Physics-Informed Graph Neural Networks for Power-Grid Stability

by

Caleb Kojo Eshun

Thesis submitted to the University of Ottawa
in partial fulfillment of the requirements for the
Master of Applied Science

School of Electrical Engineering and Computer Science
Faculty of Engineering
University of Ottawa

Declaration of Authorship

I hereby certify that this thesis is entirely my own original work except where otherwise indicated. I am aware of the University's regulations concerning plagiarism, including those concerning consequent disciplinary actions. Any use of the works of any other author, in any form, is properly acknowledged at their point of use.

Abstract

The transition of modern transmission grids toward inverter-based generation is steadily eroding the rotational inertia that has historically anchored small-signal stability. As synchronous machines are displaced, frequency deviations evolve faster, damping margins shrink, and the network becomes increasingly fragile under topology changes, contingencies, and the operational stress imposed by extreme weather events. At the same time, the streaming data now available from modern grid sensors is largely underused because prediction, stability assessment, and control are still treated as separate pipelines rather than as a single decision process. This thesis closes that gap by developing a unified, closed-loop framework for online stability support in low-inertia, topology-varying power grids, addressing three coupled objectives: online stability assessment under uncertainty, adaptive virtual inertia allocation under topology change, and integration of prediction with control under formal stability guarantees. At the core of the framework is a spectral-sensitivity-based virtual inertia allocator, derived from swing-equation dynamics and eigenvalue perturbation analysis, which identifies the buses whose inertia most strongly influences the damping of critical electromechanical modes and concentrates allocation at those locations. A spatio-temporal physics-informed graph neural network (ST-GNN) forecasts near-future grid states and reproduces these allocation targets, with Kirchhoff's laws and operational limits enforced as physics-informed penalties during training. A reinforcement-learning policy then issues coordinated control actions spanning generator re-dispatch, demand response, and inertia updates, and a supervisory coordination layer blends these actions with the spectral reference and projects them onto the feasible set prior to execution. A Lyapunov-based analysis establishes that, provided inertia stays within design bounds and the network remains connected between switching events, the closed-loop electromechanical dynamics preserve uniform asymptotic stability. Validation on the IEEE 118-bus system under N–2 contingencies not seen during training demonstrates accurate state tracking, topology-aware inertia redistribution, and sustained operational security across extended simulations, while saliency-based diagnostics expose the structural drivers of each control decision. The result is an interpretable, computationally efficient framework that brings physics-based guarantees and learning-based adaptability into a single coherent control loop for stability-aware operation of future smart grids.

Acknowledgements

This work would not have been possible without the mentorship and full support of my supervisor, Professor Javad Fattahi. His expertise in power-system stability, control under uncertainty, and the integration of physics-informed methods with machine learning shaped the foundation on which this thesis is built, and his guidance is reflected in every chapter of this work. Professor Fattahi is rigorous in his thinking and generous with his time. He has built an environment in which careful theory, strong engineering, and clear writing are held to the same high standard, and he never once let me settle for less than I was capable of. His patience with me throughout this journey, the harder questions he asked when I needed them, and his commitment to seeing the work through with me are things I will carry long after this thesis is bound. Professor Fattahi will remain my role model in research and a standard I aspire to throughout my career.

To my colleagues at the University of Ottawa, thank you for the long conversations and shared work that made the difficult stretches feel lighter. I also extend my sincere gratitude to the faculty of the School of Electrical Engineering and Computer Science for their guidance throughout this study.

Dedication

To my father, Rev. Prof. John Frank Eshun, who invested in this dream with quiet faith when others would not. His belief in me was worth more than any award could ever be.

To my mother, Mrs. Josephine Eshun, whose love and steady support held me together through it all.

To Pastor Joseph Amoah, my mentor, whose words found me on the days when doubt had taken hold. He reminded me that the work was always worth finishing.

And to my family and loved ones, whose prayers carried me through every step of this journey.

Table of Contents

Publications Arising from This Thesis

The research presented in this thesis has resulted in the following publications:

1. C. K. Eshun, N. Fatemi, and J. Fattahi, "Spectral sensitivity and physics informed GNN-RL for real time power grid stability," *Sustainable Energy, Grids and Networks*, vol. 46, Art. no. 102168, 2026, doi: 10.1016/j.segan.2026.102168.

2. C. K. Eshun, S. Bahrami, and J. Fattahi, "Robust local contingency planning in power grids using spatio-temporal graph neural networks," in *Proceedings of the 2025 IEEE Electrical Power and Energy Conference (EPEC)*, Waterloo, ON, Canada, 2025, pp. 152–156, doi: 10.1109/EPEC65543.2025.11230436.

List of Tables

List of Figures

Nomenclature

Symbol	Description
Index	
n	Bus number
t	Timestep
k	Eigenmode index
$i,\ j$	Bus indices
τ	Temporal index
ℓ	Loop index
Sets	
$\mathcal{G}$	Spatio-temporal graph
$\mathcal{V}$	Set of buses
$\mathcal{E}$	Set of edges
$\mathcal{H}$	Feasible set of inertia allocations
$\mathcal{K}$	Set of critical (non-rigid) eigenmodes
$\mathcal{N}(i)$	Set of buses adjacent to bus i
Parameters	
T	Length of the prediction horizon
$\eta \in [0, 1]$	Blending parameter
ρ_{s}	Spatial smoothing weight
γ	Discount factor
$\Theta_V,\ \Theta_\delta$	GNN model parameters
ω_s	Synchronous electrical angular speed
S_{base}	System MVA base
J	Rotor moment of inertia

Symbol	Description
D_i	Effective damping coefficient at bus i
K_{f_i}	Frequency-dependent load coefficient at bus i
K_{v_i}	Voltage-dependent load coefficient at bus i
B_{ij}	Susceptance of line (i, j)
$V_{\min},\ V_{\max}$	Lower and upper voltage operating bounds
$\delta^*,\ \omega^*$	Equilibrium operating point (rotor angle and speed)
Variables	
A_{ij}	Connectivity of buses i and j
D_{ii}	Degree of bus i
L	Combinatorial graph Laplacian
$L_B(t)$	Susceptance-weighted Laplacian
$K(t)$	Synchronizing coupling matrix
$A,\ B$	State and input matrices of the linearized system
$P^n_{d,t:t+T}$	Active power demand at bus n over horizon T
$Q^n_{d,t:t+T}$	Reactive power demand at bus n over horizon T
$P^n_{g,t}$	Active power generation at bus n
$Q^n_{g,t}$	Reactive power generation at bus n
$P_{m_i},\ P_{e_i}$	Mechanical input and electrical output power at bus i
$T_m,\ T_e$	Applied mechanical and electromagnetic torque
$\bar{M}^n_t$	Virtual inertia at bus n at time t
H_i	Inertia constant at bus i
ΔH^n_t	Inertia update at bus n
$\theta(t)$	Vector of bus voltage phase angles
$\dot{\theta}(t)$	Bus frequency deviation vector
$\delta_i,\ \omega_i$	Rotor angle and frequency deviation at bus i
$\Delta\delta,\ \Delta\omega$	Small-signal angle and frequency deviations
x	State vector $[\Delta\delta^\top\ \Delta\omega^\top]^\top$
u_i	Active-power disturbance/control input at bus i
s_t	RL state
a_t	RL action
r_t	Reward signal
$V^\pi(s_t)$	Value function for the RL policy
α_τ	Attention weight at temporal index τ

Symbol	Description
$\mathcal{J}$	Power-flow Jacobian
$V_i^{(t+1)}$	Ground-truth voltage magnitude at bus i
$\hat{V}_i^{(t+1)}$	Predicted voltage magnitude at bus i
$\delta_i^{(t+1)}$	Ground-truth voltage angle at bus i
$\hat{\delta}_i^{(t+1)}$	Predicted voltage angle at bus i
$\hat{V}^{\pi}(s_t)$	Predicted value function
$\bar{H}_i^{(t+1)}$	Ground-truth virtual inertia at bus i
$\hat{\bar{H}}_i^{(t+1)}$	Predicted virtual inertia at bus i
$\mathcal{L}$	Total training loss
$\mathcal{L}_k$	KCL (Kirchhoff current law) loss
$\mathcal{L}_v$	Voltage-limit loss
$\mathcal{L}_{nd}$	Nodal-demand loss
$\mathcal{L}_{kv}$	KVL (DC power-flow consistency) loss
s_{ij}	Edge-level saliency score for line (i, j)
Operators	
$\mathrm{Re}(\cdot)$	Real-part operator
$\mathrm{diag}(\cdot)$	Diagonal-matrix construction
$\mathbb{E}[\cdot]$	Expectation operator
$\nabla_{\mathbf{z}}(\cdot)$	Gradient with respect to descriptor $\mathbf{z}$
$\|\cdot\|_2$	Euclidean norm
Metrics	
$\lambda_k(t)$	Eigenvalue of the state matrix associated with mode k
$\mathrm{Re}(\lambda_k)$	Real part of eigenvalue λ_k
ζ_k	Modal damping ratio of mode k

Chapter 1

Introduction

1.1 Motivation

The global energy landscape is undergoing a profound transformation driven by decarbonization mandates, technological advancement, and changing consumption patterns [1]. This evolution is characterized by the rapid displacement of conventional synchronous generation by inverter-based resources (IBRs), primarily wind and solar photovoltaic systems. While this transition is essential for sustainability goals, it fundamentally alters the dynamic behaviour of power systems [2].

Traditional power grids derived inherent stability from the rotational inertia of synchronous machines. The kinetic energy stored in their rotating masses provided an immediate physical response to power imbalances, slowing the rate of change in frequency and buying crucial time for control systems to activate [3]. The shift toward converter-dominated networks creates a low-inertia environment where this natural buffer is absent. Consequently, frequency and voltage become more volatile, leading to an increased Rate of Change of Frequency

(RoCoF) following disturbances, reduced damping of electromechanical oscillations, and narrower stability margins, rendering traditional assurance mechanisms inadequate [4–6].

This transformation is further complicated by the increasing frequency and severity of extreme weather events, which introduce unprecedented stress on grid infrastructure [7]. The convergence of low-inertia conditions, heightened volatility from renewable generation, and growing infrastructure vulnerability creates a critical need for advanced operational frameworks that can maintain stability online under uncertain and rapidly evolving conditions [8, 9]. These exogenous shocks affect both supply and demand, and their effects can unfold rapidly. As shown in Figures 1.1 and 1.2, wildfire smoke and airborne dust can reduce solar irradiance and depress photovoltaic output, while extreme cold events can drive sharp increases in heating demand, and extreme heat events can drive sharp increases in cooling demand, and stress system operating margins. Such conditions increase the likelihood of large power imbalances and abrupt network reconfiguration, highlighting the need for stability-oriented decision-making under non-stationary conditions.

The mechanisms in Figures 1.1 and 1.2 are introduced into the framework through specific modelling choices in the chapters that follow. Wildfire smoke and dust reduce solar output, and this effect is reflected in the random renewable generation profiles used to train and test the model in Chapter 4. Extreme cold drives up heating demand, which enters the system as time-varying load profiles for residential, commercial, and industrial buses, and is also the period during which the demand-response controller operates most. Wind and ice damage transmission lines, and these events enter as the line outages, including the N–1 and N–2 cases, that change the network topology and the coupling matrix $K(t)$ in Chapter 2. In this way, the same drivers illustrated in the two figures shape the uncertainty model, the contingency scenarios, and the online control needs developed throughout the rest of the thesis.

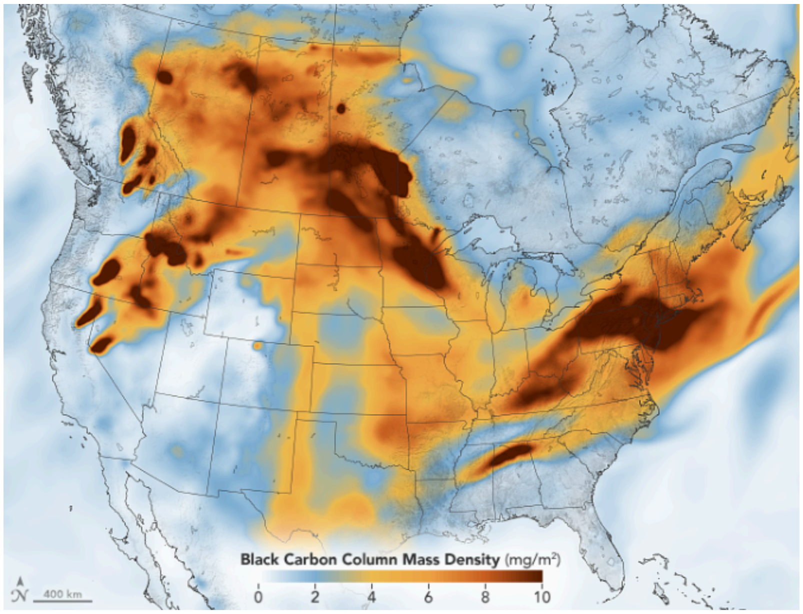

Figure 1.1: Wildfire smoke and airborne dust reduce solar irradiance across North America. Source: [10].

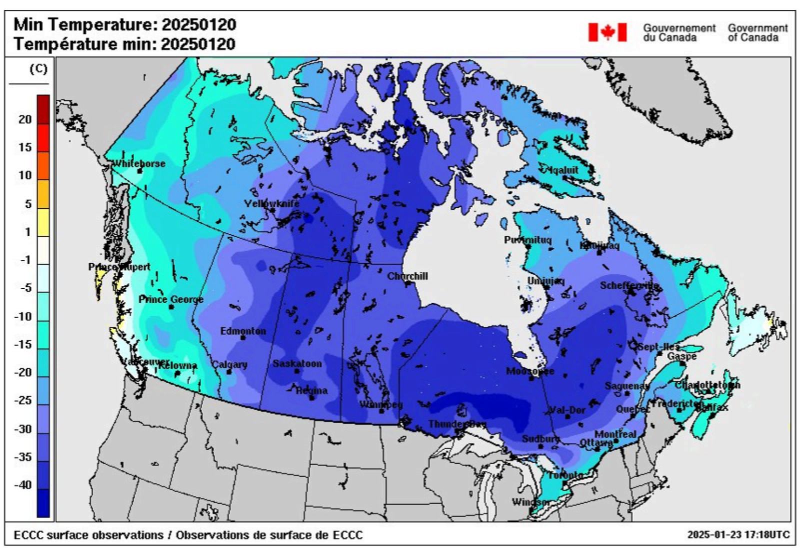

Figure 1.2: Extreme cold across Canada increases heating demand and stresses grid operation. Minimum temperature surface observations, 20 January 2025. Source: [11].

These conditions also carry an operational implication. Under non-stationary operation and changing topology, handling prediction, stability assessment, and control as separate steps leaves the system slow to respond and exposed to inconsistencies between the three. In existing work, the three are typically developed in isolation. Predictors are trained on a fixed network and say little about stability. Spectral and sensitivity methods provide stability margins at a single operating point but do not specify how to update virtual inertia as the system changes. Reinforcement-learning controllers can adapt online, but often ignore the network laws and limits required for safe operation. Each approach is useful, but on its own, none is enough.

The gap addressed in this thesis is therefore more specific than the common observation that prediction, stability assessment, and control are studied separately. The deeper issue is that very few existing frameworks bring together topology-aware stability prediction, spectral-sensitivity guidance for virtual inertia, and safety-constrained online control within a single closed-loop architecture that runs continuously as conditions and topology change. Closing this gap calls for a single architecture in which the predictor supplies short-horizon state information, the spectral-sensitivity module supplies a stability-oriented inertia reference, and the controller is kept inside the limits set by network laws and device capabilities. Building this architecture and showing that it remains stable under bounded inertia updates and admissible topology changes is the central technical objective of this thesis.

Modern power systems are increasingly instrumented with high-resolution sensors, such as Phasor Measurement Units (PMUs), providing streaming data at sub-second rates [12]. This creates unprecedented visibility into system dynamics. However, there is a significant disconnect between the availability of this data and our analytical capabilities. Traditional frameworks often assume quasi-steady-state conditions, employ linearized models around fixed operating points, and treat prediction, analysis, and control as sequential, disconnected processes, thereby missing an opportunity to enhance grid resilience [13].

Artificial intelligence and machine learning offer powerful tools for extracting insights from this complex, high-dimensional data and have demonstrated success in applications from load forecasting to fault detection [14, 15]. Yet, the direct application of conventional ML methods to stability-critical grid operations faces fundamental limitations that hinder their adoption. The "black-box" nature of many deep learning models creates trust deficits, as operators require an understandable rationale for control actions in safety-critical environments [16, 17]. Furthermore, models trained solely on historical data can produce predictions that violate fundamental physical laws, such as Kirchhoff's laws or power-flow constraints, leading to physically inconsistent and potentially dangerous control decisions [18, 19]. A final critical challenge is the difficulty ML models face in generalizing under the constantly changing conditions of a power grid, including topology changes, unusual operating conditions, and novel contingencies not seen during training [20, 21].

1.2 Thesis Objectives

The framework in this thesis is built around three challenges that follow directly from the conditions described above. The first is online stability assessment under uncertainty, capable of tracking operating points that evolve continuously with renewable variability, stochastic load behaviour, and changing network topology. The second is adaptive virtual inertia allocation, which involves deciding where to place limited inverter-based resources so they best damp the critical electromechanical modes, given topology changes and per-bus device limits. The third is the coupling of prediction and control within one closed-loop process, so that the actions remain physically consistent, carry a formal stability guarantee, and remain computationally efficient for sequential implementation. Together, these three points define the gap stated at the end of the previous section, namely that very few frameworks integrate topology-aware stability prediction, spectral-sensitivity-based inertia allocation, and safety-constrained learning-based control into a single continuously running closed-loop architecture [22, 23].

The thesis pursues four objectives that together address this gap, with each one tied to a target the framework can be checked against rather than only to a component that has been built.

The first objective is to build a fast method that uses spectral sensitivity to decide where to place virtual inertia. The target is an allocator that, for any valid network configuration, identifies the critical electromechanical modes, shows how bus-level inertia changes affect their damping, and produces a feasible inertia profile within budget and per-bus limits, fast enough to run at minute-scale control intervals on a standard transmission benchmark.

The second objective is to build a physics-informed spatio-temporal Graph Neural Network that predicts grid states as the topology and operating point change. The target is a predictor that matches AC optimal power flow values for voltage magnitudes, angles, and active and reactive power injections with low error under both N–1 and N–2 contingencies, respects Kirchhoff-based constraints in its training loss, and admits interpretation, through saliency analysis, of which buses and lines drive its predictions.

The third objective is to build a reinforcement-learning controller that selects coordinated actions across generator redispatch, demand response, and virtual inertia adjustments. The target is a controller that, over 24-hour rollouts with stochastic loading and contingency-driven topology changes, uses demand response sparingly and only where needed, tracks the spectral-sensitivity inertia reference near stressed corridors, and keeps voltages and angles within operating limits without large-scale load shedding.

The fourth objective is to build a coordinator that connects the controller to the spectral-sensitivity reference and to the network feasibility set. The target is a blending-and-projection coordinator, backed by a Lyapunov-based argument that the closed loop remains stable under bounded inertia updates and piecewise-constant, connected topologies, and supported by visualization tools that show the operator why each action was taken.

Together, these four objectives reduce the broad requirement of safe, topology-aware, online stability support to a small set of concrete targets that can be checked. The chapters that follow build the components needed to hit each target and then test them on the IEEE 118-bus benchmark, including under N–2 contingencies, which the model was not trained on.

The scope of this work is transmission-level power systems, focused on small-signal stability on time scales of seconds to minutes, where frequency stability and electromechanical oscillations matter most. The control actions covered are generator active and reactive power dispatch, demand response, and virtual inertia allocation, with disturbances coming from line outages, load variations, and inertia fluctuations. The work does not cover protection system design, electricity markets, distribution-system dynamics, fast electromagnetic transients, or long-term voltage stability.

This thesis makes several original contributions to the field of power system stability and control, which can be categorized into theoretical and methodological advancements. On the theoretical front, this work introduces a Spectral-Sensitivity Virtual Inertia Allocation Framework that leverages eigenvalue sensitivity analysis as both an input feature for predictive models and a direct control mechanism, enabling targeted allocation that adapts to topology changes. A second key contribution is a Stability-Guaranteed Coordination Architecture, a hierarchical control framework that blends reinforcement learning actions with physics-based reference signals and enforces operational constraints through projection, accompanied by formal proofs of uniform asymptotic stability. Finally, to enable topology-aware state forecasting within the closed-loop framework, we present a Physics-Informed Spatio-Temporal GNN Formulation, a novel graph neural network architecture that incorporates physical laws directly into the learning objective, ensuring physically consistent predictions while maintaining the expressive power of deep learning.

