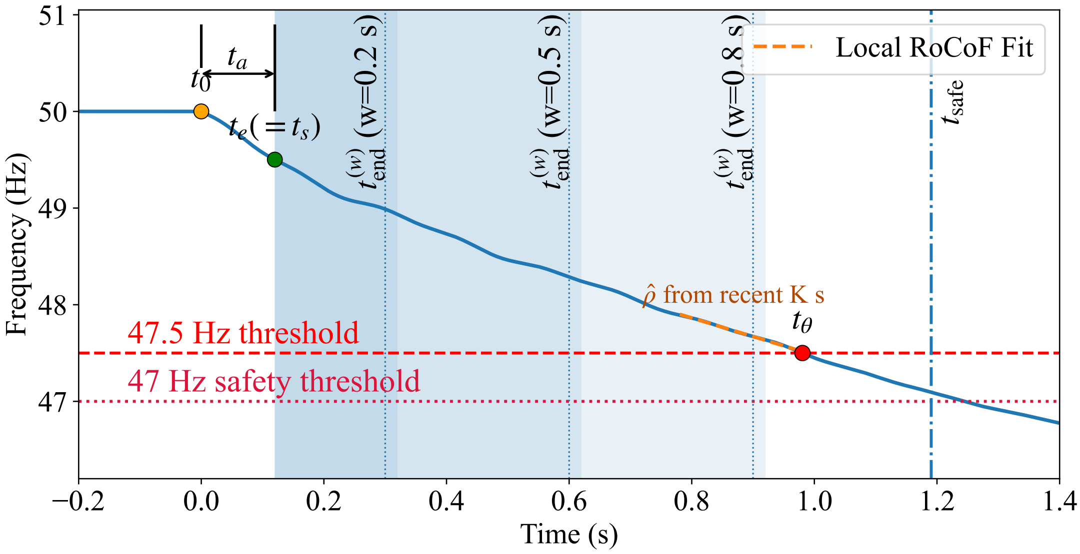

Fig. 5.2.1: Timing structure and windowing scheme for the f_{trig} threshold crossing prediction

Let t_0 denote the physical disturbance onset. A disturbance is detected at time t_e when the locally estimated RoCoF exceeds its threshold. The detection latency is defined as

$$t_a = \; t_e - t_0 \; > \; 0 \tag{5.2.1}$$

representing the offset between the true onset and the detection instant. In implementation, the first reliable post-event sample coincides with detection

$$t_s = t_e \tag{5.2.2}$$

thus, prediction starts at t_s.

Let $f(t)$ be the measured frequency and Δt the sampling interval. With the trigger and safety thresholds defined as $f_{trig} = 47.5\,\mathrm{Hz}$ and $f_{safe} = 47\,\mathrm{Hz}$, the binary crossing indicator z and first crossing time t_θ are defined as

$$z = \mathbf{1}\big(\exists\, t \in [t_s,\, t_s + T_{\max}] : f(t) < f_{safe}\big) \tag{5.2.3}$$

$$t_\theta = \inf\{\, t \geq t_s : f(t) < f_{trig} \,\} \tag{5.2.4}$$

In Eq. (5.2.3), $\mathbf{1}(\cdot)$ is the indicator function. Thus, $z = 1$ if the trajectory $f(t)$ falls below f_{safe} within the maximum observation horizon $T_{\max}$ (0.8 s), and $z = 0$ otherwise. When no crossing occurs, t_θ is undefined; such cases are utilised for the classification task only.

Prediction is based on a set of fixed observation window lengths, denoted as $W = \{w_1, w_2, w_3\}$. Each window starts at t_s, and the corresponding window end time is defined as

$$t_{end}^{(w)} = t_s + w, \quad w \in W \tag{5.2.5}$$

In this work, these parameters are set as $w_1 = 0.2\,\mathrm{s}$, $w_2 = 0.5\,\mathrm{s}$, and $w_3 = 0.8\,\mathrm{s}$. This selection is intended to capture the multi-stage temporal characteristics of the frequency response following a disturbance.

The number of discrete samples N_w and the last available sample time $t_{end-samp}^{(w)}$ are determined by the sampling interval Δt as specified in Eq. (5.2.6) and Eq. (5.2.7).

$$N_w = \left\lceil \frac{w}{\Delta t} \right\rceil + 1 \tag{5.2.6}$$

$$t_{end-samp}^{(w)} = t_s + \Big(\left\lceil \frac{w}{\Delta t} \right\rceil - 1 \Big) \Delta t \tag{5.2.7}$$

Each window produces a feature vector $x^{(w)}$ consisting of system descriptors and the corresponding windowed frequency samples

$$x^{(w)} = [H,\, D,\, f(t_s),\, f(t_s + \Delta t),\, \ldots,\, f(t_s + (N_w - 1)\Delta t)] \tag{5.2.8}$$

For each observation window, two RF models are trained: a classifier $h_{cls}^{(w)}$ that estimates the probability $p(z = 1 \mid x^{(w)})$ and a regressor $h_{reg}^{(w)}$ that predicts the first crossing time t_θ conditional on $z = 1$. Simulated events where the frequency remains above the safety threshold ($z = 0$) are used only as negative samples for classification. These are excluded from the regression stage to prevent the model from learning ill-defined targets, thereby ensuring higher accuracy for critical cases.

As each predictor is implemented as an RF, every tree during training recursively partitions the feature space by selecting splits that minimise the node impurity. For a node containing sample set S, the impurity is defined as

$$Impurity(S) = \frac{1}{|S|} \sum_{i \in S} \sum_{j} \left(y_i^{(j)} - \bar{y}^{(S)}\right)^2 \tag{5.2.9}$$

where $\bar{y}^{(S)}$ is the mean of the output dimension j over the samples in S. The tree continues splitting until a standard stopping rule is reached, such as limits on depth, the minimum number of samples per leaf, or an insufficient reduction in impurity.

After training, the combined output vector for window w is defined as:

$$\hat{y}^{(w)} = [\,\hat{p},\, \hat{t}_\theta\,] \tag{5.2.10}$$

where $\hat{p}$ denotes the estimated crossing probability $P(z = 1 \mid x^{(w)})$ from $h_{cls}^{(w)}$, and $\hat{t}_\theta$ represents the predicted crossing time from $h_{reg}^{(w)}$ conditional on $z = 1$. Simulated events where $z = 0$ are used only as negative samples for classification and are excluded from the regression stage.

This prediction task is formulated around the f_{trig} intervention threshold so that the model outputs can be used directly for emergency frequency control. The predicted quantities z and $\hat{t}_\theta$ provide the key information needed for rapid assessment of frequency security, where z indicates whether the trajectory enter the emergency band, and $\hat{t}_\theta$ gives the earliest expected entry time. Evaluating these quantities on short post-event windows allows the controller to anticipate unsafe behaviour before the nadir forms, thereby enabling timely intervention and limiting decision latency to within the window duration.

Once the predicted quantities z and $\hat{t}_\theta$ are available, they will be transferred to a latest allowable intervention time for emergency load shedding algorithm as described in Section 5.2.3. The classifier output is compared with a decision threshold of 0.5. When the classifier output exceeds the decision threshold ($z = 1$), the predicted crossing time $\hat{t}_\theta$ is combined with an estimate of the local RoCoF $\hat{\rho}$ to determine the latest safe intervention time. The slope $\hat{\rho}$ is obtained from a backward LS fit to the most recent K samples, providing an estimate of the current rate of frequency decline. This slope indicates how long the trajectory would take to fall from f_{trig} to the safety boundary f_{safe}. To avoid producing unrealistically long deadlines when the slope estimate is noisy or nearly flat, a minimum-slope requirement $|\hat{\rho}| \geq 0.5\,\mathrm{Hz/s}$ is imposed. Incorporating this slope and the actuation latency τ_{exec} creates the latest permissible control time

$$t_{safe} = \hat{t}_\theta + \frac{f_{trig} - f_{safe}}{|\hat{\rho}|} - \tau_{exec} \tag{5.2.11}$$

If the current time within the window has already exceeded t_{safe}, or if the remaining time to $\hat{t}_\theta$ is shorter than the execution delay τ_{exec}, the load shedding command determined by the adjustment model is issued immediately. Otherwise, sampling continues until the deadline is reached, avoiding unnecessary waiting for the window to complete. When the classifier output indicates ($z = 0$), no action is taken in the current window, and monitoring proceeds to the subsequent, longer window as a fallback. The offline training process for the classifiers and regressors used across these windows is detailed in Algorithm 5.2.1.

The combination of probabilistic risk detection and a timing rule that accounts for execution latency provides a balance between responsiveness and safety margin. This structure enhances data efficiency and supports reliable early intervention under rapidly

Algorithm 5.2.1 Training algorithm for f_{trig} threshold-crossing risk predictors

```
Require: Event set $N$ with frequency samples $f(k)$ (20 ms); inertia $H$; demand $D$;
    window set $W = \{w_1, w_2, w_3\}$ (e.g. $\{0.2, 0.5, 0.8\}$ s); trigger and safety thresholds
    $(f_{trig}, f_{safe}) = (47.5, 47)$ Hz; longest window $T_{\max}$; RoCoF trigger $p_{trig}$
Ensure: $\{\, h_{cls}^{(w)},\, h_{reg}^{(w)},\, \mu^{(w)},\, \sigma^{(w)} \,\}_{w \in W}$
for $w \in W$ do
    Build $x^{(w)}$ using (5.2.2)-(5.2.8); construct targets for classification and regression.
    Fit scaler $(\mu^{(w)}, \sigma^{(w)})$ on $x^{(w)}$ and standardise.
    Train classifier $h_{\mathrm{cls}}^{(w)}$ and regressor $h_{reg}^{(w)}$.
    Save $\{h_{cls}^{(w)}, h_{reg}^{(w)}, \mu^{(w)}, \sigma^{(w)}\}$.
end for
return all models and scalers.
```

evolving operating conditions.

5.2.2 Short-Term Frequency Trajectory Prediction Model

The trajectory prediction module supplements the threshold-crossing assessment by providing a short forecast of the frequency evolution following detection. Its purpose is not to trigger control actions, but to offer operators an interpretable view of the expected post-event behaviour when no shedding action is taken. The task is formulated as supervised regression, where early post-event measurements and system descriptors are used to estimate several near future frequency values sampled at the same interval Δt. All time quantities follow the same reference (t_e, t_a) established in the threshold-crossing module to maintain consistency across predictive components.

Using the same reference t_a as the threshold-crossing module, the feature vector and prediction target segment are

$$x_{traj} = \big[\, H,\, D,\, f(t_e + t_a + k\Delta t) \,\big]_{k=0}^{9} \tag{5.2.12}$$

$$f_{future} = \big[, f(t_e + t_a + s), \big]_{s \in w_1, w_1 + \Delta t, \ldots, w_2} \tag{5.2.13}$$

so that the inputs contain ten samples after $t_e + t_a$ with the system descriptors (H, D), and the output consists of the uniformly spaced future values in the interval $[\, t_e + t_a + w_1,\, t_e + t_a + w_2 \,]$.

A direct multi-horizon design is adopted, in which an independent regressor is trained for each prediction horizon

$$\begin{gathered}\hat{f}(t) = M_t(x_{traj}), \\ t \in \{\, t_e + t_a + w_1, \, \ldots, \, t_e + t_a + w_2 \,\}\end{gathered} \tag{5.2.14}$$

This avoids the error accumulation that arises in recursive forecasting, enables fully parallel estimation, and provides more stable calibration across the short prediction range. Each model M_t is an XGBoost gradient boost tree regressor with a squared error loss and ℓ_2 regularisation term

$$\mathcal{L}^{(t)} = \sum_{i=1}^{N} \big(y_i^{(t)} - \hat{y}_i^{(t)}\big)^2 \; + \; \sum_{k=1}^{K} \Omega(f_k) \tag{5.2.15}$$

where $\Omega(f_k)$ the regularisation term penalises model complexity

$$\Omega(f_k) = \gamma \, T_k \; + \; \tfrac{1}{2} \, \lambda \, \|w_k\|_2^2 \tag{5.2.16}$$

with T_k is the number of leaves in tree k, w_k the vector of leaf weights, and γ and λ hyperparameters controlling complexity.

At run time, while trajectory prediction is performed continuously, its results are particularly prioritised for display when the threshold-crossing module issues a no-shed decision. This provides the operator with the necessary evidence to verify that the frequency remains within safe limits despite without control action. The feature vector x_{traj} is assembled using the same time reference t_a with (H, D), and the set $\{\hat{f}(t)\}$ is obtained from the individual horizon regressor. When the realised nadir over the interval $[\, t_e + t_a, \, t_e + t_a + w_2 \,]$ remains above 48 Hz, the actual and predicted trajectories are displayed for operator awareness.

5.2.3 Adaptive UFLS Model

Reliable estimation of the minimum shedding requirement is essential for maintaining frequency security following severe disturbances. Instead of relying on fixed stage-based settings, which do not reflect variations in operating conditions, the shedding requirement is learned from data through a supervised regression model that maps operating descriptors and a candidate shedding level to the resulting frequency nadir.

For a candidate shedding amount $N \geq 0$, the model predicts the corresponding post-shedding nadir

$$\hat{f}_{min} = h(D, H, N) \tag{5.2.17}$$

where $h(\cdot)$ is an RF regressor trained on historical disturbance cases with input vector $[D, H, N]$ and target f_{min}. The model parameters are obtained by minimising the mean squared error

$$min \sum_{i=1}^{N} \left(f_{min,i} - \hat{f}_{min,i}\right)^2 \tag{5.2.18}$$

To determine the required shedding level, the controller searches over the set of available load blocks $N \in \{0, 1, \ldots, K\}$ and selects the smallest N for which the predicted nadir satisfies the safety requirement f_{safe}

$$\begin{aligned} N^*(D,H) = min\Big\{ & N \in \{0,1,\ldots,K\} : \\ & \hat{f}_{min}(D,H,N) \geq f_{safe} \Big\} \end{aligned} \tag{5.2.19}$$

The search is monotonic in most operating conditions, allowing it to be implemented efficiently. When no shedding level meets the constraint, a conservative fallback $N^* = K$ is used. Given N^*, the shedding vector $s^* \in \{0, 1\}^K$ follows a predefined priority order

$$s_i^* = \begin{cases} 1, & i = 1, 2, \ldots, N^* \text{ (shed } N^* \text{ loads)} \\ 0, & i = N^* + 1, \ldots, K \text{ (keep the rest)} \end{cases} \tag{5.2.20}$$

meaning that the top-priority N^* loads are shed while the remaining loads are retained. The binary vector s^* is then converted to integers through

$$u_k = \begin{cases} -1, & s_k^* = 1 \text{ (shed)} \\ 0, & s_k^* = 0 \text{ (keep)} \end{cases} \tag{5.2.21}$$

5.3 Emergency Load Shedding Based on Predicted Frequency Behaviour

This section presents how the three developed prediction models are coordinated to support emergency load shedding decisions. Rather than introducing additional predictive models, the focus here is on explicating the decision logic that links model outputs to control actions under strict timing constraints. All models are assumed to be trained offline, enabling their immediate invocation during real-time operation. An overview of the integrated approach is illustrated in Fig. 5.3.1.

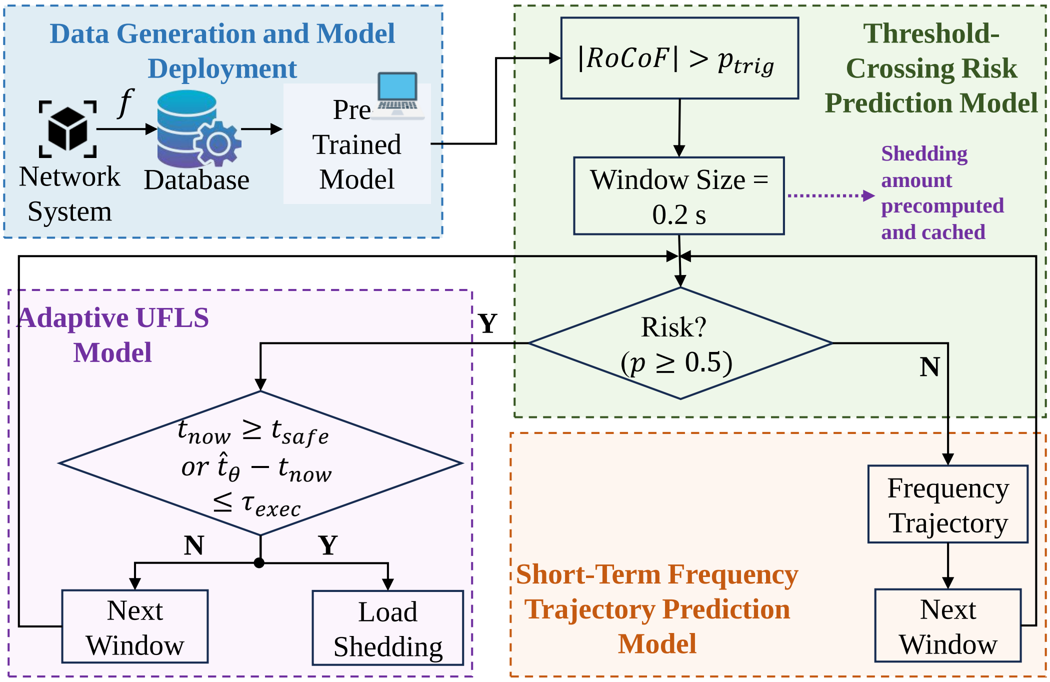

Fig. 5.3.1: Overview of proposed approach for real-time load shedding

The scheme is initiated by a RoCoF-based disturbance detection mechanism, which identifies the event occurrence time t_e and establishes a analysis starting point t_a. Simultaneously, the nadir regression model $\hat{f}_{\min}(D, H, N)$ is evaluated once to determine the minimum amount of load shedding required to maintain the system frequency above the safety limit f_{safe}. Then, the resulting shedding level N^* is translated into a breaker command and stored in advance. By separating this preparatory computation from the subsequent time-critical phase, the scheme avoids any need for online optimisation during emergency operation.