Methodologically, this research presents an Integrated Prediction-Control Pipeline that combines state prediction, stability assessment, and control action synthesis into a single cohesive architecture, thereby overcoming the traditional separation among these functions. We also develop novel Interpretability Tools for Grid AI based on edge saliency and gradient analysis, providing insight into how graph neural networks make stability-related decisions to enhance operator trust. A final methodological contribution is a Contingency-Aware Training Methodology, a comprehensive approach for training machine learning models across diverse contingency scenarios to enable robust performance under conditions not encountered during training.

1.3 Background and Related Literature

The increasing penetration of inverter-based resources (IBRs) has changed the dynamic behaviour of modern transmission systems in a fundamental way. For much of the twentieth century, bulk power-system stability was shaped primarily by the electromechanical dynamics of synchronous machines, whose rotating masses provided an inherent inertial buffer against sudden active-power imbalance. Classical stability theory, including the swing-equation representation of rotor-angle dynamics, the equal-area criterion, modal analysis, and small-signal linearization, was developed under this synchronous-machine paradigm and remains the conceptual basis for much of power-system stability analysis [24–26]. However, as synchronous generation is displaced by converter-interfaced renewable generation, battery energy storage systems, and power-electronic grid-support devices, the physical assumptions underlying these classical methods no longer hold without modification. Inertia, damping, and frequency response are no longer automatic consequences of machine physics; they increasingly depend on converter control laws, control bandwidths, measurement delays, device ratings, energy availability, and supervisory operating modes [27].

This shift has made low-inertia stability one of the central operational challenges in modern power systems. When supply and demand fall out of balance, the system frequency response unfolds across several time scales. The first stage is the inertial response, occurring within the first few hundred milliseconds to seconds, during which synchronous machines release or absorb kinetic energy and thereby slow the rate of change of frequency. Primary control follows within seconds, usually through governor droop or fast frequency response, arresting the frequency deviation but typically leaving a steady-state offset. Secondary control then acts on the minute scale through automatic generation control or equivalent coordination mechanisms to restore frequency to its nominal value [28–30]. In converter-dominated systems, the first stage is the most severely affected because fewer synchronous machines remain online to provide immediate inertial support. The resulting increase in rate of change of frequency (RoCoF), deeper frequency nadirs, reduced damping, and narrower stability margins make the system more sensitive to faults, line trips, generation loss, load changes, and protection interactions [6, 31]. Consequently, stability assessment, virtual inertia allocation, and corrective control have become tightly coupled problems rather than separate planning or operating tasks.

The stability problem itself is multidimensional. Small-signal stability concerns the response of the system to small perturbations around an operating point, while transient stability concerns the ability of the system to remain synchronized and within acceptable operating bounds following large disturbances such as short circuits, generator outages, line trips, or sudden load changes. Frequency stability, voltage stability, and converter-control stability further complicate the analysis in low-inertia grids because fast converter dynamics interact with slower electromechanical modes and network constraints [32]. These interactions are especially important in weak grids, where low short-circuit strength and high converter penetration can amplify control interactions, reduce damping, and create oscillatory behaviour not captured by traditional machine-dominated models [33]. As a result, modern stability assessment must account not only for the operating point, but also for topology, converter control modes, measurement availability, protection behaviour, and the time scale over which corrective action must be taken.

A particularly important distinction in converter-dominated systems is the difference between grid-following and grid-forming control. Grid-following inverters synchronize to an externally established voltage waveform, typically through a phase-locked loop, and inject current according to active- and reactive-power references. This operating mode has been widely deployed because it fits naturally into grids where synchronous machines provide the main voltage and frequency reference. However, at high IBR penetration or under weak-grid conditions, grid-following behaviour can become problematic because the inverter depends on a strong and stable voltage reference that may no longer exist. Grid-forming inverters, by contrast, actively establish or regulate voltage magnitude and frequency, and can therefore contribute more directly to system strength, synchronization, and frequency support [34]. This distinction matters for virtual inertia and corrective control because the same physical resource may provide very different stability services depending on whether it is operated as a grid-following unit, a grid-forming unit, a droop-controlled source, a virtual synchronous machine, or a fast-frequency-response resource [35].

The converter-control layer also introduces stability mechanisms that have no direct analogue in classical synchronous-machine systems. Phase-locked loops, inner current controllers, outer voltage or power controllers, current limiters, protection functions, and supervisory mode switching can all affect the dynamic response of an IBR during disturbances. In weak grids, PLL dynamics may interact with network impedance and other converter controls, causing poor damping or loss of synchronization [36]. During faults, converters are constrained by semiconductor current limits and cannot sustain the large short-circuit currents typical of synchronous machines. Current-limiting strategies, fault ride-through requirements, positive- and negative-sequence control, and post-fault recovery behaviour therefore become central to stability analysis [37]. These issues are directly relevant to learning-based control because a corrective action that is feasible in an abstract power-flow model may still be infeasible at the converter level if it violates current limits, energy limits, control bandwidths, or protection constraints. For this reason, any stability-support framework for low-inertia grids must treat virtual inertia and fast corrective action as constrained control services rather than unlimited ideal actuators.

Small-signal stability analysis remains one of the most established tools for studying local dynamic behaviour around an operating point. The standard procedure is to linearize the nonlinear differential-algebraic equations of the power system and examine the eigenvalues of the resulting state matrix. If all relevant eigenvalues lie in the stable region, small perturbations decay; if poorly damped or unstable modes are present, the system becomes vulnerable to oscillations or loss of stability. Eigenvalue analysis provides physically interpretable information because each mode can be associated with oscillation frequency, damping ratio, and participation factors that identify the states, generators, buses, or controllers most involved in that mode [38]. However, full eigenvalue analysis can become computationally demanding for large systems, especially when repeated across many contingencies, operating points, topology changes, and controller settings.

Sensitivity-based methods address part of this computational burden by estimating how critical eigenvalues move when system parameters are perturbed. Using left and right eigenvectors of the linearized system, first-order eigenvalue sensitivity provides a direct way to quantify how inertia constants, damping coefficients, controller gains, line parameters, or load levels affect modal damping and stability margins [39, 40]. This approach is particularly relevant for converter-dominated systems because many stability-support properties are controllable rather than fixed. For example, Altawallbeh *et al.* use sensitivity analysis to tune the gains of a virtual synchronous generator in a low-inertia grid, identifying which control parameters most improve damping and stability [41]. Sensitivity analysis therefore provides a useful bridge between classical modal theory and modern control design: it preserves physical interpretability while reducing the need for repeated exhaustive simulation. Its limitation, however, is that first-order sensitivities are usually local around a given operating point and topology. When the system experiences large disturbances, discrete switching events, or substantial changes in loading and renewable output, the validity of a locally linear sensitivity model may degrade unless the model is updated or embedded within a more adaptive framework [42].

The need for topology awareness is especially important because real power grids do not remain in a single operating mode. Transmission lines are switched, protection systems isolate faults, generation commitment changes, and power-electronic resources may transition between control modes. These discrete events interrupt the continuous dynamics of voltages, currents, angles, and frequencies, producing a hybrid dynamical system in which stability must be preserved both within each mode and across mode transitions [43]. Hybrid-system theory, switched-system analysis, and common Lyapunov function methods provide formal tools for studying such behaviour [44]. However, applying these tools directly to large-scale grids with learning-augmented controllers remains difficult. The controller may adapt its actions based on measured or predicted states, the network topology may change after faults or switching actions, and the feasible set of corrective actions may depend on operational constraints. This creates a central challenge for modern grid control: stability certificates and safety constraints must remain meaningful even when the underlying network graph and operating point change.

In addition to these physics-based methods, data-driven stability assessment has become an active area of research. Early approaches relied on statistical models, autoregressive methods, Kalman filtering, and feature-engineered classifiers to estimate system state, forecast load, or detect abnormal operating conditions [45, 46]. Later methods introduced support vector machines, decision trees, random forests, and other supervised learning techniques for transient stability classification, contingency screening, and event detection [47, 48]. More recent deep-learning approaches use convolutional neural networks, recurrent neural networks, temporal convolutional networks, autoencoders, transformers, and hybrid architectures to learn nonlinear mappings from measurements, simulation data, or phasor measurement unit streams to stability labels, voltage profiles, frequency metrics, or corrective recommendations [14, 49, 50]. These methods can be much faster at inference time than repeated time-domain simulation or nonlinear optimization, making them attractive for online security assessment and decision support.

The rise of data-driven stability assessment is also connected to the increasing availability of high-resolution measurements and large-scale simulation data. Phasor measurement units, digital relays, inverter telemetry, supervisory control and data acquisition systems, and advanced metering infrastructure provide richer visibility into grid behaviour than was historically available [51]. At the same time, offline contingency studies, dynamic simulations, and synthetic operating scenarios can generate labelled datasets for supervised learning. However, the quality of a learned stability model depends strongly on how these data are produced. A model trained only on routine operating conditions may not generalize to rare but severe disturbances, while a model trained on simplified simulations may inherit the assumptions and omissions of the simulator. This creates a close connection between data-driven stability assessment, simulation fidelity, and validation methodology.

Despite their promise, conventional data-driven methods face important limitations when used in stability-critical applications. The first limitation is physical consistency. A neural network may predict voltages, flows, angles, or control actions that fit training data statistically but violate Kirchhoff's laws, power balance, voltage limits, ramp limits, converter current limits, or device capability constraints. Such violations may be tolerable for offline analytics but are unacceptable inside a corrective-control loop [19, 49]. The second limitation is generalization under distribution shift. Historical data and simulation data may not adequately cover rare contingencies, unusual renewable profiles, changing load patterns, topology reconfigurations, cyber-physical disturbances, extreme weather events, or changes in protection and control settings [52]. The third limitation is interpretability. Operators are unlikely to trust black-box recommendations in emergency control settings unless the model can explain which buses, lines, modes, or constraints are driving a prediction or action [17]. A fourth limitation, particularly relevant for corrective control, is that prediction accuracy alone does not guarantee safe closed-loop behaviour. A model may be accurate on average but still produce unacceptable errors near stability boundaries, precisely where reliable decisions matter most [53].

Robustness and uncertainty therefore occupy a central place in learning-based power-system operation. Renewable generation uncertainty, load uncertainty, topology uncertainty, measurement noise, communication delays, forecasting errors, and imperfect knowledge of inverter control parameters can all affect the reliability of a learned predictor or controller [54]. Distribution shift is especially important in low-inertia systems because the most critical events are often rare, fast, and poorly represented in historical data. A controller trained on one set of contingencies may face a different line outage, operating condition, renewable dispatch pattern, or protection response in practice. This motivates the use of uncertainty-aware training, robust optimization, domain randomization, ensemble models, out-of-distribution detection, conservative safety filters, and stress-test validation [55]. For the present thesis, this issue is important because the proposed framework is intended not merely to predict typical operating trajectories, but to support corrective action under topology changes and contingency conditions.

These limitations have motivated the use of graph-based learning methods, because power systems are naturally structured as networks. A power grid can be represented as a graph whose nodes correspond to buses, generators, loads, or inverter resources, and whose edges correspond to transmission lines, transformers, or electrical couplings. Node features may include voltage magnitude, voltage angle, active and reactive demand, generation, inertia, damping, frequency, device status, or converter mode, while edge features may include impedance, admittance, thermal rating, power flow, switching status, or protection state [56]. Graph Neural Networks (GNNs) exploit this structure by defining learnable message-passing operations over the network, allowing each node to update its representation by aggregating information from neighbouring nodes and connected edges. This architecture naturally reflects the locality of many power-system interactions while preserving permutation invariance with respect to bus ordering [57, 58].

The development of GNNs for power systems has followed the broader evolution of graph deep learning. Spectral graph methods initially defined convolutions using graph Fourier transforms and Laplacian eigenvectors, linking graph learning to modal and spectral properties of the network. Spatial message-passing methods later became more widely used because they scale more naturally to large graphs, handle sparse network structure efficiently, and provide a more intuitive interpretation in terms of information exchange between neighbouring nodes [58]. In power-system applications, this message-passing view is especially attractive because physical interactions are mediated by network connectivity. Disturbances propagate through electrical couplings, voltage and angle changes are constrained by the admittance matrix, and topology changes directly alter the pattern of information flow. GNNs can therefore encode structural priors that are absent from conventional fully connected, convolutional, or recurrent architectures.

GNNs have been applied to power flow approximation, state estimation, topology identification, fault detection, contingency screening, transient stability prediction, voltage control, and dynamic security assessment. In power flow and state estimation, they can learn mappings from injections and topology to voltage solutions with lower inference cost than iterative solvers in repeated-query settings [59]. In fault detection and localization, they combine measurement patterns with network structure to identify the location and type of disturbance [60]. In stability assessment, they allow rapid screening of contingency scenarios by predicting stability margins, damping measures, or post-disturbance system states directly from graph-structured inputs [61]. Related methods also infer network topology from limited measurements, which is valuable when observability is incomplete or switching events are not fully reported [62]. Another related thread uses graph-based learning to predict basin stability or structural robustness directly from network topology, showing that non-trivial stability information can be extracted from the graph itself even before detailed dynamic simulation is performed [63].

Spatio-temporal GNNs extend this idea by modelling both network structure and time evolution. This is important because stability is not a static property of a single operating snapshot; it depends on how voltages, angles, frequencies, flows, and controls evolve after disturbances. Temporal dependencies may be captured using recurrent units, temporal convolutional networks, attention mechanisms, neural ordinary differential equations, or transformer-based message passing [64–66]. Recurrent models are intuitive for sequential data but may be difficult to train over long horizons. Temporal convolutional methods offer parallelizable training and stable gradients, while attention-based models can represent long-range dependencies in both space and time. These architectures are promising for dynamic grid applications, but they also increase model complexity and can make physical interpretation more difficult. Earlier work by the present authors applied a physics-informed spatio-temporal GNN to local contingency planning under N–k outages on the IEEE 118-bus system, providing a conference precursor to the framework developed in this thesis [67].

The usefulness of GNNs in power systems does not remove the need for physical constraints. On the contrary, as graph-based predictors move closer to operational decision-making, physical consistency becomes more important. Physics-informed machine learning addresses this issue by incorporating domain knowledge into learning through architectural inductive biases, physics-based loss terms, constrained layers, differentiable solvers, projection operators, or hybrid model-based/data-driven formulations [68, 69]. Physics-Informed Neural Networks (PINNs) are one important example, using residuals of governing equations as part of the training objective so that the learned model balances data fit with physical consistency [70]. In power systems, related ideas have been used to enforce power balance, Kirchhoff's laws, voltage constraints, dynamic equations, and energy-function conditions in learned models for power flow, state estimation, and stability assessment [71].

Physics-informed methods can be broadly distinguished by whether they impose constraints softly or hard. Soft-constraint methods add penalty terms to the loss function, encouraging the model to satisfy physical relationships but not guaranteeing that every output will be feasible. These methods are flexible and easy to combine with standard deep-learning architectures, but the resulting constraint satisfaction depends on penalty weights, data coverage, optimization quality, and the relative scale of loss terms. Hard-constraint methods, in contrast, enforce feasibility by construction or by projection. Examples include differentiable optimization layers, power-flow-consistent output parameterizations, Lagrangian and augmented-Lagrangian methods, constrained decoding, and projection onto feasible operating sets [72,73]. For safety-critical control, hard or post-processed feasibility enforcement is often preferable because even a small number of constraint violations can be operationally unacceptable. However, hard constraints can increase computational complexity, require differentiable formulations of physical constraints, or introduce approximation errors when the true feasible set is nonconvex [74].

Safety constraints in learning-based power-system control extend beyond algebraic feasibility. A corrective action must respect voltage limits, thermal limits, generator ramp rates, inverter current limits, state-of-charge constraints, reserve constraints, frequency limits, protection settings, and sometimes market or contractual constraints. It must also avoid destabilizing the dynamic system after it is applied. Lyapunov-based methods provide one route to formal stability reasoning by showing that a candidate energy or Lyapunov function decreases under the proposed control law. Control barrier functions, invariant-set methods, robust model predictive control, and action-shielding methods provide related tools for maintaining trajectories within safe regions [75]. In the context of learning-augmented controllers, these methods are often used not to replace learning, but to supervise it: the learning model proposes an action, and a safety layer modifies or rejects the action if it violates feasibility or stability conditions. This distinction is important because it allows a data-driven controller to retain adaptability while ensuring that the final implemented action remains physically admissible.

Reinforcement learning (RL) provides a natural framework for corrective control because it directly addresses sequential decision-making under uncertainty. In the Markov Decision Process formulation, the controller observes a state, applies an action, receives a reward, and learns a policy that maximizes expected cumulative reward over time [76]. For power systems, the state may include voltages, angles, frequencies, flows, generation levels, load levels, inertia values, topology indicators, storage states, and converter operating modes. The action may include generation redispatch, load shedding, demand response, reactive power support, tap-changing, storage dispatch, inverter mode selection, or virtual inertia adjustment. The reward can be designed to penalize voltage violations, frequency deviations, line overloads, excessive control effort, unserved energy, instability, or economic cost [77]. This flexibility makes RL attractive for emergency control, voltage regulation, frequency support, energy management, and coordinated operation of distributed energy resources.

Continuous-control algorithms such as Deep Deterministic Policy Gradient (DDPG) [78], Proximal Policy Optimization (PPO) [79], and Soft Actor-Critic (SAC) [80] have been used in power-system applications because many control variables are continuous rather than discrete. RL has shown promise in voltage control through reactive power coordination, frequency regulation through generation and demand-side action, emergency load shedding, microgrid energy management, and corrective control under contingencies [78,81,82]. However, several obstacles limit direct deployment. RL can be sample inefficient, requiring many interactions with a simulator before learning a reliable policy. The learned policy may exploit simulator artefacts and fail under model mismatch. Reward design can be difficult because a poorly weighted objective may encourage actions that improve one metric while degrading another. Exploration is also problematic in real power systems because unsafe trial actions cannot be allowed. For these reasons, RL for power systems is usually studied in simulation, digital twins, or offline training environments before any operational use is considered [83].

Safe reinforcement learning has emerged to address these concerns. Constrained policy optimization incorporates operational limits directly into the policy-learning problem. Lyapunov-based safe RL uses stability functions to restrict policy updates or certify safe regions. Shielding methods place a supervisory filter between the RL agent and the physical system, modifying actions that would violate safety constraints. Projection methods map the raw policy output to the closest feasible action under physical and operational constraints [84, 85]. These methods are essential for power systems because the controller must be evaluated not only by average reward, but also by worst-case constraint satisfaction, stability preservation, and behaviour under rare contingencies. In this thesis, RL is therefore not treated as a standalone black-box controller. Instead, it is combined with physics-informed prediction, spectral-sensitivity guidance, and feasibility projection so that the learned policy remains anchored to the physical structure and stability requirements of the grid.

The development and validation of learning-based controllers depend heavily on simulation environments. Because unsafe exploration cannot be performed on real grids, data-driven predictors and RL policies are usually trained and tested using offline dynamic simulations, benchmark systems, contingency libraries, and digital-twin-style environments [86]. Standard test systems such as the IEEE 39-bus, IEEE 68-bus, and IEEE 118-bus networks provide useful benchmarks because they allow reproducible comparison across methods, but they also simplify many features of real transmission systems, including detailed protection logic, communication delays, renewable forecasting errors, inverter-level controls, and operator intervention [87]. Digital twins and high-fidelity simulators can reduce this gap by combining network models, dynamic component models, measurement streams, and scenario generation [88]. However, even digital-twin validation must be interpreted carefully because model fidelity, parameter uncertainty, and scenario coverage strongly influence the apparent performance of a learned controller. This reinforces the need for stress testing, out-of-distribution evaluation, and physics-based safety supervision before learning-based methods can be considered operationally credible.

Among the available control variables in low-inertia systems, virtual inertia is especially important. Virtual inertia emulation allows converter-interfaced resources to mimic part of the inertial behaviour traditionally provided by synchronous machines. Implementation methods include derivative-based frequency control, virtual synchronous machine control, synchronverter control, droop-based fast frequency response, matching control, and energy-storage-supported inertial response [5,89]. These approaches differ in response speed, noise sensitivity, energy requirement, current-limiting behaviour, damping characteristics, and interaction with inner converter loops. Derivative-based methods can respond rapidly to RoCoF but may amplify measurement noise. Virtual synchronous machines emulate swing-equation dynamics more explicitly but require careful tuning of virtual inertia and damping. Storage-supported fast frequency response can provide strong short-term support but is limited by state of charge, converter rating, and energy availability [90].