Following the disturbance, the system response is assessed progressively over a sequence of post-event observation windows $w \in \{0.2, 0.5, 0.8\}$. Within each window, the threshold-crossing risk classifier evaluates whether the current frequency evolution is likely to lead to unsafe behaviour. When the estimated risk remains below the predefined threshold, no shedding action is taken, and the scheme may optionally invoke the short-term trajectory prediction model to provide additional situational awareness. The process then advances to the next, longer observation window to refine the assessment.

When a high risk is identified, focus shifts from risk detection to timing considerations. The corresponding regression model estimates the first expected threshold-crossing time $\hat{t}_\theta$. This estimate combined with the local RoCoF $\hat{\rho}$ and the known actuation latency τ_{exec}, defines a latest allowable intervention time t_{safe} according to Eq. (5.2.11). A time-based decision rule then determines whether the stored shedding command should be dispatched immediately or whether further monitoring within the current window remains feasible. Consequently, once the remaining margin becomes insufficient to accommodate the execution delay, the precomputed shedding command is issued without further deliberation. Otherwise, the scheme retains the ability to defer action and revise the assessment as additional measurements become available in subsequent windows. This strategy allows the methodology to incorporate more post-event data, thereby improving the accuracy of the risk and timing estimates before committing to a control action. Furthermore, executing the control at the latest permissible time avoids excessive load shedding, ensuring that load shedding is only executed when frequency recovery is truly non-viable.

The following algorithm summarises the operational procedure for the real-time predictive scheme.

Algorithm 5.3.1 Real-time emergency frequency control loop

```
Require: Frequency $f(k)$ (20 ms); inertia $H$; demand $D$; threshold crossing models
    $\{h_{cls}^{(w)}, h_{reg}^{(w)}\}_{w \in W}$; trajectory models $\{M_t\}$; nadir regressor $\hat{f}_{min}(D, H, N)$; safety
    limit $f_{safe}$; risk threshold $p_{thr}$; RoCoF trigger $p_{trig}$; actuation latency $\tau_{exec}$
Ensure: Breaker command $u \in \{0, -1\}^K$ (shed/hold)
Event detection: if $|\text{RoCoF}| > p_{trig}$ then declare disturbance; record $t_e$ and $t_a$.
Precompute shedding: find minimal $N^*$ such that $\hat{f}_{\min}(D, H, N^*) \geq f_{safe}$; map
    to $u$ and cache.
for $w \in \{0.2, 0.5, 0.8\}$ s do
    Form feature vector $x^{(w)} \leftarrow [\, H, D, f(t_a : t_a + w) \,]$
    $p \leftarrow h_{cls}^{(w)}(x^{(w)})$
    if $p < p_{thr}$ then
        No shed at window $w$; Log trajectory using $\{M_t\}$
        if $w$ is final then return HOLD
        elsecontinue to next window
        end if
    else
        $\hat{t}_\theta \leftarrow h_{reg}^{(w)}(x^{(w)})$; estimate local RoCoF $\hat{\rho}$
        Compute permissible deadline (5.2.11)
        if now $\geq t_{safe}$ or $\hat{t}_\theta$ − now $\leq \tau_{exec}$ then
            return DISPATCH $u$
        else
            if $w$ is NOT final then
                continue to next window for improved prediction accuracy
            else
                Wait until now $\geq t_{safe}$, then return DISPATCH $u$
            end if
        end if
    end if
end for
```

5.4 Case Studies

5.4.1 Overview of the Test Configuration

The proposed frequency behaviour prediction models and the emergency load shedding scheme are evaluated using the real-time HiL test configuration as illustrated in Fig. 5.4.1. The platform consists of an RTDS simulator running a detailed microgrid model and an external Python-based controller implemented on a PC. The two components exchange measurements and control commands through a GTNET communication card and a database interface, enabling a realistic representation of communication and processing delays.

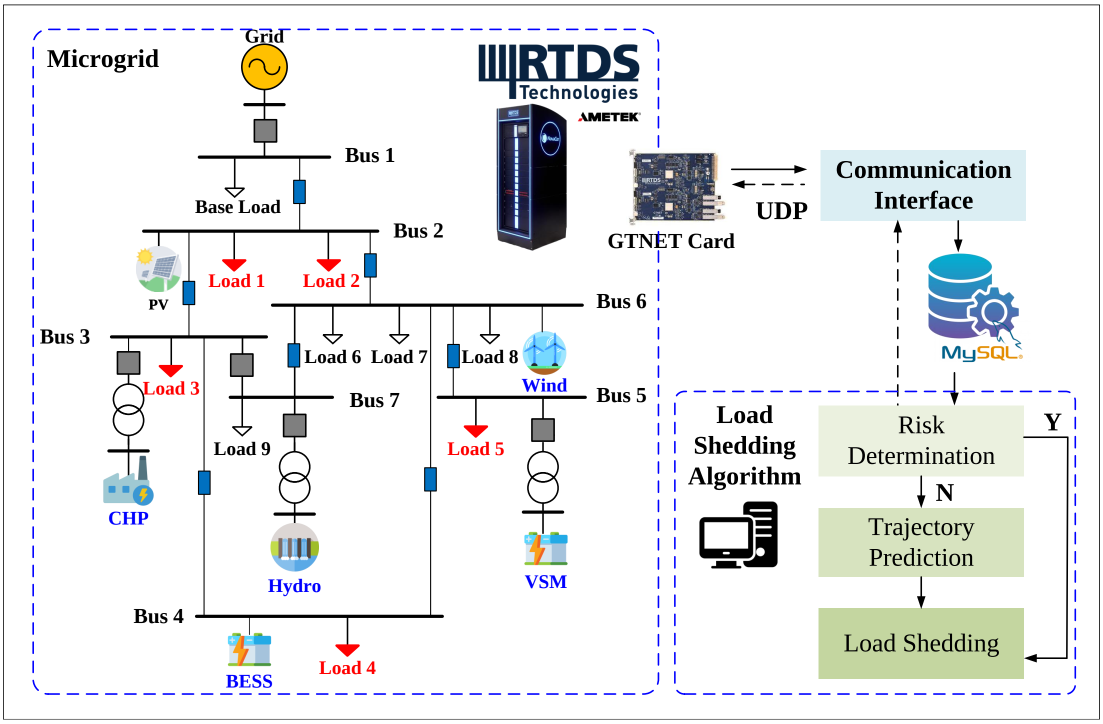

Fig. 5.4.1: HiL test setup for validating emergency frequency control scheme

Within the RTDS, a modified microgrid model is implemented to represent the network and generation dynamics. The upstream utility grid is emulated using an equivalent voltage source connected at the PCC. The microgrid includes 11 loads, of which 6 are non-sheddable and 5 are controllable at the PCC. The DER portfolio comprises a gas turbine-based CHP, a mini hydro unit, a BESS, and PV and wind units, as shown in Fig. 5.4.1.

Real-time measurements are streamed from the RTDS through the GTNET card using the User Datagram Protocol (UDP) communication protocol, which is widely used for PMU data streaming. A communication interface receives the UDP packets and writes the time-stamped data to a MySQL database, which acts as a buffer between the simulator and the controller.

The load shedding controller module, implemented on a separate PC, periodically retrieves the most recent measurement window from the database, executes the risk-classification and short-horizon trajectory prediction modules, and computes the required shedding action. The corresponding breaker commands are then transmitted back to the RTDS over the same GTNET/UDP channel, where they are applied to the appropriate breakers in the microgrid model. This HiL test explicitly captures measurement latency, communication delay, database access time, computation time, and breaker actuation delay, thereby enabling a realistic validation of the proposed adaptive emergency frequency control scheme.

To quantify the performance of the predictive models within this real-time environment, the approach is assessed from three perspectives: scalar prediction accuracy, risk classification performance, and time-related accuracy. Let y_i and $\hat{y}_i$ denote the true and predicted scalar targets for sample i in a test set of size N.

(1) Scalar prediction accuracy. For regressions task such as nadir prediction or minimum post-shedding frequency estimation, accuracy is measured by the MAE and R^2 as shown in Section 3.4.2.

Since operational risk increases as the nadir approaches the critical under-frequency limit f_{crit}, a Threshold-Normalised Absolute Error (TNAE) is also reported. TNAE scales the error relative to the distance between the target and the threshold

$$TNAE = \frac{1}{N}\sum_{i=1}^{N} \frac{|y_i - \hat{y}_i|}{\max\left(|y_i - f_{crit}|,\, \varepsilon\right)} \tag{5.4.1}$$

where a small constant $\varepsilon > 0$ ensures numerical stability when y_i is extremely close to f_{crit}.

(2) Risk classification performance. To assess the model's ability to trigger timely alerts, the continuous frequency predictions are converted into binary risk indicators (e.g. whether f_{nadir} violates f_{crit}). The performance is evaluated using the Brier score to measure the calibration of predicted risk probabilities. With $z_i \in \{0, 1\}$ indicating whether a crossing occurred and $p_i \in \{0, 1\}$ the predicted probability, the score is

$$Brier = \frac{1}{N}\sum_{i=1}^{N}(p_i - z_i)^2 \tag{5.4.2}$$

where lower values correspond to more accurate and better calibrated probability estimates.

(3) Time-related accuracy. For first-crossing time regression, accuracy is summarised using the Empirical Cumulative Distribution Function (ECDF) of absolute time errors. For each event, let $e_i = |\hat{t}_i - t_i|$ denote the absolute difference between the predicted and true first-crossing times. For any tolerance $\tau \geq 0$, the ECDF is

$$F_e(\tau) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}\{e_i \leq \tau\} \tag{5.4.3}$$

where $1\{\cdot\}$ is the indicator function and gives the fraction of events whose prediction is within τ seconds of the true value. Viewing $F_{|e|}(\tau)$ as a curve provides an interpretable accuracy profile, where curves that rise more quickly and approach one indicate smaller errors and more reliable time.

5.4.2 Implementation and Evaluation of Learning-Based Frequency Behaviour Prediction Models

To evaluate the proposed emergency frequency control scheme, four learned components are trained to support: (i) a classifier that predicts the probability of crossing 47.5 Hz, (ii) a regressor that estimates the time to threshold t_θ, (iii) a short-term frequency trajectory predictor, and (iv) a nadir-based load-shedding adjustment model. All models are trained under three observation windows which are short (0.2 s), nominal (0.5 s), and long (0.8 s). The training and validation sets were generated from simulation studies under the operating conditions described earlier.

Fig. 5.4.2 evaluates the calibration of the threshold-crossing classifier using reliability diagrams. Test samples are grouped into probability bins according to the predicted risk of dropping below 47.5 Hz. For each bin, the average predicted probability is plotted against the empirical crossing frequency, curves close to the diagonal indicate well calibrated probabilities.

Across the three observation windows, the curves lie close to the diagonal, and the nominal 0.5 s window achieves the lowest brier score (approximately 0.006), matching the performance of the 0.8 s window. The 0.2 s window is slightly conservative in the mid-range (around predicted probabilities of $0.3 - 0.4$) but remains accurate at the extremes, with an overall Brier score near 0.007. These results indicate that probability estimates remain stable across window lengths, with the nominal window offering the best balance between responsiveness and calibration.

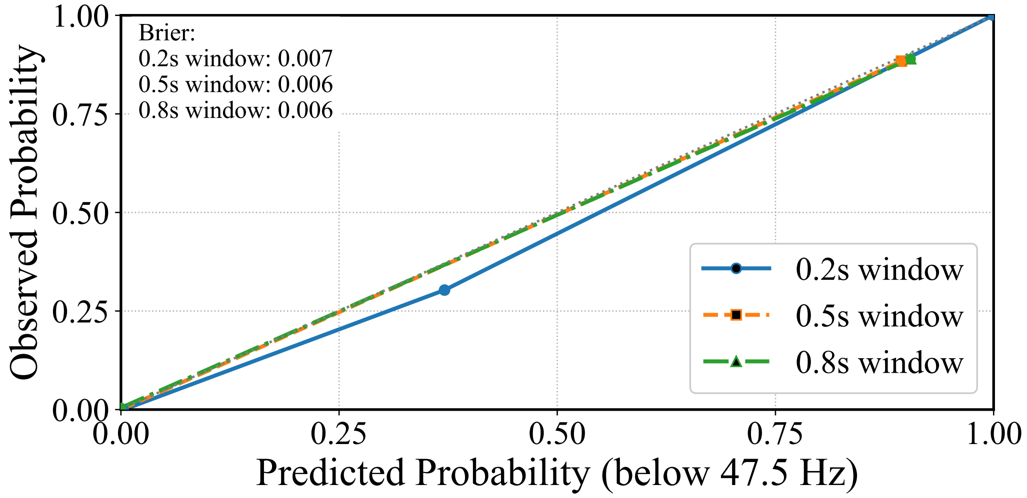

Fig. 5.4.2: Reliability of the threshold-crossing classifier under different observation window sizes

Fig. 5.4.3 illustrates the ECDF of the absolute time error $|e|$ for the first-crossing regressor under the three observation windows. All curves rise steeply from the origin and approach near unity by $0.2 - 0.3$ s, indicating that large time errors are rare and the estimates are generally reliable. In the practically relevant range $x \in [0, 0.12]$ s, the nominal window have the strongest performance: its ECDF lies consistently above those of the short and long windows, providing the highest coverage for any tolerance x. At $x = 0.1$ s the coverage probabilities are approximately 96% (nominal), 95% (short), and 94% (long). A similar ordering holds at $x = 0.05$ s, whereas for $x \geq 0.2$ s the three curves nearly coincide.

The differences follow from the amount of temporal context available in each window. The short window uses fewer samples and is therefore more sensitive to measurement noise, whereas the long window may include early recovery behaviour, which can bias the estimate. The nominal window balance these effects, generating the most favourable ECDF profile. With a tolerance of 0.1 s (act 0.1 s before the predicted crossing), the probability of acting too late is about $4-6\%$. If at least 97% of actions are required on time, the guard band should be increased slightly. The ECDF therefore provides a direct mapping from any chosen guard band to the probability of on time action and indicates that the nominal window performs best in the key operating range.

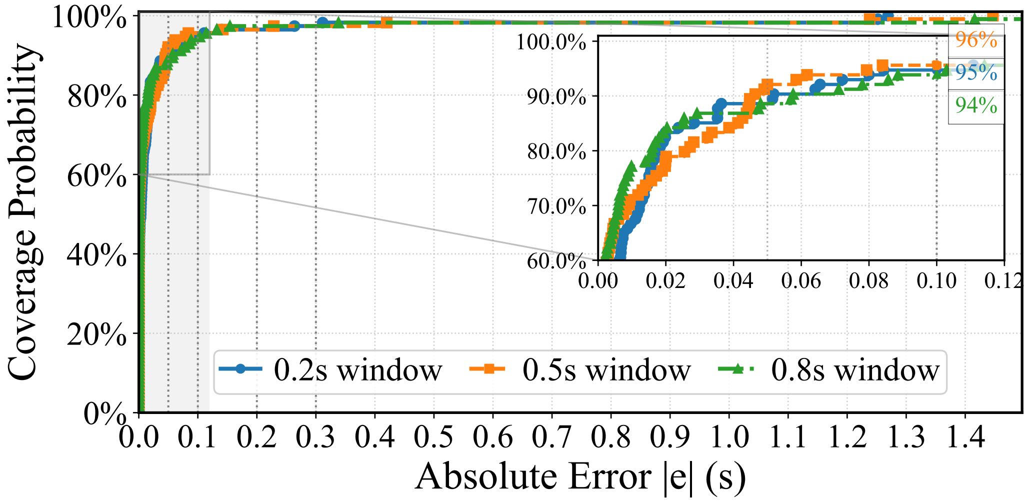

Fig. 5.4.3: Accuracy of the first-crossing time regressor under different window sizes

Fig. 5.4.4 evaluates the short-horizon trajectory predictor at a 1 s prediction horizon using features from the 0.2 s window. The upper panel compares predicted and observed frequencies with the line of equality, with points tightly concentrated around the diagonal over the full $40-50$ Hz range. No curvature or saturation is visible, indicating a highly linear and well-behaved mapping. The lower panel shows the residuals, which are narrowly distributed and centred near zero, revealing small errors and no systematic bias. These results demonstrate that the predictor produces accurate and unbiased prediction, making it suitable for real-time use within the emergency frequency control.