The class of IBRs that can provide grid-support behaviour includes photovoltaic inverters, wind power converters, battery energy storage systems, STATCOMs, and other power-electronic devices. Unlike synchronous machines, these resources do not provide inertial response as a byproduct of physical rotating mass. Their contribution depends on the control architecture and on whether the device is operated in grid-following, grid-forming, droop-support, fast-frequency-response, or virtual-synchronous-machine mode [91]. Grid-following inverters generally synchronize to an existing voltage waveform through a phase-locked loop and inject commanded current, while grid-forming inverters actively regulate voltage and frequency and can contribute more directly to system strength and frequency support [92]. This distinction matters because virtual inertia is not merely a scalar parameter to be added to a bus. It is an engineered behaviour constrained by converter rating, current limits, measurement delays, energy availability, and the interaction between inner current loops, outer voltage or power loops, and supervisory control.

Once the device-level mechanism for virtual inertia is defined, the system-level problem becomes one of allocation. Not every bus benefits equally from additional inertia. The effect of inertia depends on modal participation, electrical distance, local generation and load patterns, network topology, damping, converter control mode, and the location of disturbances. Early allocation strategies often used heuristic rules, such as placing support near large disturbances, weak areas, renewable-rich regions, or buses with large frequency deviations. Optimization-based methods later formulated inertia placement as an objective-driven problem, minimizing RoCoF, maximizing damping ratios, improving frequency nadir, reducing control effort, or satisfying security constraints subject to device and budget limits [6]. Sensitivity-based methods provide a more interpretable alternative by estimating how critical stability modes respond to incremental changes in inertia or damping at candidate buses. Sarwar *et al.*, for example, use sensitivity scores to allocate grid-forming resources for voltage support under contingency conditions [93].

Spectral sensitivity analysis provides the mathematical foundation for this allocation problem. If a critical eigenvalue determines the dominant oscillatory mode or stability margin, then the derivative of that eigenvalue with respect to a virtual inertia parameter indicates how beneficial an incremental inertia change at a given location may be. By combining this sensitivity with constraints on total inertia budget, per-bus limits, and device availability, the allocation problem can be transformed into a structured decision problem that is physically interpretable and computationally efficient [9, 40]. This is attractive because it links local controllable parameters to global stability behaviour. However, conventional spectral allocation methods are often static: they assume a fixed topology, a fixed linearization point, or a limited set of precomputed scenarios. Modern grids require allocation strategies that remain effective when the topology changes, operating conditions vary, and corrective control must be updated online. Recent work on topology-aware sensitivities begins to address this limitation by accounting for changing network configurations in the sensitivity calculation [94]. The remaining challenge is to integrate such interpretable sensitivity information into a closed-loop learning and control framework rather than treating it as a separate offline planning tool.

Interpretability is therefore not an optional addition but a requirement for practical grid AI. Operators need to know why a model predicts instability, why an inertia allocation is recommended, why a control action is selected, and which parts of the grid are most responsible for the decision. Interpretability also supports debugging, model validation, regulatory review, operator training, and post-event analysis. For GNNs, interpretability methods may identify important nodes, edges, features, time steps, or subgraphs. Attention mechanisms can indicate which neighbours or measurements influence a node representation. Gradient-based saliency can estimate the sensitivity of outputs to input features. Perturbation-based explanations can test how predictions change when buses, lines, or features are removed. Prototype and surrogate methods can summarize learned behaviour using simpler representative cases or approximate models [95]. These methods are particularly relevant in power systems because the graph structure has physical meaning: an important edge may correspond to a stressed transmission corridor, an important node may correspond to a weak bus, and an important feature may correspond to inertia, voltage, demand, converter status, or line loading.

Nevertheless, interpretability methods must be used carefully. In fast control loops, however, the practical use of these explanations is constrained. Interpretability outputs are most valuable for offline audit, post-event analysis, operator training, and model validation rather than for online intervention, since the loop timescales do not allow a human operator to inspect and act on an explanation before the next control step is executed. Attention weights do not always provide faithful explanations of model behaviour. Gradient saliency can be noisy or sensitive to scaling. Perturbation explanations may create unrealistic operating points if they remove graph elements without preserving physical feasibility. Surrogate models may oversimplify nonlinear behaviour near stability boundaries [96]. For this reason, interpretability in grid AI should be connected to physical quantities wherever possible. Explanations based on modal participation, eigenvalue sensitivity, power-flow constraints, voltage violations, line overloads, current limits, or inertia scarcity are more useful to operators than purely abstract neural-network attributions. A central motivation of the present thesis is therefore to combine learned representations with physically meaningful diagnostics, so that the framework can provide not only predictions and actions, but also explanations that correspond to recognizable grid phenomena.

Taken together, the literature does not point to a lack of methods for stability assessment, learning-based prediction, virtual inertia placement, or corrective control. Rather, the main limitation is that these methods are usually developed around different assumptions, time scales, and operational objectives. Several specific gaps follow from this fragmentation: (i) GNN-based models are effective for topology-aware prediction, but they are often evaluated primarily as estimators of post-contingency states or stability labels rather than as components of a closed-loop control architecture; (ii) spectral and eigenvalue-sensitivity methods provide physically interpretable guidance for damping improvement and inertia placement, but they are commonly formulated as offline or locally linear allocation tools; (iii) RL-based controllers can learn coordinated corrective actions, yet their practical use is limited when feasibility, stability preservation, and behaviour under rare topology changes are not explicitly enforced; (iv) many learning-based studies simplify converter and weak-grid effects, including current limits, control-mode transitions, measurement delays, energy availability, and fault ride-through constraints, even though these factors directly affect the feasibility of virtual inertia and fast corrective control; and (v) interpretability is often treated as a post-processing step rather than being tied to physical quantities that operators can use, such as modal participation, eigenvalue sensitivity, inertia scarcity, voltage violations, current limits, and stressed transmission corridors [21–23, 97].

As a result, the open problem is not prediction, sensitivity analysis, virtual inertia allocation, or control in isolation, but their integration into a physically constrained, topology-aware, and interpretable decision-making framework for low-inertia systems. This limitation becomes more pronounced in converter-dominated grids, where feasible corrective action depends not only on network-level quantities such as voltage limits, line loading, and modal damping, but also on converter-level constraints such as current limits, control mode, measurement delay, energy availability, and fault ride-through behaviour. These considerations motivate a framework in which prediction, sensitivity-guided inertia allocation, corrective control, feasibility projection, and physically meaningful diagnostics are designed as coupled components rather than independent modules [98].

This thesis addresses this problem by developing a unified framework for topology-aware stability assessment, virtual inertia allocation, and corrective control in low-inertia power systems. The framework combines a physics-informed spatio-temporal GNN for fast prediction of system evolution, a spectral-sensitivity method for interpretable virtual inertia guidance, an RL policy for coordinated corrective action, and a supervisory projection mechanism that enforces feasibility under operational constraints. The contribution is therefore not the use of GNNs, RL, or virtual inertia in isolation, but their integration into a closed-loop architecture that explicitly recognizes the structure of the grid, the changing nature of topology, the constrained behaviour of converter-interfaced resources, the uncertainty of operating conditions, and the need for interpretable stability support. By combining graph-based prediction, spectral-sensitivity allocation, reinforcement-learning control, and Lyapunov/projection-based safety supervision, the thesis aims to provide a practical pathway toward trustworthy AI-assisted operation of converter-dominated transmission systems.

1.4 Thesis Outline

This thesis is organized into five chapters. Chapter 1 introduces the motivation, research gap, objectives, and scope of the thesis.

Chapter 2 lays the foundation for the power-system model and the proposed spectral-sensitivity virtual-inertia allocation method. It formalizes the swing-equation dynamics, derives the associated small-signal state-space representation, and extends the model to account for topology changes. The chapter then develops the eigenvalue sensitivity metric used to guide constrained virtual inertia placement under total budget, per-bus limits, and device-availability constraints.

Chapter 3 presents the integrated closed-loop framework. It details the physics-informed spatio-temporal GNN architecture and training objective, the reinforcement-learning formulation for coordinated corrective actions including generation adjustment, demand response, and virtual inertia updates, and the supervisory coordination mechanism that blends the policy output with the sensitivity-based reference. The chapter also develops the feasibility projection layer and provides Lyapunov-based analysis to establish stability preservation under bounded inertia updates and piecewise-constant connected topologies.

Chapter 4 describes the experimental setup and evaluates the proposed framework on the IEEE 118-bus system under diverse operating conditions and contingency scenarios. The results assess prediction accuracy, topology-aware virtual inertia allocation, closed-loop control behaviour, constraint satisfaction, interpretability through saliency-based and sensitivity-based diagnostics, and computational performance. Additional stress-test scenarios are included to evaluate robustness beyond the journal case studies.

Finally, Chapter 5 summarizes the main findings and contributions, discusses limitations, and outlines directions for future work, including extensions to larger systems, richer inverter control models, uncertainty-aware operation, hardware-in-the-loop validation, and practical deployment considerations for AI-assisted stability support in converter-dominated grids.

Chapter 2

Power System Model and Analysis

2.1 System Modelling Framework

The physical and analytical foundations for understanding small-signal stability and designing topology-aware virtual inertia allocation in transmission power systems are developed by focusing on electromechanical behaviour. The central objective is to describe how the generator rotor angles and frequencies evolve following small disturbances about a steady operating equilibrium. Rotor angles encode the relative positions of generator shafts, while frequencies measure how fast they rotate with respect to the nominal system frequency. Together, these variables capture how well the system remains synchronized under modest changes in generation and demand.

The corresponding low-frequency electromechanical dynamics evolve through interactions among generators, inverter-based resources, and the transmission network. Their behaviour is highly sensitive to three closely related aspects of the system. First, the distribution of inertia across the network determines how effectively the grid resists and absorbs disturbances. Inertia, whether provided by synchronous machines or emulated as virtual inertia by inverter-based resources, acts as a physical buffer against rapid frequency deviations, and its spatial allocation at different buses can substantially shape the overall dynamic response. Second, the network topology, understood as the pattern of connections between buses and transmission lines, governs how power flows and oscillatory signals propagate. Changes in line status, like outages or switching operations, alter these pathways and can shift the dominant electromechanical modes. Third, the operating conditions, represented by the underlying power-flow solution, specify the equilibrium point around which the system is linearized. Different loading levels and generation dispatches can shift this equilibrium, making the grid more or less sensitive to a given disturbance.

Because electromechanical dynamics depend on the distribution of inertia, network topology, and operating conditions, they provide a natural basis for assessing stability in modern low-inertia power systems. As conventional synchronous generation is increasingly replaced or supplemented by renewable sources and inverter-based resources, the amount and placement of physical inertia change, and virtual inertia must be allocated at specific buses chosen with the network layout in mind, rather than spread evenly across the grid, to keep small-signal behaviour acceptable. The analysis of these low-frequency dynamics therefore underpins the subsequent development of stability metrics, eigenvalue sensitivity measures, and constrained virtual-inertia placement strategies.

The physical and analytical development begins with the classical swing equation. It extends to a network-coupled, multi-machine dynamic model that captures the influence of inertia, damping, and transmission-level power exchanges. This swing-based description provides a physically interpretable model of how rotor angles and frequencies evolve in time in response to disturbances. These nonlinear swing dynamics are then linearized about a steady-state operating point, yielding a state-space representation that is suitable for small-signal stability analysis. Linearization approximates the full nonlinear model with a local linearization that accurately describes the system's response to sufficiently small perturbations, thereby making analytical tools from linear systems theory applicable.

In this linearized setting, system stability is assessed by the spectral properties of the state matrix; its eigenvalues and associated damping ratios identify the dominant and least-damped electromechanical modes, which determine how oscillations grow, decay, or persist following small disturbances. This eigenvalue-based (spectral) viewpoint, often referred to as small-signal stability analysis, is standard in power-system dynamics because it provides quantitative measures of how quickly oscillations die out and which parts of the system participate most strongly in each mode.

Building on this linearized model, first-order sensitivities of the electromechanical eigenvalues with respect to inertia and network-coupling parameters are derived. These sensitivities describe how small changes in physical parameters (such as a bus's inertia or the strength of a transmission corridor) shift the eigenvalues and, consequently, modal damping and frequency. They provide a physics-based measure of structural criticality by revealing which buses and transmission corridors most strongly influence the dominant oscillatory modes. The resulting sensitivity information is then used to pose a constrained virtual inertia allocation problem under limited resources, ensuring that additional inertia (virtual) is placed where it is most effective. In this way, spectral analysis not only diagnoses stability weaknesses but also guides the deployment of control resources. The optimized allocation serves as a reference signal for the learning-based closed-loop control frameworks developed in subsequent chapters, linking the underlying physics to the higher-level control and decision-making layers.

Within this small-signal representation, virtual inertia is modelled as a modification of the effective inertia matrix $\mathbf{M}(t)$. Its physical realization, however, is achieved through

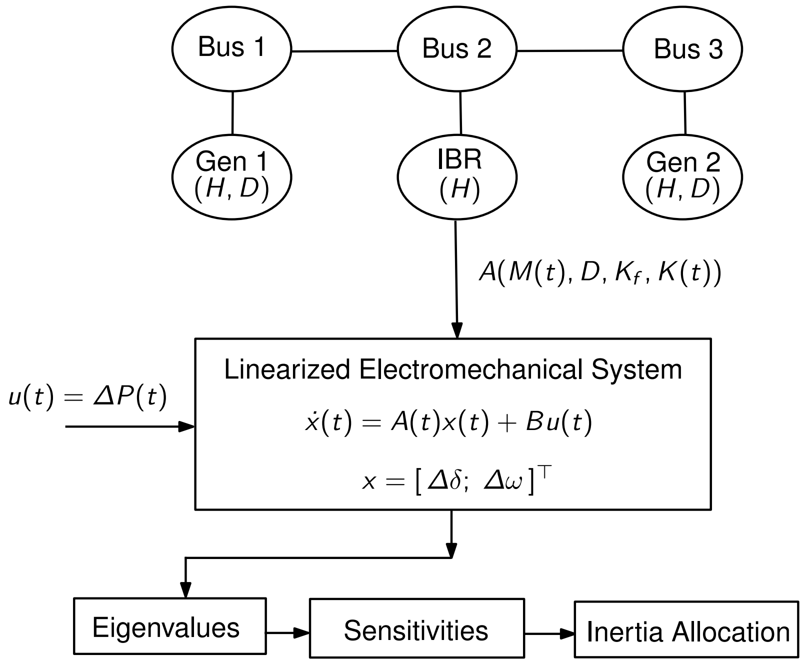

Figure 2.1: Electromechanical modelling framework for small-signal stability and virtual inertia allocation. The physical transmission network and generator parameters determine the time-varying linearized state matrix $\mathbf{A}(t) = \mathbf{A}(\mathbf{M}(t), \mathbf{D}, \mathbf{K}_f, \mathbf{K}(t))$, where $\mathbf{K}_f = \mathrm{diag}(K_{f_i})$ is the frequency-dependent load matrix defined in (2.14). The resulting electromechanical state-space model (2.22) forms the basis for eigenvalue analysis, modal sensitivity evaluation, and topology-aware inertia allocation.

converter-based power injection into the system via $\mathbf{u}(t)$, as explicitly developed in Chapter 3.

2.1.1 Implementation of Virtual Inertia in the Framework

Virtual inertia in this framework is not added to the synchronous generators, which already provide physical rotational inertia through their rotating masses. It is instead applied at inverter-enabled buses, where converter-interfaced resources emulate an inertial response that these buses would not otherwise provide. Generator buses therefore retain their fixed physical inertia throughout topology changes, while inverter buses receive an updated virtual inertia allocation whenever the network configuration changes, such as during line outages or switching operations.

The spectral-sensitivity allocator produces a per-bus virtual inertia constant H_i, which enters the swing dynamics through the inertia coefficient $M_i = 2H_i/\omega_s$ defined in (2.11). The allocated virtual inertia therefore appears as a modification of the effective inertia matrix $\mathbf{M}(t)$, subject to a total inertia budget and a per-bus cap $H_{\max}$, and is assigned only to inverter-enabled buses.

Physically, this allocated inertia is realized through the converter-based active-power injection represented by the control input $\mathbf{u}(t)$, applied only at inverter-enabled buses. The abstract entry M_i in the inertia matrix therefore corresponds to a concrete converter-level control action rather than to added physical mass.

Within the closed-loop framework, this mechanism serves two roles. First, the spectral-sensitivity allocator uses the eigenvalue sensitivities of the linearized state matrix to determine where virtual inertia most effectively damps the critical electromechanical modes, producing a topology-aware reference allocation. Second, the reinforcement-learning coordinator adjusts this reference over the operating horizon and combines it with the other control actions, subject to the same budget and per-bus limits, so that the realized inertia distribution tracks evolving topology and operating conditions while remaining physically feasible. The device-level realization of this inertia emulation is described in Section 2.5.2.

In summary, we consider a network-coupled, multi-machine transmission system operating near a steady-state equilibrium, as defined by a power-flow solution. Around this equilibrium, the electromechanical dynamics are described by rotor-angle and frequency deviations. They are approximated by a linear state-space model obtained by linearizing the nonlinear power-transfer coupling. External disturbances and controllable actions enter this model through an input vector $\mathbf{u}$, interpreted as an additive real-power injection or withdrawal that captures net active-power imbalances and fast corrective actions (e.g., generator redispatch, demand-response signals, or injections from aggregated inertia resources) applied at selected buses.

As illustrated in Fig. 2.1, variations in inertia distribution and network coupling modify the state matrix and therefore reshape the electromechanical spectrum. The compact state-space representation serves as the analytical bridge between physical system parameters and the spectral quantities used for stability assessment and inertia allocation.

To make the state-space representation explicit, the state is defined as

$$x \triangleq \begin{bmatrix} \Delta\delta^\top & \Delta\omega^\top \end{bmatrix}^\top, \tag{2.1}$$

where $\Delta\delta = \delta - \delta^*$ and $\Delta\omega = \omega - \omega^*$ are the deviations of rotor angle and frequency from a steady operating point (δ^*, ω^*) determined by the power-flow solution. Angles and frequencies are measured in a frame that rotates at the nominal system frequency (50/60 Hz), so at this operating point, the frequency deviation is zero ($\omega^* = 0$), and δ^* is the corresponding set of steady rotor angles in that frame.

The next section derives these electromechanical dynamics from the classical swing equation and their network-coupled form, providing the nonlinear basis for the linearized model $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$.

2.2 Electromechanical Dynamics and Network Coupling

The electromechanical model used throughout this chapter is derived from the basic rotational dynamics of a synchronous generator. Physically, the generator rotor behaves as a spinning mass that accelerates when the mechanical torque exceeds the electromagnetic torque and decelerates when the reverse holds. Starting from this rotational equation of motion, the classical swing equation is derived to describe how the rotor angle and speed respond to imbalances between mechanical input and electrical output power. This single-machine description is then generalized to account for network power coupling between buses, yielding a nonlinear, network-coupled multi-machine model in which inertia, damping, and transmission-level interactions jointly shape the evolution of frequency and rotor-angle dynamics.

The rotor dynamics of a synchronous generator follow from Newton's second law for rotational motion,

$$J \frac{d^2\theta_m(t)}{dt^2} = T_m(t) - T_e(t), \tag{2.2}$$

where J is the rotor moment of inertia, $\theta_m(t)$ is the mechanical rotor angle, and T_m and T_e are the applied mechanical torque and opposing electromagnetic torque, respectively. The kinetic energy stored in the rotating mass at mechanical angular speed $\omega_m = \frac{d\theta_m}{dt}$ is

$$E_k = \frac{1}{2} J \omega_m^2. \tag{2.3}$$

In power-system analysis, inertia is commonly expressed using the inertia constant H, defined as the stored kinetic energy at rated speed divided by the system MVA base.

$$H = \frac{E_{k,\text{rated}}}{S_{\text{base}}} = \frac{\frac{1}{2}J\omega_s^2}{S_{\text{base}}}, \tag{2.4}$$

where ω_s is the synchronous electrical angular speed and S_{base} is the system MVA base. Rearranging (2.4) yields

$$J = \frac{2HS_{\text{base}}}{\omega_s^2}. \tag{2.5}$$

With $\omega_m = \dot{\theta}_m$, (2.2) can be written in speed form as

$$J\frac{d\omega_m}{dt} = T_m - T_e. \tag{2.6}$$

Multiplying (2.6) by ω_m gives the power-balance form

$$J\omega_m\frac{d\omega_m}{dt} = P_m - P_e, \tag{2.7}$$

where $P_m = T_m\omega_m$ and $P_e = T_e\omega_m$ denote the mechanical input and electrical output power, respectively. Under the standard small-deviation assumption $\omega_m \approx \omega_s$ and using (2.5), (2.7) reduces to the classical per-unit swing equation in speed form,

$$\frac{2H}{\omega_s}\frac{d\omega}{dt} = P_m - P_e. \tag{2.8}$$

Let δ denote the electrical rotor angle corresponding to the mechanical rotor angle θ_m, related by $\delta = \frac{p}{2}\theta_m$, where p is the number of poles. In synchronous reference-frame coordinates, the kinematic relation $\dot{\delta} = \omega$ holds, where ω is the frequency deviation. Generalizing to a multi-bus system, δ_i and ω_i denote the rotor angle and frequency deviation at bus i, with the corresponding kinematic relation $\dot{\delta}_i = \omega_i$. In deviation form about an operating point (δ^*, ω^*), this becomes $\Delta\dot{\delta}_i = \Delta\omega_i$, which is used in the small-signal analysis of Section 2.3.