Fig. 5.4.5 evaluates the RF frequency nadir predictor used to determine the load shedding requirement. The plot compares predicted nadirs with observations on the test set and shows tight alignment with the line of equality. The residuals are narrowly concentrated around zero and display no trend with the observation level, indicating low bias and stable variance across the operating range. These results demonstrate that the predictor provides reliable guidance for determining the shedding step.

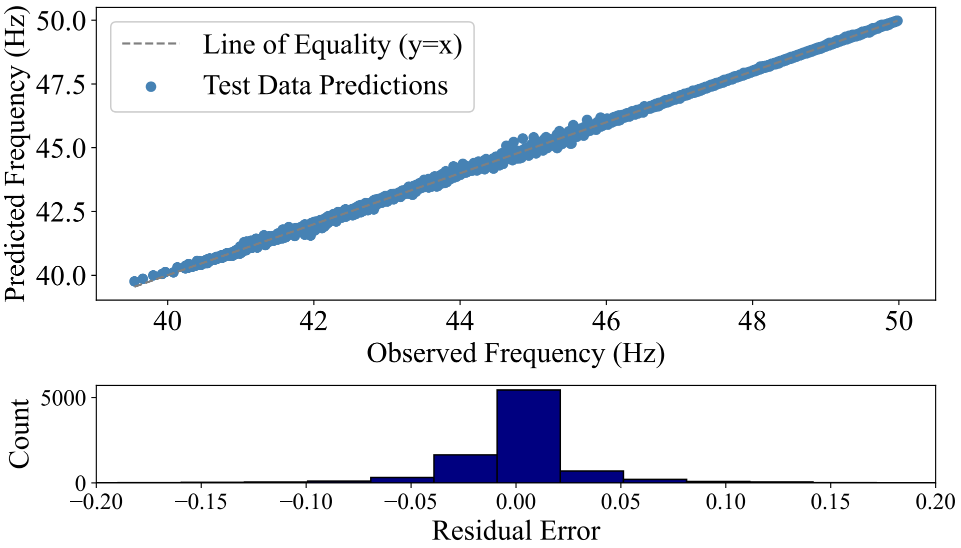

Fig. 5.4.4: Evaluation of a short-horizon frequency trajectory predictor

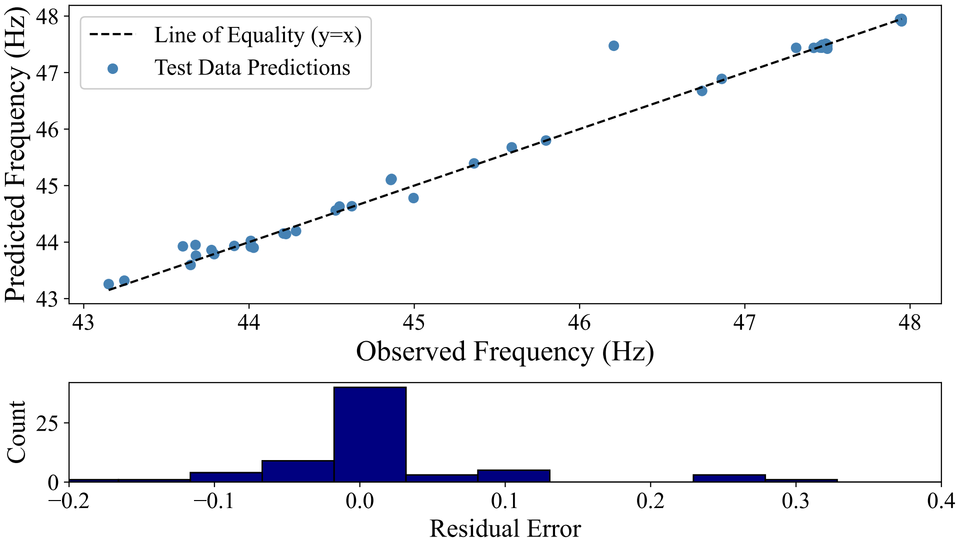

Fig. 5.4.5: Prediction accuracy of the load shedding adjustment model

The three learning modules deliver accurate and low bias performance across all observation windows. The consistency of these results across a wide range of operating conditions also highlights the robustness of the proposed framework. Since the training and validation datasets include diverse system states with varying demand levels, generation mixes, and disturbance magnitudes, the models develop a strong generalisation capability rather than being limited to specific operating points.

With respect to input sensitivity, the high accuracy observed in these evaluations is rooted in the models' ability to capture the underlying physical principles of power system dynamics. Specifically, the disturbance size and system inertia emerge as the most influential factors, as they directly determine the magnitude and rate of frequency deviation. While the impact of network topology is efficiently accounted for through the CoI aggregation, short-term frequency features such as measured RoCoF play a crucial role in capturing the evolving state during transients. The use of multi-window structures further enhances this robustness by reducing sensitivity to measurement noise. Overall, the model behaviour aligns with well-established physical insights, ensuring that the predictive framework responds to critical inputs in a physically meaningful and reliable manner. These results support their deployment in the subsequent HiL experiments and demonstrate their suitability for real-time frequency stability assessment and adaptive emergency frequency control.

5.4.3 Case with No Load Shedding Required (Safe Operation)

To evaluate the performance of the frequency trajectory predictor under realistic operating conditions, 10 disturbance scenarios of varying severity were generated. These scenarios range from low to high disturbance magnitudes.

Table 5.4.1 summarises the actual and predicted nadirs along with the MAE and TNAE values. Errors remain small across all disturbance levels, demonstrating accurate nadir prediction throughout the safe region.

To further assess the trajectory predictor, three representative scenarios of increasing severity were selected. Fig. 5.4.6 compares predicted trajectories (point) with measured frequency responses for three cases: case 1 (low), case 4 (medium), and case 7 (high). After event detection, the predictor issues update every 20 ms and provide predictions up to 0.5 s. In low severity scenario (case 1), deviations between predicted and measured trajectories are negligible, showing that the model captures small dynamics accurately.

Table 5.4.1: Prediction accuracy of frequency nadir for selected disturbance scenarios

Case	Actual Nadir (Hz)	Predict Nadir (Hz)	MAE (Hz)	TNAE (%)
1	**49.86**	**49.85**	**0.01**	**0.42**
2	49.60	49.61	0.01	0.48
3	49.42	49.39	0.03	1.56
4	**49.28**	**49.29**	**0.01**	**0.56**
5	49.13	49.15	0.02	1.23
6	49.05	49.02	0.03	1.94
7	**48.89**	**48.91**	**0.02**	**1.44**
8	48.79	48.77	0.02	1.55
9	48.62	48.65	0.03	2.68
10	48.56	48.58	0.02	1.89

With medium severity (case 4), small but acceptable differences appear while accuracy remains adequate for decision making. For high severity (case 7), deviations increase with dynamic complexity. However, the trend is captured, and the nadir stays at or above 48 Hz. In all scenarios, emergency frequency control is not triggered.

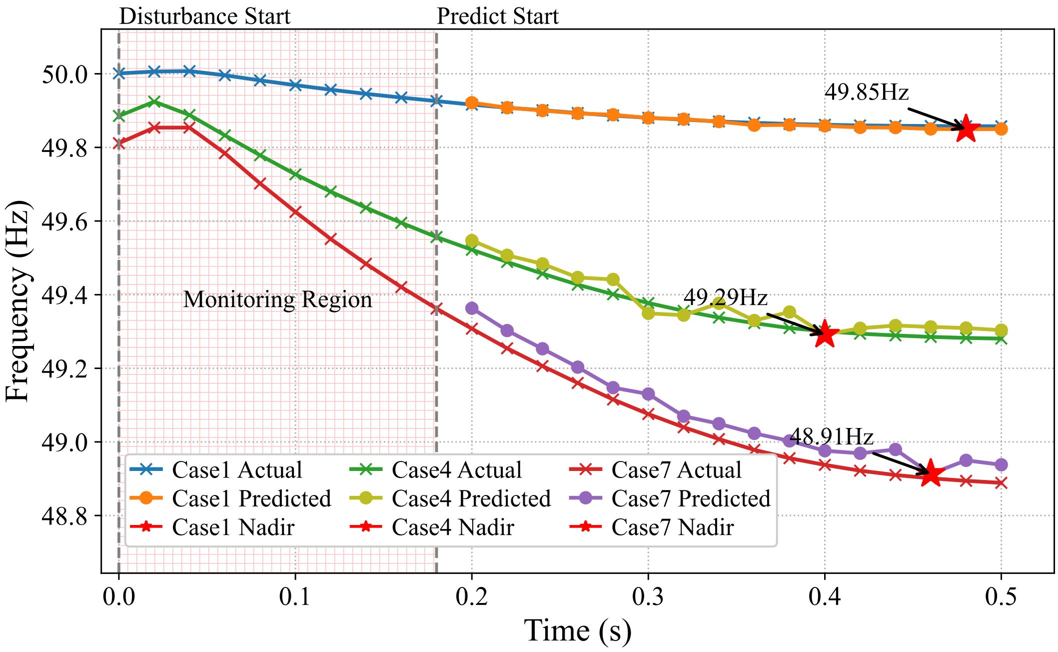

Fig. 5.4.6: Comparison between measured and predicted frequency trajectories under selected disturbance scenarios

Overall, the trajectory predictor delivers accurate and timely prediction across the full range of disturbance severities. These results support its use for real-time decision support in proactive grid control, enhancing situational awareness.

5.4.4 Unsafe Contingency Requiring Minimum Load Shedding

This section considers under-frequency contingencies in which the predicted nadir falls below 48 Hz and corrective action through emergency load shedding is required. Real-time measurements from the PMU stream into a database that feeds both the RTDS simulation and the controller. A disturbance is declared when the RoCoF exceeds 0.05 Hz/s, after which the controller enters multi-window decision process.

Following detection, the load shedding controller retrieves the instantaneous frequency, system inertia, and demand. Predictions are generated at three decision windows using pretrained models. Each prediction provides a short-term trajectory, an associated nadir value and its timing. When the predicted nadir is below 48 Hz, the control logic determines the minimum number of load units required to maintain the post event nadir at or above the 47 Hz safety boundary and prepares the corresponding command.

Table 5.4.2 summarises the results of 10 representative contingency events evaluated with the proposed prediction based control method. The results demonstrate effective corrective action across a range of disturbance magnitudes and inertia levels. Fig. 5.4.7 and Fig. 5.4.8 illustrate detailed comparisons of frequency trajectories for selected case studies (case 5 and case 8).

Fig. 5.4.7 illustrates a low-load contingency with a generation imbalance and demonstrates the UFLS workflow driven by the $\hat{t}_\theta$ predictor. The shaded regions denote three progressively expanding decision windows used for risk assessment and scheduling. Within each window, gated frequency measurements and system operational data are used to estimate the probability of breaching 47.5 Hz and, where relevant, to predict the first crossing time $\hat{t}_\theta$. Combined with the locally estimated RoCoF, $\hat{t}_\theta$ defines a 47 Hz safety deadline that guides the dispatch timing. In this case, the scheme manages the event by updating predictions sequentially across the three windows.

During the initial short-term assessment, the predicted frequency nadir remains

Table 5.4.2: Minimum load shedding requirements derived from multi-window prediction (10 test cases)

No.	P_{loss} **(MW)**	$f_{\text{nadir}}^{\text{without-shed}}$ **(Hz)**	Load$_{\text{pred}}$	$f_{\text{nadir}}^{\text{shed}}$ **(Hz)**	**Decision Window (s)**
1	1.61	44.95	5	48.08	0.8
2	1.63	44.89	5	47.99	0.8
3	1.66	44.81	5	47.88	0.8
4	1.71	44.23	5	47.80	0.8
5	1.75	44.54	3	47.66	0.5
6	1.78	44.48	3	47.62	0.5
7	1.83	43.59	4	47.80	0.8
8	1.87	44.25	4	47.92	0.5
9	1.92	43.75	2	47.88	0.2 deadline
10	1.96	43.47	2	47.91	0.2 deadline

Note: Each load unit is equivalent to 0.3 MW.

above critical thresholds, indicating that immediate intervention is unnecessary. As the prediction horizon is extended, the predicted trajectory nears the statutory limit, motivating continued observation rather than premature action. In the final decision window, the predicted trajectory clearly indicates a violation of the 47 Hz limit in the absence of corrective measures, triggering the UFLS decision-making mechanism at approximately 1.3 s to determine the required load shedding action.

The figure further contrasts the system response with and without load shedding to illustrate the effect of the proposed strategy. Without corrective action, the frequency trajectory breaches the 47 Hz statutory limit, whereas the load shedding level identified by the predictive scheme (2 units in this case) maintains the system frequency above the critical threshold. This comparison confirms that the UFLS strategy can identify a sufficient load shedding magnitude to ensure regulatory compliance while avoiding excessive intervention.

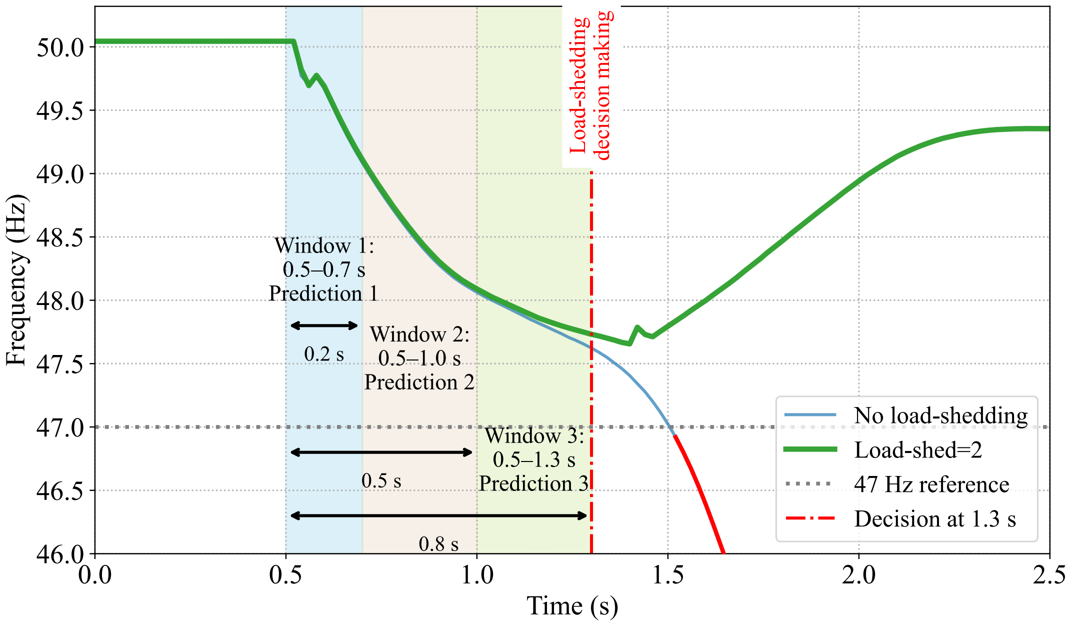

Fig. 5.4.7: Effect of alternative load shedding levels on frequency nadir (case 5)

Fig. 5.4.8 illustrates a severe under-frequency contingency. Without load shedding, the frequency trajectory would reach a nadir of approximately 43.5 Hz. As the observation window expands, the multi-window predictors progressively identify the emerging risk, with the final window confirming an inevitable violation of the 47 Hz limit in the absence of intervention. Based on this update and the locally estimated RoCoF, the controller schedules the load shedding command at $t = 1.23$ s, ensuring sufficient lead time before the anticipated threshold crossing. Applying the selected action (5 load units) maintains the frequency trajectory above 47 Hz, resulting in a controlled nadir of 47.45 Hz, whereas the unmitigated case would violate the constraint. The short red segment indicates the interval over which the uncontrolled trajectory would breach the limit. This case illustrates that the proposed scheme schedules the earliest feasible intervention using the minimum discrete shedding level required to satisfy the frequency constraint.