The swing equation can equivalently be expressed in angle form as

$$\frac{2H}{\omega_s}\ddot{\delta} = P_m - P_e. \tag{2.9}$$

To capture frequency-sensitive load effects and damper-winding contributions, the swing equation for generator (or aggregated inertia resource) i is written as

$$M_i\ddot{\delta}_i + D_i\dot{\delta}_i = P_{m_i} - P_{e_i}, \tag{2.10}$$

where D_i is an effective damping coefficient and

$$M_i = \frac{2H_i}{\omega_s}. \tag{2.11}$$

Throughout, the speed (frequency) deviation in synchronous reference coordinates is defined as $\omega = \dot{\delta}$, so that at an equilibrium operating point $\omega_i^* = 0$.

Figure 2.2 summarizes the nonlinear swing dynamics at a single bus in control-oriented form. This structure, with inertia, damping, and network power feedback appearing explicitly, is the starting point for the linearization and state-space representation developed in Section 2.3.

To represent inverter-based resources (IBRs) and demand behaviour, the electrical power term is decomposed into generation and demand components at each bus. The net electrical power injection is written as

$$P_{e_i} = P_{g_i}^s + P_{g_i}^I - P_{d_i} \tag{2.12}$$

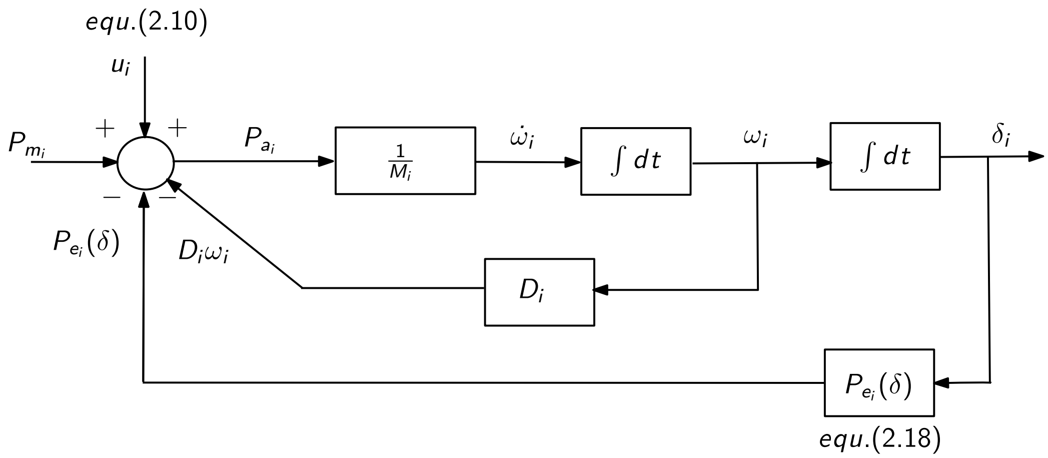

Figure 2.2: Control-oriented block diagram of the generator electromechanical dynamics. The summation node implements the nonlinear swing dynamics in (2.10), written in speed form, with inertia defined by (2.11) and nonlinear network electrical power given by (2.18). The network power feedback block references the lossless active-power coupling in (2.18), corresponding to equation (2.18) in the thesis. The inertia block and cascaded integrators generate (ω_i, δ_i) in synchronous coordinates, forming the nonlinear basis for the linearized state-space model (2.22).

where $P^s_{g_i}$ denotes synchronous-machine generation, $P^I_{g_i}$ represents power injections from inverter-based resources, and P_{d_i} is the real-power demand at bus i.

The demand P_{d_i} is modelled as a nominal component perturbed by endogenous deviations of frequency and voltage about the operating point, together with an exogenous component:

$$P_{d_i} = P^0_{d_i}\left[1 + k_{f_i}\,\frac{\Delta\omega_i}{\omega_s} + k_{v_i}\,\frac{V_i - V^*_i}{V^*_i}\right] + P^e_{d_i}(t), \tag{2.13}$$

where $P^0_{d_i}$ is the baseline active-power demand at the operating point and carries units of power, $\Delta\omega_i = \omega_i - \omega_s$ is the absolute frequency deviation, and ω_s and V^*_i are the nominal synchronous frequency and the operating-point voltage magnitude at bus i, respectively. The coefficients k_{f_i} and k_{v_i} are dimensionless frequency- and voltage-load sensitivities that quantify the percentage change in active-power demand per unit of normalized frequency or voltage deviation; this is the standard frequency- and voltage-dependent load representation in small-signal power-system analysis. Each term inside the bracket is therefore dimensionless, the bracketed factor is a unitless multiplier of the baseline power $P^0_{d_i}$, and $P^e_{d_i}(t)$ enters directly in power units, so (2.13) is dimensionally consistent.

For the small-signal electromechanical analysis adopted in this chapter, voltage magnitudes are treated as quasi-steady at their operating-point values, so $V_i \approx V_i^*$, and the voltage-dependent term in (2.13) drops out. This modelling choice is consistent with the focus on rotor-angle and frequency dynamics introduced in Section 2.1 and is revisited in Section 2.3, where voltage magnitudes enter the linearized model implicitly through the synchronizing coefficients rather than as dynamic states. Defining the absolute frequency-load coefficient

$$K_{f_i} \triangleq \frac{P^0_{d_i}\, k_{f_i}}{\omega_s}, \tag{2.14}$$

which carries units of power per unit of frequency; the demand deviation about the operating point reduces to

$$\Delta P_{d_i} = K_{f_i}\, \Delta\omega_i + P^e_{d_i}(t), \tag{2.15}$$

where each term now has consistent units of power, since K_{f_i} explicitly carries the dimensional factor required to convert frequency deviations into equivalent power deviations.

Substituting (2.12) and (2.15) into (2.10) and recalling that $\dot{\delta}_i = \Delta\omega_i$ in synchronous coordinates, the swing dynamics become

$$M_i\ddot{\delta}_i + (D_i + K_{f_i})\dot{\delta}_i = P_{m_i} - P^I_{g_i} - P^n_{e_i} - P^e_{d_i}(t), \tag{2.16}$$

where $P^n_{e_i}$ denotes the network-coupled electrical power exchange. The frequency-load coefficient K_{f_i} therefore augments the mechanical damping D_i, while exogenous demand variations enter as a forcing term on the right-hand side.

The network-coupled component $P^n_{e_i}$ introduced in (2.16) is determined by power transfer over the transmission lines connected to bus i. It can be expressed as

$$P^n_{e_i} = \sum_{j=1}^{N} V_i V_j |Y_{ij}| \cos(\delta_i - \delta_j - \gamma_{ij}), \tag{2.17}$$

where V_i and V_j are bus voltage magnitudes, and $Y_{ij} = |Y_{ij}|\angle\gamma_{ij}$ denotes the complex line admittance between buses i and j.

For transmission networks, line reactance typically dominates resistance ($R_{ij} \ll X_{ij}$), so the series admittance is approximately imaginary. Writing $Y_{ij} = G_{ij} + jB_{ij}$, this implies $G_{ij} \approx 0$ and $\gamma_{ij} \approx 90°$. Substituting $\gamma_{ij} \approx 90°$ into (2.17) gives

$$\cos(\delta_i - \delta_j - \gamma_{ij}) \approx \sin(\delta_i - \delta_j).$$

Under this lossless approximation, active-power coupling is therefore governed primarily by angle differences across lines, and the susceptance matrix B_{ij} captures the strength of each synchronizing line.

Accordingly, (2.17) reduces to

$$P^n_{e_i} = \sum_{j=1}^{N} V_i V_j B_{ij} \sin(\delta_i - \delta_j), \tag{2.18}$$

and substituting (2.18) into the swing dynamics yields a nonlinear, network-coupled multi-machine model in which angle differences across transmission corridors provide the synchronizing interactions that govern collective electromechanical behaviour.

This enables us to distinguish endogenous electromechanical dynamics, shaped by inertia, damping, IBR control actions, and network coupling, from exogenous demand variations that enter the system as external forcing terms. Such a separation enables subsequent analysis to isolate the impact of inverter-based resources and demand uncertainty on frequency stability and oscillatory behaviour.

The resulting equations describe a nonlinear, network-coupled multi-machine swing model in which synchronizing interactions arise from angle differences across transmission corridors and are modulated by inertia and damping at each bus. While this formulation captures the essential electromechanical physics governing the evolution of frequency and angle deviations, direct stability analysis of the nonlinear model is not tractable for large-scale systems. To enable systematic assessment of oscillatory behaviour and its dependence on inertia distribution and network structure, the dynamics are linearized about an operating point in the following section, yielding a state-space representation suitable for modal and sensitivity analysis.

2.3 State-Space Representation for Spectral Analysis

The nonlinear network-coupled swing dynamics derived in Section 2.2 are accurate but inconvenient for stability analysis. The eigenvalue and sensitivity tools used throughout this thesis to rank buses, design allocations, and certify stability are defined for linear systems and apply only locally to the nonlinear model around a chosen operating point. For this reason, the swing equations are linearized about an equilibrium (δ^*, ω^*), with perturbations defined as $\Delta\delta = \delta - \delta^*$ and $\Delta\omega = \omega - \omega^*$. Defining the state vector $x = [\Delta\delta^\top \ \Delta\omega^\top]^\top$, the linearized dynamics can be written in a first-order state-space form, whose explicit structure is derived below and used for the spectral analysis of the next section.

Studying the nonlinear network-coupled swing dynamics directly is impractical for large transmission systems, for three related reasons. First, every bus is coupled to every other bus through nonlinear sine terms of the form $V_i V_j B_{ij} \sin(\delta_i - \delta_j)$, and these coupled equations have no closed-form solution except in a few special cases. Second, the system is large: a grid with hundreds of buses has hundreds of rotor angles and frequencies evolving together, which is too much to reason about by hand or to characterize analytically. Third, the dynamics depend on the operating point, so any analysis is valid only for one specific dispatch and topology and would have to be redone whenever conditions change. To handle these difficulties, oscillatory behaviour and its dependence on inertia and topology are studied using a linear approximation, which exposes the system's modal properties and can be updated as the operating point evolves.

Linearization about an equilibrium operating point yields a state-space model whose system matrix admits spectral decomposition. The eigenvalues of this matrix characterize the dominant electromechanical modes, including their oscillation frequencies and (for oscillatory modes) damping characteristics. At the same time, the associated eigenvectors describe the spatial structure of these modes across buses and generators.

In power system dynamics, the stability of small perturbations is determined by the eigenstructure of the linearized state matrix. Eigenvalues with negative real parts correspond to decaying perturbations, while those with small magnitude real parts indicate weakly damped behaviour. For complex-conjugate pairs, the imaginary component determines the oscillation frequency, and the ratio of the real to the imaginary parts characterizes modal damping.

Consequently, modes located near the imaginary axis impose the most restrictive stability margins and are of primary interest in inertia-sensitive electromechanical analysis. Such spectral information provides a quantitative basis for evaluating how variations in inertia and topology influence the behaviour of small perturbations. Consequently, eigenvalue sensitivity analysis provides the analytical bridge between physical parameter variations and topology-aware virtual inertia allocation strategies developed in subsequent sections.

Here, an operating point denotes a steady-state equilibrium of the network-coupled dynamics for a given topology and dispatch. In synchronous reference coordinates, it corresponds to $\omega^* = 0$ and a set of steady-state angle differences δ^* that satisfy the real-power balance implied by the underlying power flow solution. Linearization about (δ^*, ω^*) provides a locally valid approximation of the nonlinear dynamics for small perturbations, enabling stability assessment via the eigenstructure of the resulting state matrix.

The state vector is restricted to rotor-angle and frequency deviations, (2.1), reflecting a deliberate focus on electromechanical timescale dynamics. At this timescale, the dominant stability-limiting phenomena following small disturbances are governed by the interaction among inertia, damping, and real-power transfer across the transmission network.

Voltage magnitudes and reactive-power dynamics primarily evolve on faster electromagnetic or slower secondary-control time scales. They are therefore treated as quasi-steady when forming the small-signal electromechanical model. In particular, steady-state voltage values V^* enter the analysis implicitly through the synchronizing coefficients $K_{ij} = \partial P_{e_i}^n / \partial \delta_j \big|_{\delta^*, V^*}$, but voltage deviations are not retained as dynamic states.

This modelling choice does not imply that voltage and reactive power are unimportant; rather, it reflects the objective of isolating the electromechanical modes that dominate low-frequency oscillatory behaviour and are most sensitive to the distribution of inertia and the network structure. Extensions that explicitly incorporate voltage dynamics are possible but fall outside the scope of the present small-signal stability framework.

Applying a first-order Taylor expansion to the network-coupled swing equations yields a set of local linear balance relations. Each relation corresponds to a bus i that possesses electromechanical inertia (synchronous machines) or control-defined inertia (inverter-based resources),

$$M_i \Delta\dot{\omega}_i + D_i \Delta\omega_i + \sum_{j=1}^{N} K_{ij} \Delta\delta_j = u_i, \tag{2.19}$$

where u_i represents the incremental active-power disturbance and/or fast control action at bus i. The synchronizing coupling coefficients are defined by

$$K_{ij} = \left. \frac{\partial P_{e_i}^{n}}{\partial \delta_j} \right|_{\delta=\delta^*}. \tag{2.20}$$

Under the lossless coupling approximation in (2.18), the off-diagonal terms satisfy $K_{ij} = -V_i^* V_j^* B_{ij} \cos\bigl(\delta_i^* - \delta_j^*\bigr)$ for $i \neq j$, while $K_{ii} = -\sum_{j \neq i} K_{ij}$. Consequently, $\mathbf{K}$ has a weighted Laplacian structure evaluated at the operating point. Intuitively, $\mathbf{K}$ quantifies the extent to which buses synchronize across the transmission network. Large entries correspond to strongly coupled transmission lines, while weak or zero entries indicate weak or absent synchronizing paths. Topology changes modify the set of active edges and therefore update $\mathbf{K}$ by removing or altering the corresponding coupling terms; in this thesis, such changes are modelled as piecewise-constant updates of $\mathbf{K}(t)$ across decision intervals.

Stacking all buses gives the vector form

$$\mathbf{M}\Delta\dot{\boldsymbol{\omega}} + (\mathbf{D} + \mathbf{K}_f)\Delta\boldsymbol{\omega} + \mathbf{K}\Delta\boldsymbol{\delta} = \mathbf{u}, \tag{2.21}$$

where $\mathbf{M} = \mathrm{diag}(M_i)$ is the diagonal inertia matrix, $\mathbf{D} = \mathrm{diag}(D_i)$ is the damping matrix, $\mathbf{K}_f = \mathrm{diag}(K_{f_i})$ represents the frequency-dependent load contribution, $\mathbf{K} = [K_{ij}]$ is the coupling matrix, and $\mathbf{u}$ aggregates incremental disturbances and actuation across buses.

Using the kinematic relation $\Delta\dot{\delta} = \Delta\omega$ and defining the state vector as $\mathbf{x} = [\Delta\boldsymbol{\delta}^\top \; \Delta\boldsymbol{\omega}^\top]^\top$, the dynamics can be written in first-order state-space form as

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}(\mathbf{D}+\mathbf{K}_f) \end{bmatrix}}_{\mathbf{A}(\mathbf{M},\mathbf{D},\mathbf{K}_f,\mathbf{K})} \mathbf{x} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1} \end{bmatrix}}_{\mathbf{B}} \mathbf{u}. \tag{2.22}$$

This representation makes explicit how inertia ($\mathbf{M}$), damping ($\mathbf{D}$), frequency-dependent load effects ($\mathbf{K}_f$), and topology/operating-point coupling ($\mathbf{K}$) jointly determine the system matrix $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ used for stability assessment and sensitivity computations in subsequent sections.

Equivalently, (2.22) may be written compactly as

$$\dot{\mathbf{x}} = \mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})\mathbf{x} + \mathbf{B}\mathbf{u}.$$

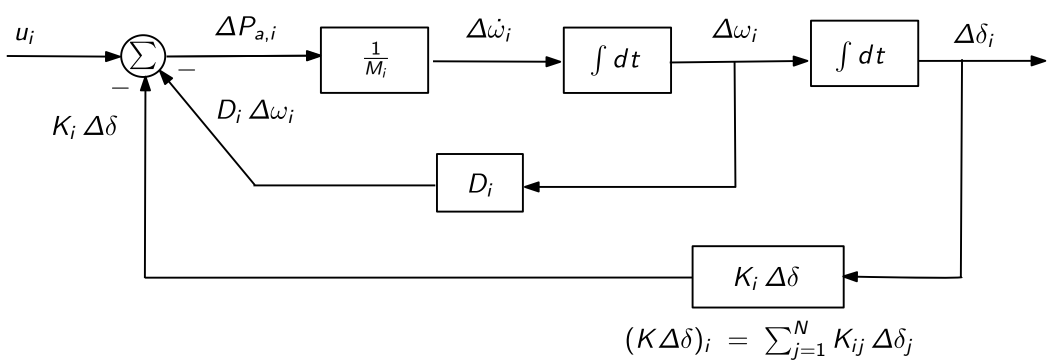

Figure 2.3: Block-diagram realization of the linearized swing dynamics at bus i. The damping block represents the combined effective damping $D_i + K_{f_i}$, where K_{f_i} is the frequency-dependent load coefficient defined in (2.14). For visual clarity, the block is labelled D_i; the full expression including K_{f_i} is used in the state-space model (2.22).

Figure 2.3 summarizes the linearized swing dynamics at bus i in control-oriented form. The summation node implements the deviation-form scalar balance in (2.19), with inertia M_i, damping, and the linearized network-coupling term $\sum_j K_{ij}\Delta\delta_j$ derived in (2.20) appearing as distinct functional blocks. The cascaded integrators generate $(\Delta\omega_i, \Delta\delta_i)$, and stacking the per-bus relations yields the first-order state-space model in (2.22). The damping block is drawn as D_i for visual simplicity; in the full small-signal model, the frequency-load contribution K_{f_i} defined in (2.14) adds to D_i, so the effective damping seen by the electromechanical mode at bus i is $D_i + K_{f_i}$. This control-oriented view complements the algebraic state-space form and motivates the spectral analysis developed in the next section.

The Laplacian structure of $\mathbf{K}$ reflects the fact that real-power transfer is driven primarily by angle differences across transmission corridors. As a result, a uniform shift of all angles does not change power transfers, implying invariance to the absolute reference frame and producing a rigid reference mode. Since this mode does not affect relative-angle dynamics, stability assessment focuses on the non-rigid electromechanical modes, using grounded/reference-bus or center-of-inertia coordinates.

With the linearized plant model established, the next section introduces the modal stability quantities used throughout the thesis. In particular, the eigenvalues of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ determine oscillation frequencies and damping, and the least-damped modes form the critical set used to evaluate how inertia and coupling perturbations affect stability margins.

2.4 Stability Metrics and Sensitivity-Based Criticality

With the linearized state-space representation (2.22) established in Section 2.3, small-signal behaviour is characterized through the eigenstructure of the system matrix $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$.

For the homogeneous electromechanical dynamics obtained by setting $\mathbf{u} = \mathbf{0}$,

$$\dot{\mathbf{x}} = \mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})\,\mathbf{x}, \tag{2.23}$$

The eigenvalues λ_k of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ determine whether perturbations around the operating equilibrium decay or grow. Local asymptotic stability of the equilibrium requires $\mathrm{Re}(\lambda_k) < 0$ for all stability-relevant electromechanical modes.

Eigenvalues commonly appear as complex-conjugate pairs $\lambda_k = \sigma_k \pm j\omega_k$, where $\sigma_k = \mathrm{Re}(\lambda_k)$ governs exponential decay or growth and $\omega_k = \mathrm{Im}(\lambda_k)$ sets oscillation frequency. To compare damping across oscillatory modes, the modal damping ratio is defined as

$$\zeta_k = -\frac{\mathrm{Re}(\lambda_k)}{|\lambda_k|}. \tag{2.24}$$

Larger ζ_k corresponds to stronger damping, whereas smaller ζ_k indicates weakly damped oscillations that can persist and therefore constrain stability margins.

Under the lossless network approximation, the coupling matrix $\mathbf{K}$ exhibits a Laplacian-like structure, implying invariance to uniform shifts of all rotor angles. This invariance induces a rigid reference mode that does not affect the dynamics of relative angles. Stability assessment is therefore conducted on the non-rigid electromechanical modes obtained after accounting for this reference subspace (e.g., through grounded/reference-bus coordinates or center-of-inertia coordinates). Let $\mathcal{K}$ denote the set of critical non-rigid electromechanical modes. In practice, $\mathcal{K}$ is selected as the least-damped subset, since these modes dominate post-disturbance behaviour and impose the most restrictive stability margins. Equivalently, $\mathcal{K}$ may be formed by selecting modes with the largest real parts $\mathrm{Re}(\lambda_k)$ (closest to the imaginary axis) or the smallest damping ratios ζ_k defined in (2.24).

In quantifying how modal stability varies with changes in operating point and localized structural perturbations, we next use first-order eigenvalue sensitivity theory. This is a standard tool that describes how the eigenvalues of a matrix shift as its entries change slightly. The main result is that the shift of a simple eigenvalue can be expressed in terms of the right and left eigenvectors of the original matrix, so modal changes can be predicted without recomputing the full Eigen decomposition whenever a parameter changes. A comprehensive treatment of power-system small-signal stability can be found in [99]. Let $(\lambda_k, \mathbf{v}_k)$ denote a simple eigenpair of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ with associated left eigenvector $\mathbf{w}_k$, satisfying

$$\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})\,\mathbf{v}_k = \lambda_k \mathbf{v}_k, \tag{2.25}$$

$$\mathbf{w}_k^{\mathrm{H}} \mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K}) = \lambda_k \mathbf{w}_k^{\mathrm{H}}, \tag{2.26}$$

together with the normalization

$$\mathbf{w}_k^{\mathrm{H}} \mathbf{v}_k = 1, \tag{2.27}$$

where $(\cdot)^{\mathrm{H}}$ denotes the conjugate transpose.

For a small perturbation of the system matrix, $\mathbf{A} \mapsto \mathbf{A} + \Delta\mathbf{A}$, the corresponding first-order eigenvalue variation is

$$\Delta\lambda_k \approx \mathbf{w}_k^{\mathrm{H}} (\Delta\mathbf{A})\,\mathbf{v}_k. \tag{2.28}$$

Equivalently, if a scalar parameter θ affects the dynamics through $\mathbf{A}(\theta)$, the first-order eigenvalue sensitivity is

$$\frac{\partial \lambda_k}{\partial \theta} = \mathbf{w}_k^{\mathrm{H}} \left(\frac{\partial \mathbf{A}}{\partial \theta} \right) \mathbf{v}_k. \tag{2.29}$$

The bilinear form $\mathbf{w}_k^{\mathrm{H}}(\cdot)\,\mathbf{v}_k$ makes explicit how localized changes in system parameters project onto mode k through the associated left and right eigenvectors.

Two classes of variations are emphasized. First, inertia variations modify $\mathbf{M}$ and hence the inertia-weighted blocks of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$, enabling sensitivities of the form $\partial\lambda_k/\partial M_i$ that quantify how inertia adjustments at bus i shift the damping and frequency of mode k. Second, structural and topological variations modify the coupling matrix $\mathbf{K}$ (and hence $\mathbf{A}$), thereby altering the synchronizing strength between buses. This motivates sensitivity analysis with respect to coupling parameters associated with individual transmission lines, or directly with respect to entries that contribute to $\mathbf{K}$, thereby identifying transmission lines whose weakening or removal most strongly affects the least-damped modes. Because the least-damped modes determine the stability margins, sensitivities are evaluated primarily on the critical set $\mathcal{K}$. For an oscillatory mode $\lambda_k = \sigma_k \pm j\omega_k$, changes in $\mathrm{Re}(\lambda_k)$ indicate stabilizing or destabilizing trends. The geometric picture is illustrated in Figure 2.4.

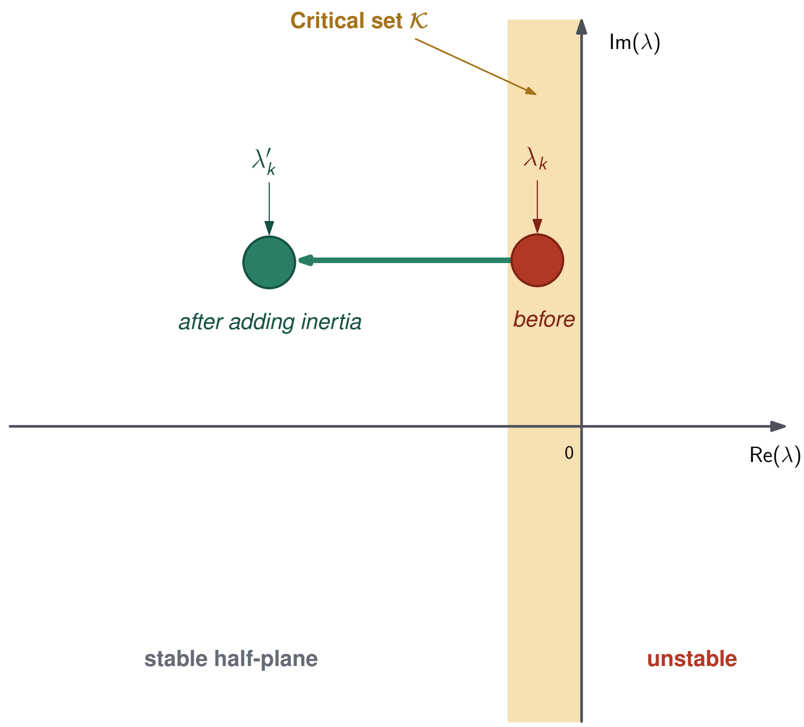

Figure 2.4: Geometric meaning of eigenvalue sensitivity: adding inertia at a strongly participating bus shifts a critical eigenvalue λ_k leftward to λ'_k, improving the damping of the corresponding mode.

Figure 2.4 shows what eigenvalue sensitivity means geometrically. An eigenvalue λ_k of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ is plotted as a point in the complex plane: its real part controls how fast a post-disturbance oscillation dies out, and its imaginary part sets the oscillation frequency. Eigenvalues near the imaginary axis (the critical set $\mathcal{K}$, shown shaded) are dangerous because their oscillations decay slowly and can leave the system ringing for many seconds, whereas eigenvalues deeper in the left half-plane settle almost immediately. Adding inertia at a strongly participating bus shifts the critical eigenvalue λ_k leftward to λ_k', which is exactly the improvement in damping the rest of this section quantifies through the sensitivity $\partial\lambda_k/\partial M_i$. The same picture applies to coupling perturbations $\partial\lambda_k/\partial K_{ij}$: corridors that produce the largest leftward motion are the ones whose strengthening (or whose loss, in the opposite direction) most strongly affects stability margins. The overall pathway from modal sensitivities to bus-level criticality scores and to the resulting reference inertia allocation is summarized visually in Figure 2.7 (Section 2.5.1).

Eigenvalue sensitivities, therefore, enable bus-level and transmission line-level rankings based on their influence on the critical electromechanical modes. A bus-level criticality score for inertia-based actuation is obtained by aggregating real-part sensitivities over the critical set $\mathcal{K}$:

$$g_i = \sum_{k\in\mathcal{K}} \mathrm{Re}\left\{\frac{\partial\lambda_k}{\partial M_i}\right\}. \tag{2.30}$$

Negative values of g_i indicate that increasing inertia at bus i tends to shift the critical eigenvalues leftward (improving damping). In contrast, positive values indicate a tendency toward reduced damping under increased inertia at that location.

The same eigenvalue sensitivity machinery applies to the transmission network itself. Each line (i, j) contributes a coupling weight to $\mathbf{K}$ physically, the synchronizing strength of that line, equal under the lossless approximation to $V_i^* V_j^* B_{ij} \cos\big(\delta_i^* - \delta_j^*\big)$. Differentiating λ_k

with respect to this weight yields the sensitivity $\partial\lambda_k/\partial K_{ij}$, which measures how the critical mode responds when the line's contribution to the network changes.

The connection to contingency analysis is direct. A line outage is the extreme version of weakening a coupling weight: the weight is set to zero, and the corresponding entries of $\mathbf{K}$ are removed. To first order, the induced eigenvalue shift is $\Delta\lambda_k \approx \mathbf{w}_k^{\mathrm{H}}(\Delta\mathbf{A})\mathbf{v}_k$, where $\Delta\mathbf{A}$ encodes the change to the coupling block of (2.22). Lines with the largest $|\partial\lambda_k/\partial K_{ij}|$ are the ones whose outages would produce the largest rightward shift of the critical mode, as illustrated in Figure 2.5. These are, therefore, the contingencies an operator-level screening procedure should examine first.

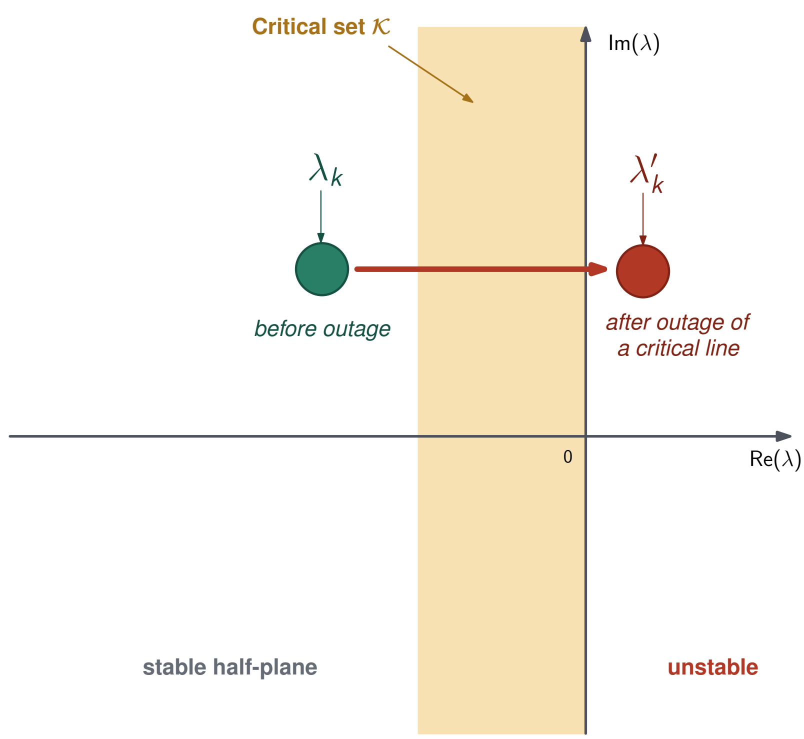

Figure 2.5: Effect of losing a high-criticality line: the critical eigenvalue shifts rightward, reducing damping and, in severe cases, crossing into the unstable half-plane.

The bus-level and line-level rankings together provide a complete first-order map of stability-critical structure under a given operating point: where intervention is most effective, and where the system is most vulnerable. Because both rankings come from the same eigendecomposition, evaluating one is essentially free once the other has been computed, which is the practical reason both can be maintained continuously as the operating point and topology evolve.

These criticality measures, therefore, play two roles in the framework. First, they provide an interpretable diagnostic that ranks the most stability-critical buses and lines for a given operating point and topology, supplying the physics-based ordering on which any allocation or contingency-screening procedure can be built. Second, the bus-level score g_i acts as the analytical bridge to control synthesis: it is the weighting that turns the abstract sensitivity formula (2.29) into the concrete allocation problem developed in Section 2.5.1. The full pathway from sensitivities to bus-level scores to the resulting reference inertia allocation is summarized in Figure 2.7.

2.5 Virtual Inertia Allocation and Actuation

Virtual inertia is the central controllable quantity in the proposed framework. Rather than being a static system parameter, it is treated as a decision variable that can be adjusted at each control interval in response to changing topology and operating conditions. At the device level, inverter-based resources at eligible buses emulate inertial behaviour by injecting active power proportional to the measured rate of change of frequency, as described in Section 2.5.2. At the system level, the allocation problem determines how the available virtual inertia budget should be distributed across eligible buses to maximally improve damping of the critical electromechanical modes. The reference allocation H_{ref} produced here is then tracked by the closed-loop controller developed in Chapter 3, which issues inertia update commands through the input channel $\mathbf{u}(t)$ as the system evolves.

The sensitivity-based criticality measures developed in Section 2.4 provide a physics-grounded indication of where localized inertia changes have the greatest impact on the least-damped electromechanical modes. Building on this insight, limited virtual inertia resources must be allocated judiciously to improve damping of the critical modes while respecting practical feasibility constraints. To this end, a sensitivity-guided virtual inertia allocation strategy is developed based on the linearized plant matrix $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ and the critical mode set $\mathcal{K}$. The objective is to distribute controllable, converter-provided inertia across eligible buses within a total virtual-inertia budget and per-bus bounds, yielding a physics-based reference allocation that serves as both a stability benchmark and a guiding signal for the learning-based and closed-loop control frameworks developed in subsequent chapters. The section then connects these abstract allocation variables to the actuation capabilities available in the linearized model's input channel.

2.5.1 Sensitivity-Guided Virtual Inertia Allocation

The modal sensitivities developed in Section 2.4 quantify how eigenvalues move in response to localized parameter changes. In this subsection, we specialize these sensitivities to variations in bus-level inertia and use them to construct a physics-based virtual inertia allocation.

Let λ_k be a (simple) eigenvalue of $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K})$ with right and left eigenvectors $\mathbf{v}_k$ and $\mathbf{w}_k$, normalized such that $\mathbf{w}_k^\dagger \mathbf{v}_k = 1$. From (2.29), the first-order sensitivity of λ_k with respect to a scalar parameter θ is

$$\frac{\partial \lambda_k}{\partial \theta} = \mathbf{w}_k^\dagger \left(\frac{\partial \mathbf{A}}{\partial \theta} \right) \mathbf{v}_k. \tag{2.31}$$

For the sensitivity derivation that follows, the small frequency-load contribution $\mathbf{K}_f$ is absorbed into the effective damping and suppressed in the notation; the resulting expression captures the dominant dependence of the eigenvalues on inertia. The simplified swing state matrix is then

$$\mathbf{A}(\mathbf{M},\mathbf{D},\mathbf{K}) = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}\mathbf{D} \end{bmatrix}, \tag{2.32}$$

the dependence on inertia appears only through $\mathbf{M}^{-1}$ in the lower blocks. Since $\mathbf{M} = \text{diag}(M_i)$ is diagonal, a change in M_i affects only the ith diagonal entry. To isolate this effect, define the coordinate projector $\mathbf{E}_i = \mathbf{e}_i\mathbf{e}_i^\top$, which has a 1 in position (i,i) and zeros elsewhere.

Using $\mathbf{M}^{-1} = \text{diag}(1/M_i)$, the derivative of $\mathbf{M}^{-1}$ with respect to M_i is

$$\frac{\partial \mathbf{M}^{-1}}{\partial M_i} = -\frac{1}{M_i^2}\mathbf{E}_i,$$

so that the corresponding derivative of the state matrix is

$$\frac{\partial \mathbf{A}}{\partial M_i} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \frac{1}{M_i^2}\mathbf{E}_i\mathbf{K} & \frac{1}{M_i^2}\mathbf{E}_i\mathbf{D} \end{bmatrix}. \tag{2.33}$$

To make the resulting expression more interpretable, partition the right eigenvector as $\mathbf{v}_k = [\mathbf{v}_{k,\delta}^\top \;\; \mathbf{v}_{k,\omega}^\top]^\top$, consistent with the state ordering $\mathbf{x} = [\Delta\boldsymbol{\delta}^\top \;\; \Delta\boldsymbol{\omega}^\top]^\top$. The left eigenvector is partitioned analogously as $\mathbf{w}_k = [\mathbf{w}_{k,\delta}^\top \;\; \mathbf{w}_{k,\omega}^\top]^\top$. Substituting (2.33) into (2.31) and using the block structure of $\partial\mathbf{A}/\partial M_i$ yields the closed-form expression

$$\frac{\partial \lambda_k}{\partial M_i} = \frac{1}{M_i^2}\,(\mathbf{w}_{k,\omega})_i^*\Big((\mathbf{K}\mathbf{v}_{k,\delta})_i + D_i(\mathbf{v}_{k,\omega})_i\Big), \tag{2.34}$$

where $(\cdot)_i$ denotes the ith entry and $(\cdot)^*$ denotes complex conjugation. The real part $\mathrm{Re}\{\frac{\partial \lambda_k}{\partial M_i}\}$ indicates whether increasing inertia at bus i shifts mode k leftward (improved damping) or rightward (reduced damping) in the complex plane.

To obtain a bus-level score that targets the least-damped modes, sensitivities are aggregated over the critical set $\mathcal{K}$:

$$g_i = \sum_{k \in \mathcal{K}} \mathrm{Re}\left\{ \frac{\partial \lambda_k}{\partial M_i} \right\}. \tag{2.35}$$

Negative values of g_i indicate that increasing inertia at bus i tends to improve the damping margin of the critical modes. In contrast, positive values indicate a tendency toward reduced damping under increased inertia at that location.

Let H_i denote the effective inertia constant at bus i, decomposed as

$$H_i = H_i^{\mathrm{real}} + H_i^{\mathrm{vir}}, \tag{2.36}$$

where H_i^{real} is the physical inertia contribution (fixed for synchronous machines) and H_i^{vir} is the controllable virtual inertia provided by inverter-enabled resources. In the per-unit swing model, inertia enters through

$$M_i = \frac{2H_i}{\omega_s}, \tag{2.37}$$

so allocating virtual inertia corresponds to modifying M_i through H_i^{vir}.

Virtual inertia allocation is performed over a set of inverter-eligible buses $\mathcal{G}_{\mathrm{inv}}$, while synchronous generator buses $\mathcal{G}_{\mathrm{gen}}$ retain their physical inertia values. The allocation is computed under a total virtual-inertia budget and per-bus bounds:

$$\sum_{i \in \mathcal{G}_{\mathrm{inv}}} H_i^{\mathrm{vir}} \leq H_{\mathrm{total}}, \tag{2.38}$$

$$0 \leq H_i^{\mathrm{vir}} \leq H_{\max}, \qquad \forall i \in \mathcal{G}_{\mathrm{inv}}, \tag{2.39}$$

$$H_i^{\mathrm{vir}} = 0, \qquad \forall i \in \mathcal{G}_{\mathrm{gen}}. \tag{2.40}$$

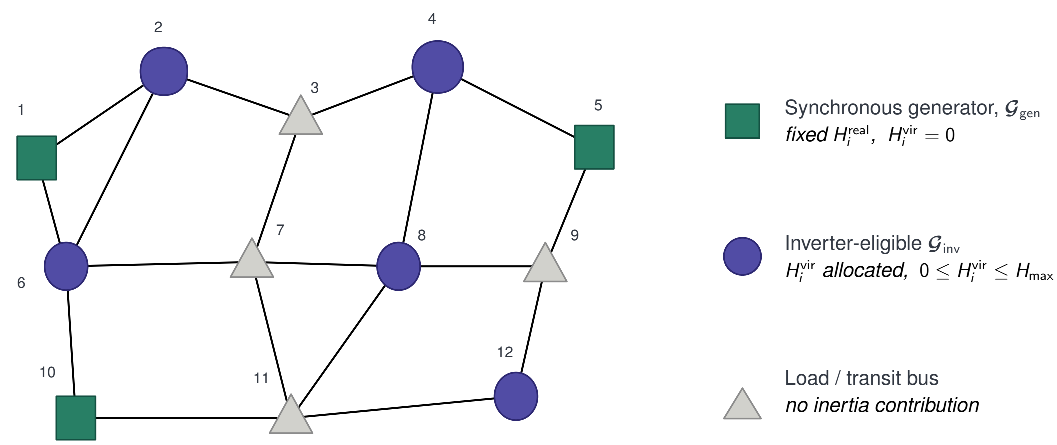

Figure 2.6: Bus classification used by the allocation problem. Synchronous-generator buses ($\mathcal{G}_{\mathrm{gen}}$, teal squares) carry fixed physical inertia and receive no virtual inertia; inverter-eligible buses ($\mathcal{G}_{\mathrm{inv}}$, purple circles) receive an allocated H_i^{vir} subject to per-bus and budget bounds; load and transit buses (gray triangles) contribute no inertia.

Figure 2.6 illustrates the partition. The three constraints in (2.38)–(2.40) act on different bus classes: the budget constraint bounds the total virtual inertia summed over $\mathcal{G}_{\mathrm{inv}}$, the box constraint enforces device-level limits at each inverter-eligible bus, and the equality constraint fixes the virtual inertia at synchronous-generator buses to zero so that physical and virtual contributions are not double-counted.