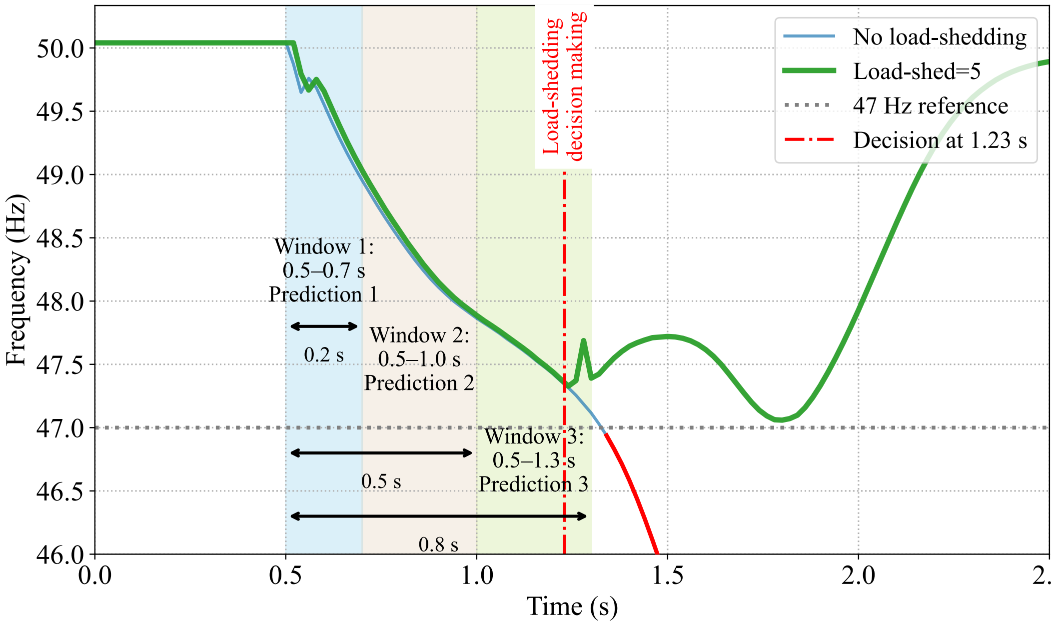

Fig. 5.4.8: Effect of alternative load shedding levels on frequency nadir (case 8)

Across all scenarios, the proposed scheme performs desirable optimisations simultaneously, curtailment cost is minimised, because the shedding point is raised only to the extent strictly required to meet the nadir constraint. Security margin is standardised, leading to a predictable post-event nadir slightly above 47 Hz regardless of disturbance magnitude, a property advantageous for system operators when coordinating with generator frequency response commitments.

The case study demonstrates that an online, multi-window frequency trajectory predictor can deliver load shedding actions that are at once adequate, necessary, and disturbance adaptive. Consequently, the controller avoids the conservatism associated with fixed step UFLS schemes and provides a credible pathway toward risk informed and cost effective protection for modern low-inertia power systems.

5.5 Summary

This chapter has presented methods for real-time, multi-window frequency prediction to enable emergency frequency control actions. The approach integrates three modules, i.e. an RF-based estimator for frequency nadir and nadir time, an XGBoost predictor for short-term frequency trajectories, and an RF model for load shedding adjustment. All components can operate with standard PMU measurement resolutions (i.e. 20 ms for 50 Hz systems) using multiple observation windows on the scale of hundreds of milliseconds to balance responsiveness and reliability. Shorter windows provide early situational awareness, whereas longer windows stabilise the estimates as more post-disturbance information becomes available.

Case studies on a realistic microgrid show that the nadir estimator and trajectory predictor maintain high accuracy across operating conditions. When frequency events that lead to the nadir remaining above 48 Hz, errors are small (TNAE $< 3\%$) and consistent across window sizes, confirming broad applicability. When a violation of the security limit (e.g. 47 Hz) is predicted, the control module selects the minimum integer load reduction that keeps the nadir above the limit, avoiding both insufficient and excessive shedding. In severe events, the multi window structure enables earlier intervention, with decisions made in the second window even though the nadir occurs in the third. This demonstrates that lead time can be exchanged for prediction precision to ensure system stability.

Overall, the proposed emergency frequency control during for extreme frequency events maintains a narrow and consistent post event margin above the critical frequency thresholds while limiting load reduction to the minimum required. These results indicate that the proposed approach can serve as a risk informed and effective solution to safeguard the security of low-inertia power systems under severe power imbalance events.

Chapter 6

Conclusion and Future Work

6.1 Conclusion

The large-scale integration of IBRs and the progressive displacement of conventional SGs have fundamentally reshaped the dynamic landscape of modern power systems. This transition has led to a significant reduction in system inertia and increasingly complex frequency dynamics, which creates major challenges for maintaining frequency stability and safeguarding the security of energy supply. Traditional frequency control frameworks, relying on simplified analytical models and offline simulations under fixed contingency assumptions, are no longer adequate to manage such highly volatile and complex system behaviour. Existing methods for estimating critical system parameters, e.g. RoCoF and power imbalance, also face substantial limitations and cannot effectively support frequency control actions. Furthermore, most existing control frameworks remain predominantly reactive, being triggered only after significant frequency deviations have already occurred, and therefore often lack the dynamic responsiveness needed to arrest the sharp frequency excursions characteristic of low-inertia grids.

Learning-based techniques provides a promising research direction for addressing a variety of challenges resulted from the massive integration of IBRs. However, existing research has largely concentrated on direct control implementations or narrowly defined prediction tasks, with comparatively limited attention given to However, existing research has largely focused on direct control implementations or narrowly scoped prediction tasks, with comparatively limited attention given to the development of proactive, situation-aware, and practically deployable decision-support methodologies. Consequently, a clear knowledge gap remains in translating learning-based estimation and prediction advances into effective real-time operational support for low-inertia power systems.

In response to this gap, this thesis makes original contributions to knowledge by developing a set of learning-based methodologies for frequency stability support in low-inertia power systems, explicitly designed to strengthen existing control and operational practices rather than replace them entirely. These contributions are organised around three interrelated research themes. The first concerns probabilistic nadir prediction and frequency response scheduling, through which the thesis demonstrates how learning-based models can capture non-linear operating dependencies, quantify uncertainty in post-disturbance security assessment, and support risk-informed and economically efficient procurement of frequency response services. The second concerns real-time estimation of RoCoF and power imbalance size, where the thesis develops measurement-driven estimators capable of extracting reliable disturbance information from raw PMU data under noisy and rapidly evolving system conditions, which provide more accurate information for critical frequency control actions. The third concerns learning-based prediction of frequency behaviour and the emergency frequency control actions that such prediction enables, in which the thesis formulates a multi-window predictive structure to support adaptive and time-critical control actions during severe disturbances. The principal contributions of this thesis are summarised in detail below.

Probabilistic Nadir Prediction and Response Scheduling

To support frequency security during the operational schedule stage, this thesis developed a learning-based nadir prediction approach designed to capture the non-linear dependencies among operating conditions. An RF estimator maps pre-event features to post-event nadirs with high accuracy. Feature importance analysis shows that equivalent inertia remains the predominant influence on nadir behaviour, consistent with swing equation physics.

The approach extends deterministic prediction to a probabilistic form by simulating variability in key features, particularly inertia. The resulting prediction interval is well calibrated, which enables operators to assess the probability of adverse events rather than relying on a single point estimate. When incorporated into an ancillary services scheduling model, this probabilistic capability improves economic efficiency by allocating frequency ancillary services only where it is most effective. In the extended tested scenarios, the proposed approach maintained nadirs above critical limits (i.e. 49.5 Hz in this study) while avoiding unnecessary ancillary services procurement, thereby demonstrating clear value for low-inertia systems in which excessive over-provisioning is increasingly costly.

Real-Time Estimation for RoCoF and Power Imbalance Sizes

Accurate situational awareness immediately after a disturbance is essential for effective frequency control. To meet this requirement, the thesis developed learning-based methods for real-time estimation of RoCoF and power imbalance, which are two critical parameters widely used to enable frequency control actions.

For RoCoF estimation, traditional RoCoF measurement methods often suffer from inherent time delays and errors caused by the adverse effects of noise filtering, particularly during the early stages of a disturbance. Rather than proposing fundamental modifications to existing hardware or measurement algorithms, this research introduces an approach that operates directly on the raw, noisy RoCoF measurement output from a PMU. This DWR-SVR approach enables effective calibration of the raw signal and significantly reduces the estimation delay compared with conventional low-pass filtering methods. In addition, the integration of the MOS mechanism allows the observation window length to be adaptively adjusted, thereby improving the balance between estimation speed and accuracy under different disturbance conditions.

For estimation of power imbalance size, the thesis developed an XGBoost-based estimator to address the limitations of conventional swing equation-based approaches. In low-inertia systems with rapidly varying inertia conditions, analytical methods frequently produce inaccurate estimates of the actual disturbance size because they depend on precise knowledge of equivalent system inertia and RoCoF. The proposed estimator overcomes this limitation by inferring power imbalance directly from the initial post-disturbance frequency trajectory, thereby avoiding the need for explicit and accurate inertia information. In this way, the thesis establishes a more robust and practically applicable approach for real-time disturbance quantification in future low-inertia systems.

Real-Time Prediction of Frequency Behaviour and Emergency Frequency Control

To prevent system instability during severe contingencies, this thesis developed learning-based predictors that allow real-time prediction of frequency behaviour, thereby enabling a more effective emergency frequency control scheme that replaces conventional reactive, fixed-threshold methods with an adaptive multi-window structure. In low-inertia power systems, traditional reactive under-frequency load shedding schemes often act too late to arrest rapid frequency decline, or alternatively shed excessive load because of limited predictive capability. To overcome these limitations, the proposed scheme employs progressively increasing observation windows to explicitly address the balance between response speed and prediction accuracy. Shorter windows provide early situational awareness and enable rapid disturbance detection, while longer windows improve prediction reliability as more post-disturbance information becomes available.