The criticality scores g_i already tell us, to first order, how each unit of inertia at bus i shifts the critical eigenvalues. Recall from (2.35) that a more negative g_i corresponds to a stronger stabilizing effect, so a natural choice is to make the allocation problem itself first-order in H_i^{vir}, weighting each bus by its score:

$$\min_{H^{\mathrm{vir}}} \sum_{i \in \mathcal{G}_{\mathrm{inv}}} g_i \, H_i^{\mathrm{vir}} \tag{2.41}$$

subject to the feasibility constraints (2.38)–(2.40). Buses with the most negative g_i reduce the objective the most per unit of inertia, so (2.41) sends the virtual inertia where it is expected to produce the largest leftward shift of the critical modes.

The linearized state-space model (2.22) is itself a first-order Taylor expansion of the swing dynamics, so a higher-order objective would impose a precision that the underlying plant cannot deliver. The allocation determines where inertia is placed rather than predicting the exact magnitude of each eigenvalue shift, and first-order sensitivities already provide the direction and ordering that the LP relies on. The per-bus cap $H_{\max}$ and the budget H_{total} confine each H_i^{vir} to a small neighbourhood of the operating point, the regime in which the linear approximation is valid. And H_{ref} is a reference rather than a final command; residual second-order error in the predicted shift is absorbed by the closed-loop controller of Chapter 3 as the system evolves. Higher-order terms would refine the magnitude of the predicted shift but not its direction or ordering, at a substantially greater computational cost on a grid whose operating point and topology change frequently.

In practice, (2.41) may be solved directly as a linear program or implemented via a greedy/projection procedure that allocates budget to the most negative g_i buses while enforcing (2.38) and (2.39). The resulting virtual inertia map H^{vir} defines a physics-based reference allocation H_{ref} through (2.36) for the current operating point and topology.

In the greedy step (line 6 of Algorithm 2.1), the buses in $\mathcal{G}_{\mathrm{inv}}$ are sorted by g_i from most negative to least negative, since more negative scores indicate a stronger stabilizing effect. The most stabilizing bus is filled first up to the per-bus limit $H_{\max}$; the remainder of the

Algorithm 2.1 Sensitivity-guided virtual inertia reference allocation

1: **Input:** operating point and topology (hence $\mathbf{K}$), parameters $(\mathbf{M}, \mathbf{D})$, bounds $(H_{\max})$, budget H_{total}
2: Form $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K})$ using the simplified form in (2.32)
3: Compute eigenpairs $(\lambda_k, \mathbf{v}_k, \mathbf{w}_k)$ of $\mathbf{A}$ and select critical set $\mathcal{K}$
4: Compute $g_i = \sum_{k\in\mathcal{K}} \operatorname{Re}\{\partial\lambda_k/\partial M_i\}$ using (2.34) and (2.35)
5: Initialize $H_i^{\text{vir}} \leftarrow 0$ for all $i \in \mathcal{G}_{\text{inv}}$; set $H_i^{\text{vir}} \leftarrow 0$ for $i \in \mathcal{G}_{\text{gen}}$
6: Allocate budget to buses with most negative g_i (greedy), saturating at $H_{\max}$, until $\sum_{i\in\mathcal{G}_{\text{inv}}} H_i^{\text{vir}} = H_{\text{total}}$
7: Form $H_{\text{ref},i} \leftarrow H_i^{\text{real}} + H_i^{\text{vir}}$ for all buses
8: **Output:** H_{ref} satisfying the virtual-inertia constraints

budget then flows to the next bus, and so on, until the total budget H_{total} is exhausted. This produces a feasible allocation in a single pass and is an exact solution of the linear surrogate (2.41) when the feasibility set is the budget–box polytope used in this thesis.

Figure 2.7 summarizes the sensitivity-guided virtual inertia allocation developed in this subsection. The eigenanalysis of the linearized state matrix $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ produces modal sensitivities with respect to bus-level inertia, which are aggregated into the scores g_i. These scores then drive the linear allocation problem (2.41), whose solution H_i^{vir} satisfies the budget constraint (2.38), the per-bus cap $H_{\max}$ in (2.39), and the requirement that synchronous-generator buses receive no virtual inertia (2.40). The total budget constraint limits the overall amount of virtual inertia that can be deployed across the network, while the per-bus bounds ensure that each inverter-enabled bus operates within its device capability limits.

Combining the allocated virtual inertia with the existing physical inertia yields the reference inertia profile H_{ref} defined in (2.36). This reference allocation provides a physics-based target inertia distribution for the current operating point and network topology.

The next subsection describes how the allocated H_i^{vir} is physically delivered by an inverter

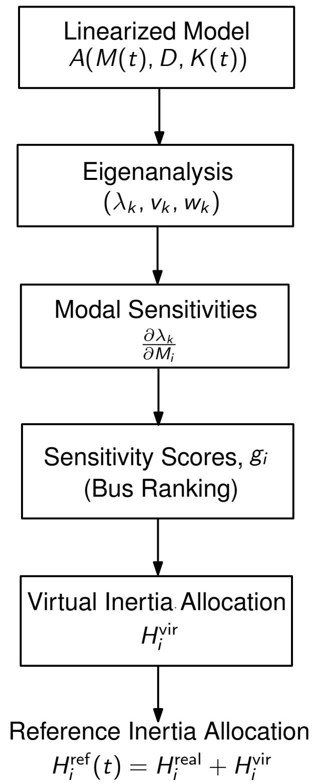

Figure 2.7: Sensitivity-guided virtual inertia allocation and update framework.

at the device level, while the closed-loop framework that dynamically tracks H_{ref} over time is developed in Chapter 3.

2.5.2 IBR Realization with Controllable Virtual Inertia

The allocation (2.41) produces a target virtual inertia H_i^{vir} for each inverter-eligible bus. This subsection describes how the inverter at bus i actually delivers that target, so that the abstract decision variable H_i^{vir} corresponds to a concrete behaviour of a physical device.

The inverter starts by measuring the local terminal frequency ω_i through its phase-locked loop (PLL). From this measurement, two signals are formed: the frequency deviation

$\Delta\omega_i = \omega_i - \omega_s$ and its time derivative $\Delta\dot{\omega}_i$, the rate of change of frequency (RoCoF). These are the two inputs that drive virtual inertia emulation. How quickly the inverter can respond to a disturbance is determined by the PLL bandwidth.

The inverter then injects an extra active power $P_i^{\text{vir}}(t)$ designed to mimic what a synchronous machine would do. The simplest and most common form, used in virtual synchronous machine (VSM) control and inertia-emulating grid-forming droop, takes the form

$$P_i^{\text{vir}}(t) = -\,M_i^{\text{vir}}\,\Delta\dot{\omega}_i(t) -\, D_i^{\text{vir}}\,\Delta\omega_i(t), \tag{2.42}$$

where M_i^{vir} sets the virtual inertia and D_i^{vir} sets the virtual damping. The first term opposes the rate of change of frequency, as a spinning rotor would; the second term opposes the frequency deviation itself, providing damping.

The link between the allocated H_i^{vir} from (2.36) and the controller gain M_i^{vir} in (2.42) comes from the per-unit conversion in (2.37). Setting

$$M_i^{\text{vir}} = \frac{2\,H_i^{\text{vir}}}{\omega_s} \tag{2.43}$$

makes the inverter's power injection act on the swing dynamics like an extra rotating mass with inertia constant H_i^{vir}. So the allocation problem (2.41) translates directly into inverter gain commands: H_i^{vir} at the system level, M_i^{vir} at the device level.

To see why this is the same as adding physical inertia, add $P_i^{\text{vir}}(t)$ to the active-power balance at bus i in (2.16) and substitute (2.42). Using $\Delta\dot{\delta}_i = \Delta\omega_i$, the result in deviation form is

$$\big(M_i + M_i^{\text{vir}}\big)\Delta\ddot{\delta}_i + \big(D_i + K_{f_i} + D_i^{\text{vir}}\big)\Delta\dot{\delta}_i + \cdots = u_i. \tag{2.44}$$

The inverter's contribution appears as an additional term in the inertia and damping at bus i. From the small-signal point of view used in this chapter, this looks the same as an actual increase in the physical M_i and D_i. That is the reason H_i^{vir} can be treated as part of the inertia matrix in the allocation problem, even though no physical mass is being added.

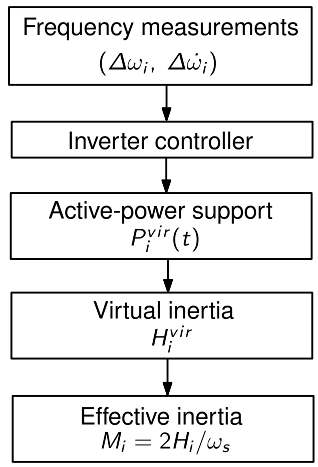

Figure 2.8: Inverter-based realization of controllable virtual inertia. Local frequency and rate-of-change-of-frequency measurements feed into the inverter controller, which generates the synthetic power injection (2.42) using the gain M_i^{vir} from (2.43).

Two practical limits determine how well the control law (2.42) can be realized. First, the inverter's current rating caps how large $P_i^{\mathrm{vir}}(t)$ can be during big disturbances; once the converter hits its current limit, the actual injection falls short of the commanded value, and the effective virtual inertia drops below H_i^{vir}. Second, the PLL and measurement chain introduce a small delay between a change in frequency and the inverter's response, which limits how quickly the inverter can usefully react to a disturbance. Both limits are captured by the per-bus bound $H_i^{\mathrm{vir}} \leq H_{\max}$ in (2.39), which represents the largest sustained virtual inertia each device can deliver within its operating range.

2.5.3 Actuation Capabilities and Operating Context

The previous subsection produced a physics-based reference inertia allocation H_{ref} for a given operating point and topology. Real power systems, however, are continually reshaped by control and operational actions that the linearized model must also account for. These include generator re-dispatch, inverter-based resource (IBR) modulation, demand response, and corrective measures triggered by contingencies. This subsection situates these actions within the linearized framework (2.22) so that the allocation, actuation, and operational drivers can be discussed in the same language.

The actions enter the linearized dynamics through one of two channels. They either modify the plant parameters $(\mathbf{M}, \mathbf{D}, \mathbf{K})$ inside $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$, which reshapes the modal structure, or they enter through the input channel $\mathbf{u}(t)$ as a corrective power injection. Inverter-based resources are the only class capable of doing both. A frequency-following injection acts through $\mathbf{u}(t)$, while a virtual-synchronous-machine controller modifies the effective $(\mathbf{M}, \mathbf{D})$ as derived in Section 2.5.2. The remainder of the subsection follows this structure. The operational drivers (renewable variability, dispatch, topology) are discussed first, followed by the actuation channels (governors, IBR control stack, demand response, fast frequency response), and finally the uncertainty under which all of this operates, which motivates the closed-loop framework developed in Chapter 3.

Figure 2.9 maps each operational mechanism to the channel through which it enters (2.22). The left column collects mechanisms that act on the plant matrix $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$ by changing one of its constituent matrices. The right column collects mechanisms that act on the input channel $\mathbf{B}\mathbf{u}(t)$ as a corrective power injection. The remainder of this subsection walks through the entries column by column, then closes with how uncertainty in either column propagates over time.

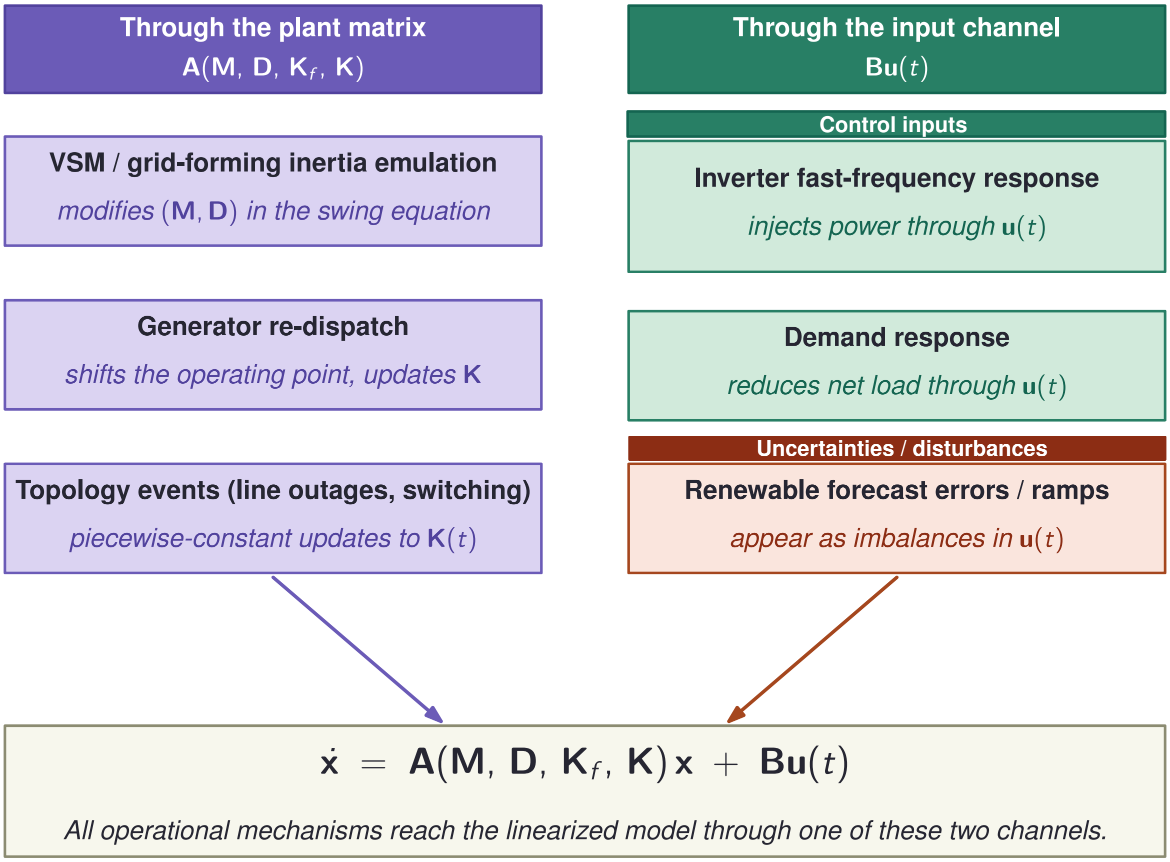

Figure 2.9: Operational mechanisms and the two channels through which they enter the linearized model (2.22). Control inputs (virtual inertia updates, generator re-dispatch, demand response, and fast-frequency response) enter through one of two channels. Disturbances (line outages and switching events) produce piecewise-constant updates to $\mathbf{K}(t)$ through the plant matrix channel. Uncertainties (renewable variability and forecast errors) enter as exogenous imbalances through $\mathbf{Bu}(t)$. The left column groups mechanisms modifying $\mathbf{A}(\mathbf{M}, \mathbf{D}, \mathbf{K}_f, \mathbf{K})$; the right column groups mechanisms acting through $\mathbf{Bu}(t)$.

As shown in the top of the left column of Figure 2.9, virtual synchronous machine (VSM) and grid-forming inertia emulation modify the effective inertia and damping parameters $(\mathbf{M}, \mathbf{D})$. Inverters configured for these control modes were derived in Section 2.5.2, where the closed-loop equivalence (2.44) shows that the inverter contribution adds directly to M_i and D_i at the bus where it is deployed. Buses equipped with such inverters constitute the eligible set $\mathcal{G}_{\text{inv}}$ of the allocation problem (2.41), and the per-bus bound $H_{\max}$ in (2.39) reflects the device-level current rating that limits how much sustained virtual inertia each unit can deliver.

The middle entry of the left column captures generator re-dispatch. Operating on minute-level time scales, re-dispatch adjusts mechanical input power P_{m_i} in (2.10) and shifts the equilibrium (δ^*, V^*) at which the synchronizing coefficients $\mathbf{K}$ defined in (2.20) are evaluated. The effect on the linearized model is structural rather than instantaneous: by the time re-dispatch has taken effect, the system has settled to a new operating point, and the linearization is re-evaluated with an updated $\mathbf{K}$ that reflects the changed power-flow pattern.

The bottom entry of the left column accounts for discrete topology events such as line outages and switching. These events change the set of active edges in the network and therefore the entries of $\mathbf{K}$ directly, producing piecewise-constant updates to $\mathbf{K}(t)$ between decision intervals. Unlike re-dispatch, these updates are not gradual: each contingency produces a step change in the grid's synchronizing structure that the modal spectrum of $\mathbf{A}$ inherits. The tolerance limits of this treatment, specifically the contingency severity beyond which the connectivity assumption and the first-order sensitivity estimates cease to hold, are discussed in Section 3.2.

Crossing to the right column of Figure 2.9, the top entry is converter-based fast-frequency response. Operating on the swing time scale, it injects active power directly into $\mathbf{u}(t)$ in response to measured frequency deviations, providing immediate corrective acceleration without modifying the plant matrix. This is the most agile actuation channel in the model and the natural counterpart to slower mechanisms in the left column.

The middle entry in the right column is demand response. Modulating controllable loads reduces the net electrical demand seen by the system, and in the swing equation (2.10), this appears as a corrective contribution to $u_i(t)$. Functionally, demand response and fast-frequency response inject the same kind of signal at a bus; what differs is the resource that delivers it. Both fold into the same channel and combine with the inertia-related actions on the left column.

The bottom entry in the right column captures renewable variability and forecast errors. Short-term ramps in wind and solar output appear as imbalances on the right-hand side of the swing equation, entering through $\mathbf{u}(t)$ as exogenous disturbances rather than as control inputs. They are not actions an operator chooses; they are perturbations that the framework must absorb. The same channel that carries fast frequency response and demand response also carries the disturbances that those actions are designed to counteract.

The two columns of Figure 2.9 together account for every operational mechanism considered in the rest of the thesis, and the same two columns account for the persistent uncertainty under which the system operates. Renewable intermittency and short-term forecast errors enter through the right column as exogenous power imbalances. Scheduled re-dispatch, commitment changes, and environmental events such as wildfire smoke or extreme cold (Figures 1.1 and 1.2) enter through the left column as shifts in the operating point and discrete updates to $\mathbf{K}(t)$. Both classes of uncertainty force the linearization itself to be re-evaluated as the system evolves, which is the operational reason a closed-loop framework is needed in the first place.

The next chapter builds on this linearized model and the reference allocation H_{ref} developed here. Where this chapter treats virtual inertia as a structural parameter inside $\mathbf{A}$, Chapter 3 treats it as a control action delivered through $\mathbf{u}(t)$. It does so by designing a spatio-temporal physics-informed graph neural network (ST-GNN) and a reinforcement-learning controller that act on the same two channels of Figure 2.9 under topology change and uncertainty.

Chapter 3

Spatio-Temporal Physics-Informed GNN with Reinforcement Learning Control

The electromechanical analysis developed in Chapter 2 shows that small-signal stability in low-inertia transmission systems is governed by quantities that are both spatially distributed and time varying. In the linearized model of (2.22), the state matrix $A(M, D, K_f, K)$ depends explicitly on the inertia distribution, damping, frequency-dependent load effects, and the coupling matrix associated with the prevailing operating point and network topology. As a result, the modal properties used to assess stability, including eigenvalues, damping ratios, and modal sensitivities, are not fixed network characteristics, but vary with loading conditions, dispatch, renewable injections, and topology changes. Consequently, any framework intended to support stability online must account for both temporal evolution of the electrical state and repeated changes in the graph structure of the network.

Furthermore, the stability support cannot be treated as a purely local control problem. The sensitivity expressions derived in Chapter 2 indicate that the effect of bus-level inertia changes depends on the global modal structure of the system, as captured by the eigenstructure of $A(M, D, K_f, K)$. In particular, the bus-level criticality measures and the constrained allocation problem in (2.41) establish that the stabilizing value of virtual inertia depends on where it is placed within the network and how that placement interacts with the current topology. This implies that inertia support should remain explicitly guided by a topology-aware spectral criterion rather than being determined only from local measurements or unconstrained learned correlations.

At the same time, the open-loop formulation of Chapter 2 is not sufficient for operational use under evolving conditions. The analysis there provides modal assessment and a reference virtual inertia profile for a given operating point and topology, but practical transmission-system operation requires these quantities to be updated sequentially as the system evolves. Load and renewable variability alter the system state over time, while outages, switching actions, and reconfiguration events modify the active set of transmission corridors and hence the coupling matrix K. Under these conditions, the stability-relevant quantities derived from the linearized model must be embedded within a closed-loop framework that can repeatedly observe the current state, anticipate its near-term evolution, and coordinate corrective actions over successive decision intervals. Figure 3.1 provides an intuitive interpretation of the transition from the open-loop analysis of Chapter 2 to the closed-loop framework developed here. The left panel represents the nominal operating condition in which the network remains well behaved around the current equilibrium. The middle panel illustrates how a contingency and the associated topology change can redistribute electromechanical stress and alter the locations at which stability support is most valuable. The right panel indicates that the corrective response is not uniform across the network, but must instead be directed to selected buses and updated as the state evolves. For this reason, the stability-support problem requires a framework that combines topology-aware reference generation, prediction of near-term system evolution, and sequential coordination of control actions.