Within this structure, an RF estimator is used to predict frequency nadir and nadir time, an XGBoost model is employed to predict short-term frequency trajectories, and an additional RF-based model is used to determine the required adjustment in load shedding. By integrating these predictive components with a decision making module, the proposed scheme identifies the minimum necessary load shedding required to maintain system stability within defined security limits. The operational feasibility of the method was further validated through real-time HiL simulation studies. These results confirmed that the full prediction and decision-support process can be executed within the timing constraints of modern power system operation, thereby demonstrating its suitability for practical emergency frequency control applications.

In conclusion, the contributions of this PhD research advance fundamental understanding of how learning-based estimation and prediction can be integrated across operational scheduling, real-time situational awareness, and emergency control to achieve enhanced frequency control in low-inertia power systems. They support a transition in frequency control from reactive and model-constrained approaches towards predictive, uncertainty-aware, and operationally compatible decision support. Therefore, the thesis establishes a learning-based methodological foundation for more effective frequency stability management in future power systems.

6.2 Future Work

As noted in the previous section, this thesis developed learning-based frameworks for frequency security that spans pre-event scheduling, real-time situation awareness, and emergency control. Although the proposed methods have demonstrated strong performance, several important directions remain for further investigation. These relate both to the limitations of the current methodologies and to the broader challenge of translating simulation-based research into practical deployment. The following areas therefore represent key priorities for extending this work.

6.2.1 Advanced Modelling and Physics - Informed Architectures

The current framework relies primarily on established ML algorithms, e.g. RF and XGBoost. Although these methods have proven effective, they do not explicitly embed physical system structure within the model formulation, which limits their interpretability. Future research should therefore investigate hybrid architectures that combine learning-based methods with power system physics. For example, physics-informed neural networks could be developed to incorporate swing equation constraints directly into the training process, thereby improving generalisation under unseen operating conditions and enhancing the physical consistency of model predictions. In addition, the input feature set could be extended to include dynamic state indicators derived from real-time measurements, such as damping-related characteristics and RoCoF trajectories, which may further improve the predictive performance of nadir estimation and disturbance quantification models.

6.2.2 Alternative Frequency Reference Selection

The present framework primarily relies on system-level frequency representations derived from aggregated models and microgrid-level measurements. While this approach is well aligned with system-level decision making, future work could investigate the use of alternative frequency reference signals. In particular, the frequency measured at the largest connected synchronous generator may provide a more stable and inertia-dominated representation of system dynamics. Such a reference could be less affected by localised network oscillations and measurement artefacts. A systematic comparison between system-level frequency measures and generator-based frequency references would therefore be valuable in assessing their respective impacts on estimation accuracy and real-time decision support performance.

6.2.3 Comprehensive Evaluation of Impact of Synthetic Inertia

The current study primarily considers the rotational inertia provided by synchronous machines and the frequency support from existing ancillary services. However, IBRs have been increasingly required to provide synthetic inertia, which significantly alters the inertial characteristics of the power system. While inertia emulation was included in the case study to a limited extent, future work should more comprehensively examine the impact of fast-acting, non-rotational responses on frequency security assessment. Furthermore, since the performance synthetic inertia is often dependent on control loop gains and local voltage conditions, the existing learning-based models should be further adapted in order to account for rotational and virtual inertial contributions under different system operating conditions and controller settings.

6.2.4 Adaptive Scheduling and Economic Dispatch Integration

The scheduling formulation developed in Chapter 3 currently uses linear effectiveness coefficients ($\beta_{\mathrm{DC}}, \beta_{\mathrm{Static}}$) to quantify the frequency support delivered by response services. In practice, however, the effectiveness of these services is inherently non-linear and depends on prevailing system conditions. Future work should therefore focus on learning these gains adaptively, so that they can be expressed as dynamic functions of variables such as system inertia, disturbance size, and demand level, rather than as fixed coefficients. Furthermore, the current scheduling model could be extended to a probabilistic or robust economic dispatch formulation that explicitly account for uncertainty. Practical factors, e.g. activation delays and ramp rate limits, should also be considered so that such a formulation could enable frequency security requirements to be co-optimised with generation cost, thereby providing a more integrated and economically grounded operational framework.

6.2.5 Uncertainty Quantification and Data Reliability

The probabilistic nadir prediction approach presented in Chapter 3 explicitly accounts for uncertainty. However, the real-time situation awareness tools in Chapter 4 currently provide only point estimates. To improve operator confidence and practical applicability, future research should extend these methods to include uncertainty quantification. Confidence intervals for quantities, e.g. RoCoF and power imbalance, would allow operators to distinguish more reliably between genuine disturbances and transient measurement artefacts, thereby supporting more robust real-time decision making.

In addition, further work is needed to assess model performance under degraded data conditions. Practical PMU-based monitoring systems may be affected by missing measurements, communication delays, timing errors, or corrupted signals. Future studies should therefore investigate the resilience of the proposed methods under such conditions and develop suitable data cleaning, reconstruction, and fault tolerant estimation strategies. Addressing these issues will be essential to ensure that the framework remains reliable when deployed in imperfect real world monitoring environments.

6.2.6 Hardware Testing and System Scaling

While limited HiL and real-time simulation studies were conducted in this PhD to validate the performance of the proposed methods, further work is needed to support their transition towards real world application. Future research should therefore focus on more comprehensive and realistic hardware prototyping and testing. Implementing the developed DWT-SVR, XGBoost, and RF models on representative substation or control hardware would provide valuable insight into their computational requirements, execution latency, and real-time feasibility under practical operating constraints.

In parallel, the scalability of the proposed emergency frequency control scheme should be investigated on larger and more complex transmission-level systems. As future power systems become increasingly interconnected and heterogeneous, it is necessary to verify that the proposed control logic remains stable, effective, and computationally tractable when applied to wide-area networks with diverse grid topologies and dynamic characteristics. Such studies would provide a stronger basis for assessing the transferability of the framework beyond the smaller-scale systems considered in this thesis.

In summary, future research should aim to progress from the development of individual predictive tools towards a fully integrated, adaptive, and operationally robust framework for frequency security support. This would address the modelling, uncertainty, deployment, and scalability challenges identified above, which allows the proposed methods to be further advanced to support the secure and stable operation of real world low-inertia power systems.

Appendix A

Technical Characteristics of PMU Measurements

PMUs are critical to the proposed framework as they provide high-speed, time-synchronised measurements of voltage and current phasors [1]. The operational process begins with the high-frequency sampling of analogue waveforms. Each voltage and current sample is synchronised to a common time source, typically the global positioning system, ensuring a precise universal time coordinated timestamp.

Following synchronisation, the PMU employs a sliding window approach, such as the DFT or recursive algorithms, to estimate the fundamental frequency components, including voltage magnitudes and phase angles [2]. From these estimated phasors, the bus frequency is derived based on the rate of change of the voltage phase angle. The RoCoF is subsequently calculated as the time derivative of the estimated frequency. Consequently, these synchronised frequency-related profiles and raw technical quantities serve as the primary foundational inputs for the learning-based models proposed in this thesis.

While PMUs offer high reporting rates, conventional estimation algorithms exhibit inherent accuracy limitations during rapid system transients, which directly motivates the RoCoF estimation framework proposed in Chapter 4.1. Standard PMU algorithms optimised for steady-state compliance typically introduce significant reporting delays, often in the range of 80 ms to 100 ms. Under fast frequency excursions following a severe disturbance, the fixed-window processing utilised in conventional PMUs leads to increased estimation errors, phase lags, and numerical oscillations. This degradation in measurement accuracy hampers the reliability of frequency-related indicators during the critical initial transient period. To mitigate these limitations without introducing the heavy phase lags associated with conventional low-pass filtering, this thesis leverages learning-based models. By capturing the underlying non-linear relationship between the noisy, transient PMU outputs and the actual system dynamics, the proposed approach restores the physical accuracy of the signals and drastically reduces time delays, thereby providing timely and robust indicators for emergency control decisions.

References for Appendix A

[1] R. F. Nuqui and A. G. Phadke, “Phasor measurement unit placement techniques for complete and incomplete observability,” *IEEE Transactions on Power Delivery*, vol. 20, no. 4, pp. 2381–2388, 2005.

[2] M. K. Penshanwar, M. Gavande, and M. R. Satarkar, “Phasor measurement unit technology and its applications-a review,” in *2015 International Conference on Energy Systems and Applications*, IEEE, 2015, pp. 318–323.