These requirements motivate the structure adopted in this chapter. Since the network state evolves over both space and time, the power system is represented as a spatio-temporal graph whose nodes and edges carry the electrical quantities required for prediction and control. Since the relevant stability information remains encoded in the modal structure of the linearized electromechanical model, the eigenvalue-sensitivity analysis of Chapter 2 is retained as a spectral module that computes a topology-aware reference virtual inertia allocation. Since the operational problem is sequential and involves multiple coupled control actions, including redispatch, demand-side intervention, and inertia adjustment, the control layer is formulated as a reinforcement learning policy acting over repeated decision intervals rather than as a one-step allocation mechanism.

Accordingly, the proposed framework combines three complementary components, each

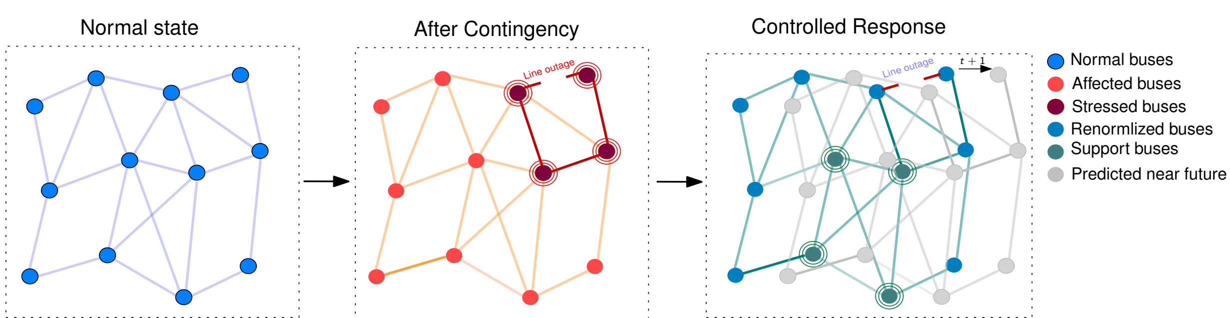

Figure 3.1: Conceptual illustration of the motivation for the proposed closed-loop framework. Under normal conditions, the network operates with balanced electromechanical behaviour. Following a contingency, a topology change redistributes power transfers and concentrates modal stress on a subset of buses and corridors.

introduced by the structure of the problem rather than by algorithmic preference alone. A spectral-sensitivity module evaluates the current operating condition and topology through the linearized model and produces the reference inertia profile H_{ref}. A spatio-temporal physics-informed graph neural network processes recent graph-structured system trajectories and predicts the short-term evolution of key electrical variables. A reinforcement learning policy then determines coordinated control actions using the present state and the predicted near-term system evolution. A supervisory coordination layer reconciles the policy output with the spectral reference and enforces operational feasibility before the executed action is applied to the grid.

Therefore, the linearized electromechanical model in (2.22) remains the physical basis of the formulation, and the sensitivity-guided allocation in (2.41) remains the source of the reference inertia signal. The additional development introduced here is the closed-loop machinery required to use these quantities repeatedly under time-varying operating conditions and network topologies. The chapter therefore extends the open-loop spectral analysis of Chapter 2 into a prediction-and-control architecture for adaptive, topology-aware stability support. Next, Section 3.1 presents the closed-loop system architecture and its operational workflow. Section 3.2 introduces the spatio-temporal graph representation used to model the evolving power network. The subsequent sections then present the control formulation, the spectral-sensitivity-based reference allocation, the spatio-temporal physics-informed graph neural network, and the supervisory coordination mechanism used to combine physics-based guidance with closed-loop decision-making.

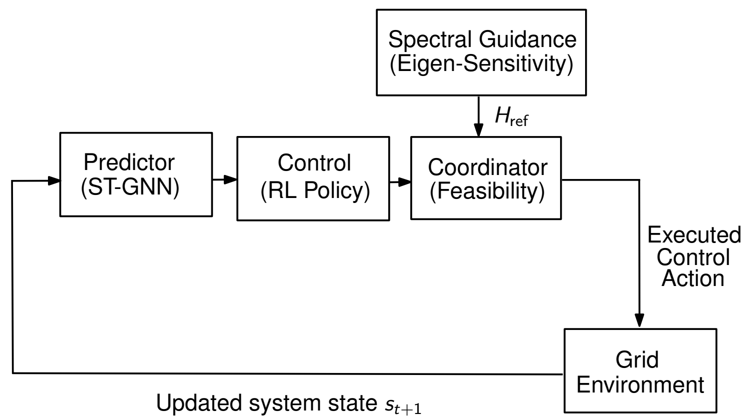

Figure 3.2: Conceptual closed-loop interaction between prediction, control, and the power grid environment.

3.1 Closed-Loop System Architecture

To integrate the identified requirements of physics-based analysis, spatio-temporal system evolution, and continuous reinforcement-learning-based control, we develop a closed-loop framework that operates as an integrated prediction-and-decision system. At each control interval, the framework observes the current grid state, predicts its near-term evolution, and determines coordinated control actions based on both the measured and predicted operating conditions. Rather than treating prediction, optimization, and control as separate processes, the proposed architecture embeds them within a unified feedback loop for topology-aware stability support under evolving operating conditions. This process is illustrated conceptually in Fig. 3.2, where components interact with the grid environment in a closed loop. The operational sequence begins with the grid environment providing the current system state. This state is then processed by the prediction module to estimate the network's near-term evolution. Based on both the observed and predicted states, the control module determines a candidate action, which the coordinator then verifies to ensure feasibility and operational constraints are met before execution. The resulting system response forms the updated state for the next decision interval, thereby completing the closed-loop cycle.

The grid environment block in Fig. 3.2 is instantiated as the IEEE 118-bus transmission network. It simulates system operation under AC power-flow conditions with line-switching contingencies. System evolution follows the electromechanical dynamics introduced in Chapter 2, described by the linearized state-space formulation (2.22), where the matrices M, D, and K characterize the inertia distribution, damping properties, and network coupling structure, respectively. At each time step, the environment provides observable bus-level state variables including voltage magnitudes $V_i(t)$, voltage angles $\theta_i(t)$, active power injections $P_i(t)$, and virtual inertia allocations $H_i(t)$, which together constitute the state information available to the control policy. These state variables are derived from the underlying AC power-flow simulation and associated load–generation models within the environment.

As shown in the spectral guidance block of Figure 3.2, the spectral-sensitivity module receives the current network Laplacian $L(t)$ and inertia matrix $M(t)$ and evaluates the eigenvalue sensitivities of the linearized system matrix $A(M, t)$ using the first-order sensitivity formulation in (2.29). These sensitivities are then aggregated into bus-level criticality scores g_i, as defined in (2.30), which quantify how much a marginal change in inertia at bus i affects the damping of the critical electromechanical modes, that is, the dominant low-frequency oscillations in the system. Using these scores, the module solves the budget-constrained virtual inertia allocation problem in (2.41) to obtain a reference inertia profile H_{ref}, thereby identifying the buses where additional inertia provides the strongest stabilizing influence on the least-damped modes.

In the predictor in Fig. 3.2, the spatio-temporal graph neural network (ST-GNN) processes a short history of system states $(s_{t-k}, \ldots, s_t)$ together with the current network topology. It interprets each time step as a graph (with buses as nodes and transmission lines as edges), so the input becomes a sequence of graph-structured observations. By learning spatial dependencies across buses and temporal correlations in system evolution, the ST-GNN produces one-step-ahead predictions of key grid variables, including bus voltage magnitudes $\hat{V}_i(t+1)$, voltage angles $\hat{\theta}_i(t+1)$, demand $\hat{D}_i(t+1)$, generation $\hat{G}_i(t+1)$, and virtual inertia $\hat{H}_i(t+1)$. These predictions provide situational awareness of the system's near-future operating conditions under evolving load patterns, control actions, and topology changes.

In the control layer of Fig. 3.2, the reinforcement learning (RL) policy receives the current system state s_t along with the ST-GNN predictions and outputs a real-valued (continuous) action vector a_t, meaning that each entry of a_t is a numerical control signal (generator dispatch, demand response or load shed, or seconds of virtual inertia) rather than a discrete on or off choice; these signals specify the next-step generator dispatch, demand-response levels, load-shedding amounts, and virtual inertia allocations across the network.

The coordinator in Fig. 3.2 is a supervisory coordination layer that combines the RL proposal with the spectral reference inertia profile H_{ref} using a blending-and-projection mechanism that enforces operational constraints, including inertia bounds, generation limits, and power-balance feasibility. The resulting feasible control action is then applied to the grid environment, which evolves to a new system state and returns a reward signal reflecting both economic costs and stability objectives. This closed-loop interaction repeats at each decision interval, enabling adaptive placement of virtual inertia and coordinated grid control under varying operating conditions and network topologies.

To complement the conceptual closed-loop interaction shown in Fig. 3.2, Algorithm 3.1 summarizes the operational workflow executed at each decision interval. The algorithm provides a stepwise representation of the interaction among prediction, physics-guided allocation, control, and system evolution within the closed-loop framework. Algorithm 3.1

Algorithm 3.1 Closed-Loop Decision Cycle of the Framework

Input: Time-varying grid $\mathcal{G}(t)$, feasible sets $\mathcal{U}$ and $\mathcal{H}$

Initialize: node features x_t, edge features e_t, inertia state $\mathbf{H}_t$

for each control interval t **do**

  Observe current state $s_t = \{x_t, e_t\}$

  Compute embeddings; predict $\hat{V}_i^{(t+1)}$, $\hat{\delta}_i^{(t+1)}$, and $\hat{H}_i^{(t+1)}$

  Compute physics-guided reference inertia allocation $\mathbf{H}_{\text{ref}}(t)$

  Determine control action a_t

  Extract inertia adjustment $\Delta\mathbf{H}_t$ from a_t

  Enforce feasibility constraints to obtain executed inertia $\mathbf{H}_{t+1} \in \mathcal{H}$

  Apply the coordinated control action, including $\mathbf{H}_{t+1}$, to the grid environment

  Observe next state s_{t+1} and reward r_t

  Record transition data (s_t, a_t, r_t, s_{t+1}) for analysis and evaluation

end for

makes explicit the temporal ordering of the decision variables and the closed-loop interaction between the predictive, physics-guided, control, and grid-environment layers at each interval t. In particular, the grid environment acts as the physical system within the feedback loop. Following execution of the coordinated control action, the system evolves to the next operating condition and returns the updated state s_{t+1} together with the reward signal r_t, which are then fed back into the subsequent decision interval. In contrast to the conceptual loop shown in Fig. 3.2, the algorithm introduces the time-indexed variables s_t, a_t, r_t, s_{t+1}, and $\mathcal{G}(t)$, thereby clarifying how the framework transitions from one operating point to the next under evolving topology and control actions. This formal stepwise representation provides the basis for the graph-based Markov decision process formulation introduced in the following section.

Next, the spatio-temporal graph representation of the power network used in the proposed framework is introduced; it provides the basis for the subsequent state, action, and reward formulations.

3.2 Spatio-Temporal Graph Representation

To enable the learning and control framework to process both the physical structure of the power network and its temporal evolution, the transmission grid is represented as a spatio-temporal graph. In this representation, each bus is treated as a node, and each transmission line is treated as an edge connecting two nodes. This graph-based view allows the framework to capture how electrical quantities evolve across both space, through network connectivity, and time, through changing operating conditions, contingencies, and switching events. The resulting representation provides the mathematical foundation for the ST-GNN predictor and the subsequent control formulation.

Formally, the power grid at time t is represented as

$$\mathcal{G}(t) = (\mathcal{V}, \mathcal{E}(t)),$$

where $\mathcal{V}$ denotes the set of buses in the transmission network and $\mathcal{E}(t)$ denotes the set of transmission lines that are in service at time t.

To mathematically describe how buses are interconnected, the network topology is represented using an adjacency matrix $A(t) \in \mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|}$. This matrix records which buses are directly connected by transmission lines at time t. Specifically, $A_{ij}(t) = 1$ if a transmission line between buses i and j is in service, and $A_{ij}(t) = 0$ otherwise.

From this connectivity structure, the corresponding degree matrix $D(t)$ is constructed, where each diagonal entry $D_{ii}(t) = \sum_j A_{ij}(t)$ gives the number of lines connected to bus i.

Using these two matrices, the combinatorial graph Laplacian is defined as

$$L(t) = D(t) - A(t). \tag{3.1}$$

This Laplacian provides a compact mathematical description of how each bus interacts with its neighbouring buses through the network structure.

For power-system analysis, the interaction between buses also depends on the electrical characteristics of the transmission lines, rather than only on whether a connection exists. To capture this, a susceptance-weighted Laplacian $L_B(t)$ is introduced, whose off-diagonal entries are given by $(L_B)_{ij}(t) = -B_{ij}$ for $i \neq j$, while the diagonal entries satisfy $(L_B)_{ii}(t) = \sum_{k \neq i} B_{ik}$. This weighted Laplacian captures both the physical network connectivity and the electrical strength of coupling between buses.

Due to changes in the transmission network during operation, events such as line outages, switching actions, and contingency responses are represented as updates to the graph connectivity over time. In this work, these changes are modelled as piecewise-constant updates to the adjacency matrix $A(t)$ at discrete-event times, so that the network topology is updated only when a switching, outage, or contingency event occurs. Each topology update changes the susceptance-weighted Laplacian $L_B(t)$ and, consequently, the electrical coupling matrix $K(t)$ introduced in Chapter 2. Changes in the physical network configuration therefore directly alter the electromechanical interactions among buses.

The framework is designed to handle N$-$1 and N$-$2 contingencies, which represent the loss of one or two transmission elements. Beyond this range, such as simultaneous loss of three or more elements (N$-x$, $x > 2$), the graph Laplacian update may disconnect portions of the network, violating the connectivity assumption underlying the Lyapunov stability guarantee of Theorem 3.1. In such cases, the spectral allocator's first-order sensitivity estimates also become unreliable, as the linearization point shifts abruptly. The N$-$2 boundary therefore reflects a practical operating envelope rather than a fundamental algorithmic limit, and extension to more severe contingencies would require explicit multi-island control logic.

As a result, the linearized swing dynamics can be written in terms of the time-varying state matrix

$$A(M, D, K(t)) = \begin{bmatrix} 0 & I \\ -M^{-1}K(t) & -M^{-1}D \end{bmatrix}, \tag{3.2}$$

where M denotes the diagonal inertia matrix and D denotes the damping matrix. The eigenstructure of $A(M, D, K(t))$ determines the oscillatory stability behaviour of the network, including the damping of critical electromechanical modes.

To represent the evolving operating condition of the grid at each decision interval, every bus $n \in \mathcal{V}$ is associated with a node-level feature vector

$$x_t^n = \{P_{d,t:t+T}^n,\ Q_{d,t:t+T}^n,\ P_{g,t}^n,\ Q_{g,t}^n,\ V_t^n,\ \delta_t^n,\ H_t^n\}, \tag{3.3}$$

which collects the key electrical quantities associated with each bus, including the forecasted active and reactive power demands $P_{d,t:t+T}^n$ and $Q_{d,t:t+T}^n$ over the prediction horizon T, the active and reactive generation injections $P_{g,t}^n$ and $Q_{g,t}^n$, the bus voltage magnitude V_t^n, the voltage phase angle δ_t^n, and the effective inertia constant H_t^n applied at bus n. In this work, the prediction horizon is set to $T = 6$ time steps, corresponding to 30 minutes at a 5-minute control resolution.

The graph-structured system state at time t is therefore defined by the pair

$$s_t = \{\mathcal{G}(t), x_t\}, \tag{3.4}$$

where $\mathcal{G}(t)$ captures the time-varying network topology and $x_t = \{x_t^n\}_{n\in\mathcal{V}}$ collects the bus-level operating variables. This graph-based state representation serves as the input to the ST-GNN predictor and as the system state for the subsequent control formulation.

In addition to this operational graph representation, graph-structural descriptors are introduced later for post-hoc interpretability analysis of the trained ST-GNN. These descriptors are used to examine how the model's predicted virtual inertia allocation depends on network structure and are not part of the online state definition.

3.3 Spatio-Temporal Physics-Informed Graph Neural Network

The spatio-temporal physics-informed graph neural network (ST-GNN) serves as the predictive core of the proposed framework and operates on the graph-structured state representation introduced in Section 3.2. It processes a temporal sequence of historical grid states together with the evolving network topology to generate one-step-ahead predictions of key state variables required for control and stability assessment. Specifically, the ST-GNN consumes a history window of T recent system states $\{s_{t-T+1}, \ldots, s_t\}$ and produces one-step-ahead predictions of bus voltage magnitudes, voltage angles, virtual inertia allocations, and auxiliary physical quantities required by the physics-informed training objective and the subsequent control and coordination stages.

Two boundary conditions arise in practice. At system startup or after a cold restart, the history window $\{s_{t-T+1}, \ldots, s_t\}$ is not fully populated. In this case, the missing entries are padded with the steady-state nominal operating point, providing a neutral initialization that avoids arbitrary extrapolation. When a previously lost transmission line reappears, the adjacency matrix $A(t)$ is updated at the next decision interval to reflect the restored connection. Entries in the history window that predate the restoration retain the pre-reappearance topology, introducing a brief transient in the spatial embeddings until the window slides past the switching event. This transient is bounded in duration by the window length T and is acknowledged as a limitation of the fixed-window architecture.

To realize this mapping from historical graph states to predicted system variables, the ST-GNN is decomposed into three components: a spatial graph encoder, a temporal encoder, and a set of task-specific prediction heads. The spatial graph encoder extracts topology-aware node embeddings at each time step, the temporal encoder processes the resulting embedding sequence over the history window, and the prediction heads map the aggregated latent representation to the required outputs.

3.3.1 Architecture Overview and Design Rationale

Figure 3.3 summarizes the construction of the network and the flow of information through it. The ST-GNN architecture consists of three sequentially composed components.The spatial graph encoder applies a stack of graph convolutional network (GCN) layers to extract topology-aware node embeddings at each time step, propagating information across neighbouring buses according to the normalized adjacency operator. The temporal encoder, implemented as a Temporal Convolutional Network (TCN), processes the ordered sequence of spatial embeddings over the history window to capture the dynamic evolution of grid states. Finally, a set of task-specific prediction heads maps the aggregated latent representation to the required output quantities: bus voltage magnitudes, voltage angles, virtual inertia allocations, and auxiliary physical quantities used in the physics-informed training objective.

The GCN backbone was selected over alternative aggregation schemes, including Graph Attention Networks (GAT) and Graph Isomorphism Networks (GIN), based on empirical evaluation on the IEEE 118-bus system. On a fixed and known topology, symmetric degree-normalized aggregation is sufficient to capture the spatial coupling structure, and learned attention weights or injective aggregation provide no measurable benefit while increasing model complexity. The TCN is preferred over recurrent architectures for the temporal component due to its parallelizable training, stable gradient behaviour over long history windows, and competitive accuracy on sequential grid-state data.

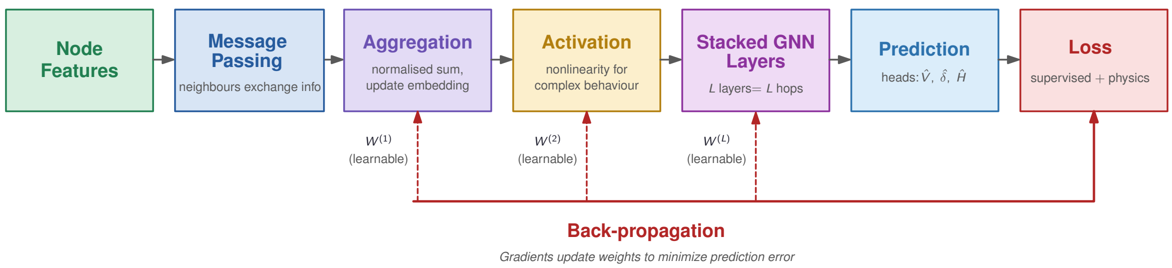

Figure 3.3: Construction of the graph neural network used in the ST-GNN predictor. Node features are propagated through message passing, degree-normalized aggregation, and nonlinear activation, with L stacked layers giving each bus an L-hop receptive field over the network. The prediction heads output V, δ, and H, and the composite supervised and physics-informed loss in (3.13) drives back-propagation, which updates the learnable weight matrices $W^{(\ell)}$ at each layer.

For each time step t, let $X_t \in \mathbb{R}^{|\mathcal{V}|\times F}$ denote the node-feature matrix and let $A_t \in \mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|}$ denote the adjacency matrix of the corresponding graph. Define the self-loop-augmented adjacency matrix $\tilde{A}_t = A_t + I$ and its associated degree matrix $\tilde{D}_t = \mathrm{diag}\big(\sum_j (\tilde{A}_t)_{ij}\big)$. The normalized graph propagation operator is then given by

$$S_t = \tilde{D}_t^{-1/2}\tilde{A}_t\tilde{D}_t^{-1/2}.$$

The spatial graph encoder is implemented as a stack of graph-convolution layers. Setting $H_t^{(0)} = X_t$, the layer-wise propagation is written as

$$H_t^{(\ell+1)} = \sigma\left(S_t H_t^{(\ell)} W_S^{(\ell)}\right), \qquad \ell = 0, \ldots, L_S - 1, \tag{3.5}$$

where $W_S^{(\ell)}$ denotes the trainable weight matrix of the ℓth spatial layer and $\sigma(\cdot)$ denotes the activation function. The final spatial embedding at time t is therefore

$$Z_t = H_t^{(L_S)} = \Phi_S(X_t, A_t). \tag{3.6}$$

The normalization $S_t = \tilde{D}_t^{-1/2}\tilde{A}_t\tilde{D}_t^{-1/2}$ implements symmetric degree scaling of the adjacency operator. Consequently, the contribution transmitted across an edge is weighted by the inverse square root of the degrees of its incident nodes, which balances message propagation across sparsely and densely connected buses, preserves symmetry for undirected graphs, and improves numerical stability. The exponent $-\frac{1}{2}$ thus arises directly from the inverse square-root degree normalization used to construct the symmetric graph filter.

The spatial encoder is applied before the temporal encoder by design. Power network topology imposes a fixed coupling structure on bus interactions: the graph Laplacian encodes how disturbances at one bus propagate to neighbouring buses through transmission lines. Resolving these spatial dependencies first produces node embeddings that are topology-consistent at each time step. The temporal encoder then tracks how these spatially coherent embeddings evolve over the history window. Reversing this order would mix temporal dynamics before the spatial structure is resolved, producing embeddings that conflate network-propagation effects with time-varying load and generation changes, making the learned representations harder to interpret and less generalizable to unseen topologies.

While the spatial encoder captures the network structure at each time step, the grid also exhibits strong temporal dependencies. To model these dynamics, the ordered sequence of spatial embeddings over the history window is denoted by $\mathbf{Z}_t = (Z_{t-T+1}, \ldots, Z_t)$ and is processed by a temporal encoder based on a Temporal Convolutional Network (TCN).

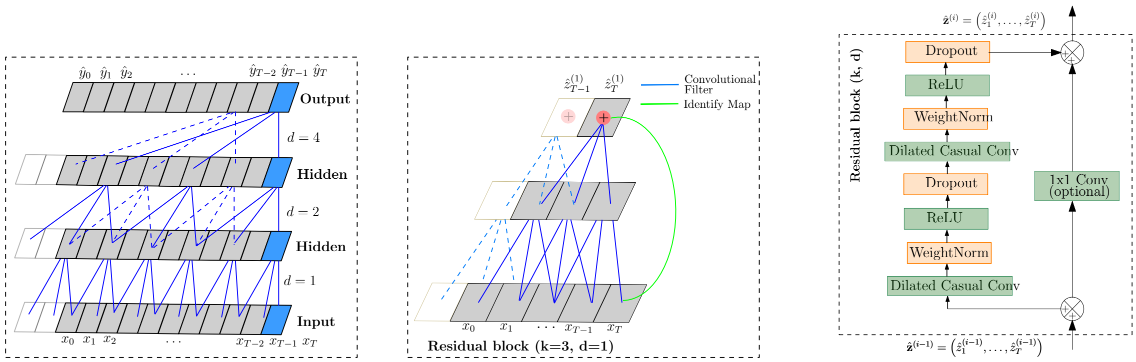

(a) (b) (c)

Figure 3.4: Temporal convolutional structure used to illustrate the temporal encoder Φ_T; (a) dilated causal hierarchy, (b) generic residual block, and (c) residual block example for $k = 3$ and $d = 1$.

Let $\Phi_T(\cdot)$ denote the temporal encoder. The corresponding temporally encoded sequence is then written as

$$\mathbf{U}_t = \Phi_T(\mathbf{Z}_t) = (u_{t-T+1}, \ldots, u_t)\,, \tag{3.7}$$

where each u_τ denotes the temporal representation associated with time index τ. In the present framework, Φ_T is implemented as a causal temporal-convolution stack whose effective receptive field increases with depth, so that the encoded representation at time t depends only on the historical embeddings in the window $(Z_{t-T+1}, \ldots, Z_t)$ and not on future information.

Figure 3.4 clarifies the structure of the temporal encoder used in the proposed ST-GNN. Figure 3.4a shows how causal convolutions with increasing dilation enlarge the effective receptive field over the history window. Figure 3.4b shows the internal composition of a generic residual block, including the dilated causal convolutions, normalization, nonlinearities, dropout, and skip path. Figure 3.4c provides a concrete example for kernel size $k = 3$ and dilation $d = 1$, illustrating how the convolutional filter and identity mapping combine locally. Together, these panels show how the temporal encoder captures both short-range and longer-range temporal dependencies before attention-based aggregation.

The encoded temporal features are then aggregated using an attention mechanism. For each $\tau \in \{t-T+1, \ldots, t\}$, define

$$\alpha_\tau = \frac{\exp\left(q^\top u_\tau\right)}{\sum_{k=t-T+1}^{t} \exp\left(q^\top u_k\right)}, \tag{3.8}$$

where q is a learnable query vector. The final temporally aggregated latent representation is then given by

$$h_{t+1} = \sum_{\tau=t-T+1}^{t} \alpha_\tau \, u_\tau. \tag{3.9}$$

This attention mechanism assigns greater weight to time steps that are most informative for prediction, such as those immediately following a disturbance or topology change.

The aggregated latent representation h_{t+1} is then passed to a collection of task-specific prediction heads. These heads produce the next-step voltage magnitude $\hat{V}_i^{(t+1)}$, voltage angle $\hat{\delta}_i^{(t+1)}$, auxiliary value estimate $\hat{V}^\pi(s_t)$, virtual inertia allocation $\hat{H}_i^{(t+1)}$, predicted demand $\hat{P}_{d,i}^{(t+1)}$, predicted net active-power injection $\hat{P}_i^{\text{inj}}(t+1)$, and predicted line flow $\hat{f}_{ij}(t+1)$. The quantities $\hat{P}_i^{\text{inj}}(t+1)$ and $\hat{f}_{ij}(t+1)$ are auxiliary physics heads introduced to support the physics-informed training objective; they are not required as primary control outputs, but they enable direct enforcement of nodal and corridor-level consistency. Collectively, these outputs provide the predictive quantities required by the subsequent control, coordination, and physical regularization stages.

Figure 3.5 summarizes the overall ST-GNN architecture. Historical graph states over a time window are first provided as input, including the node-feature matrices and the time-varying network topology. These inputs are processed by the spatial graph encoder

Φ_S, implemented as a stack of graph-convolution layers, to produce topology-aware node embeddings at each time step. The resulting embedding sequence is then processed by the temporal encoder Φ_T, implemented as a Temporal Convolutional Network, and subsequently aggregated by an attention mechanism to produce a latent representation for the next prediction step. This latent representation is finally passed to several task-specific heads that output next-step voltage magnitudes and angles, the auxiliary value estimate, the bus-wise virtual inertia allocation, and the auxiliary physical quantities used by the training objective. A physics-informed loss at the output layer encourages these predictions to satisfy basic network laws and operational constraints.

The ST-GNN is trained using a composite objective that combines supervised prediction terms with physics-informed regularization. This design is motivated by the role of the predictor within the closed-loop framework. The network must reproduce the state-transition patterns observed in the training trajectories so that its outputs remain useful for short-horizon prediction and control. At the same time, prediction accuracy alone is not sufficient in power-system applications, because a purely data-driven model may fit historical trajectories while violating fundamental physical laws or operating constraints. The training objective is therefore constructed to enforce both predictive fidelity and physical consistency.

Accordingly, the overall loss is decomposed as

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \mathcal{L}_{\text{phys}}, \tag{3.10}$$

where $\mathcal{L}_{\text{sup}}$ denotes the supervised prediction loss and $\mathcal{L}_{\text{phys}}$ denotes the physics-informed regularization term.

The supervised component is defined as

$$\begin{aligned}\mathcal{L}_{\text{sup}} = \mathbb{E}\Big[&|V_i^{(t+1)} - \hat{V}_i^{(t+1)}|^2 + \lambda|\delta_i^{(t+1)} - \hat{\delta}_i^{(t+1)}|^2 \\ &+ \mu|V^\pi(s_t) - \hat{V}^\pi(s_t)|^2 + \nu|H_{\text{ref},i}^{(t+1)} - \hat{H}_i^{(t+1)}|^2 \\ &+ \beta|P_{d,i}^{(t+1)} - \hat{P}_{d,i}^{(t+1)}|^2\Big],\end{aligned} \tag{3.11}$$

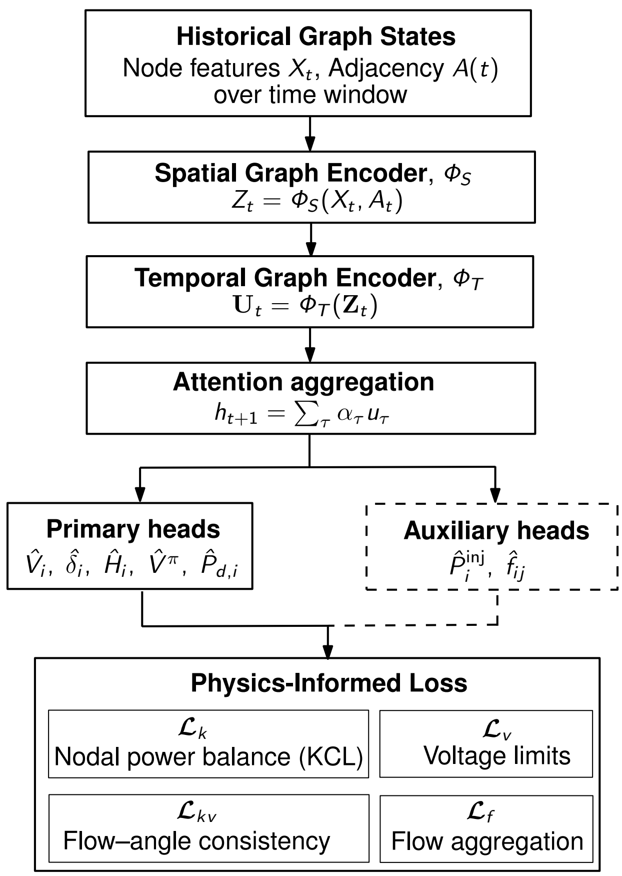

Figure 3.5: Architecture of the proposed spatio-temporal physics-informed graph neural network (ST-GNN). The forward path consists of a spatial graph encoder Φ_S, a temporal encoder Φ_T, attention-based temporal aggregation, and task-specific prediction heads. The auxiliary physics heads support the physics-informed loss, while saliency analysis is performed post hoc on the trained model. The rationale for performing spatial encoding prior to temporal encoding is discussed in Section 3.3.

where $H_{\mathrm{ref},i}^{(t+1)}$ denotes the reference inertia allocation provided by the spectral-sensitivity module introduced later in Section 3.4. The terms in (3.11) penalize errors in next-step voltage magnitude, voltage angle, auxiliary value estimation, reference virtual inertia tracking, and demand prediction, respectively. The coefficients λ, μ, ν, and β balance the relative influence of these learning targets. The physics-informed component is written as

$$\mathcal{L}_{\mathrm{phys}} = \alpha_k \mathcal{L}_k + \alpha_v \mathcal{L}_v + \alpha_f \mathcal{L}_f + \alpha_{kv} \mathcal{L}_{kv}, \tag{3.12}$$

where α_k, α_v, α_f, and α_{kv} are nonnegative weighting coefficients. These terms penalize deviations from power-balance relations and operating limits, thereby encouraging the predictor to remain compatible with the electromechanical structure of the network even when the operating condition differs from those observed during training.

Combining (3.11) and (3.12), the overall training objective becomes

$$\begin{aligned} \mathcal{L} = \mathbb{E}\Big[& |V_i^{(t+1)} - \hat{V}_i^{(t+1)}|^2 + \lambda |\delta_i^{(t+1)} - \hat{\delta}_i^{(t+1)}|^2 \\ & + \mu |V^{\pi}(s_t) - \hat{V}^{\pi}(s_t)|^2 + \nu |H_{\mathrm{ref},i}^{(t+1)} - \hat{H}_i^{(t+1)}|^2 \\ & + \beta |P_{d,i}^{(t+1)} - \hat{P}_{d,i}^{(t+1)}|^2 \Big] \\ & + \alpha_k \mathcal{L}_k + \alpha_v \mathcal{L}_v + \alpha_f \mathcal{L}_f + \alpha_{kv} \mathcal{L}_{kv}. \end{aligned} \tag{3.13}$$

First, to enforce nodal active-power consistency, a penalty term based on Kirchhoff-type active-power balance is introduced as

$$\mathcal{L}_k = \sum_{i \in \mathcal{V}} \left| \hat{P}_i^{\mathrm{inj}}(t+1) - \sum_{j \in \mathcal{N}(i)} B_{ij} \big(\hat{\delta}_i^{(t+1)} - \hat{\delta}_j^{(t+1)} \big) \right|^2, \tag{3.14}$$

where $\mathcal{N}(i)$ denotes the set of buses adjacent to bus i and B_{ij} denotes the line susceptance.

This term penalizes mismatches between the predicted net nodal injection and the network power transfer implied by the predicted voltage-angle differences. It therefore discourages state predictions that violate nodal power balance.

Next, to enforce consistency between predicted line flow and angle difference across each transmission corridor, a DC power-flow consistency term is introduced as

$$\mathcal{L}_{kv} = \sum_{(i,j)\in\mathcal{E}} \left| \hat{f}_{ij}(t+1) - B_{ij}\big(\hat{\delta}_i^{(t+1)} - \hat{\delta}_j^{(t+1)}\big) \right|^2 . \tag{3.15}$$

This term penalizes discrepancies between the predicted line-flow quantities and the flows implied by the predicted bus-angle differences, thereby enforcing consistency between nodal states and corridor-level behaviour.

In addition, the voltage-limit penalty

$$\mathcal{L}_v = \sum_{i\in\mathcal{V}} \Big[\max(0, V_{\min} - \hat{V}_i^{(t+1)})^2 + \max(0, \hat{V}_i^{(t+1)} - V_{\max})^2 \Big] \tag{3.16}$$

penalizes predicted voltages that fall below the lower operating bound $V_{\min}$ or exceed the upper operating bound $V_{\max}$. This term encodes feasibility with respect to standard voltage-security limits and discourages the predictor from producing physically implausible voltage levels.

Finally, to preserve consistency between bus-level injection and corridor-level flow descriptions, a flow-aggregation penalty is introduced as

$$\mathcal{L}_f = \sum_{i\in\mathcal{V}} \left| \hat{P}_i^{\mathrm{inj}}(t+1) - \sum_{j\in\mathcal{N}(i)} \hat{f}_{ij}(t+1) \right|^2 . \tag{3.17}$$

This term links the auxiliary nodal and edge-level physics heads and improves internal consistency of the learned electrical representation.

Taken together, the supervised terms in (3.11) and the physics-informed penalties in (3.12) ensure that the ST-GNN is trained not only to fit observed trajectories, but also to respect the structural and operational constraints of the underlying power system. This is particularly important in the present framework, since the predicted states are subsequently used by the control layer; physically inconsistent predictions would otherwise propagate directly into the closed-loop decision process.

3.3.2 Post-Hoc Saliency Analysis of the Trained ST-GNN

In addition to predictive accuracy, it is also important to examine how the trained ST-GNN uses network structure to generate its virtual inertia predictions. While the spectral-sensitivity allocator provides a physics-grounded interpretation of inertia placement, the learned model provides an implicit, data-driven mapping from graph structure and operating conditions to predicted inertia. To analyze this structural dependence, a gradient-based saliency framework is employed as a post-hoc interpretability tool.

Because the ST-GNN predicts bus-wise virtual inertia through its inertia head, the saliency analysis focuses on identifying the structural properties of the transmission network that most strongly influence these predictions. In particular, the objective is to determine which transmission corridors exhibit the greatest sensitivity of the learned inertia allocation to perturbations in graph-structural descriptors.

To expose this structural sensitivity, each bus $i \in \mathcal{V}$ is associated with a vector of graph-theoretic descriptors

$$\mathbf{z}_i = [d_i, c_i, b_i]^\top,$$

where $d_i = \deg(i)/\max_{j\in\mathcal{V}}\deg(j)$ denotes the normalized degree, c_i denotes the closeness centrality, and b_i denotes the betweenness centrality, all computed on the undirected bus-line graph. These descriptors summarize the local and mesoscopic structural role of each bus within the network.

Let the trained multitask ST-GNN produce the vector of bus-wise virtual inertia predictions

$$\hat{\mathbf{H}}(\mathcal{G}) = [\hat{H}_i]_{i\in\mathcal{V}}.$$

To quantify how this predicted inertia allocation depends on the structural descriptors, define the scalar functional

$$\ell_H(\mathcal{G}) = \sum_{i\in\mathcal{V}} \hat{H}_i, \tag{3.18}$$

which aggregates the predicted inertia over the network.

The node-level saliency vector at bus i is then defined as the gradient of this functional with respect to the descriptor vector of that bus,

$$\mathbf{g}_i = \nabla_{\mathbf{z}_i}\ell_H(\mathcal{G}) \in \mathbb{R}^3. \tag{3.19}$$

The magnitude of this gradient reflects how strongly variations in the structural role of bus i influence the model's predicted inertia distribution.

To identify structurally critical transmission corridors, an edge-level saliency score is assigned to each transmission line $(i,j)\in\mathcal{E}$ as

$$s_{ij} = \|\mathbf{g}_i - \mathbf{g}_j\|_2. \tag{3.20}$$

Large values of s_{ij} indicate interfaces at which small perturbations in the graph-structural

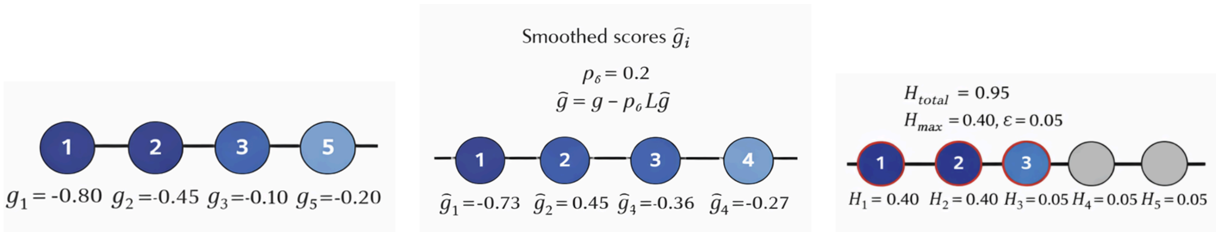

(a) (b) (c)

Figure 3.6: Illustrative 5-bus example of the spectral-sensitivity-based virtual inertia allocation procedure: (a) raw bus-level sensitivity scores obtained from aggregated eigenvalue sensitivities, (b) Laplacian-smoothed topology-aware scores, and (c) the resulting constrained reference inertia allocation H_{ref}.

descriptors at the two endpoints produce markedly different effects on the aggregated inertia prediction. In this way, edges with high saliency identify transmission corridors where the learned model is most structurally sensitive in its virtual inertia allocation.

These edge-level saliency scores are visualized as saliency maps over the network graph, enabling the identification of structurally influential corridors in the IEEE 118-bus system. The corresponding results are presented in Chapter 4.

3.4 Spectral-Sensitivity for Virtual Inertia Allocation

The virtual inertia allocation module provides the physics-guided reference signal used by the closed-loop framework. Its role is not to replace the sequential controller, but to generate, at each operating condition, a stability-oriented inertia profile derived directly from the electromechanical model of Chapter 2. In that chapter, the small-signal dynamics were written in the form $\dot{x} = A(M, D, K(t))x + Bu$, where the inertia matrix M, damping matrix D, and topology-dependent coupling matrix $K(t)$ jointly determine the modal properties of the system. Since the critical electromechanical modes are governed by the eigenstructure of $A(M, D, K(t))$, any reference allocation of virtual inertia should be derived from how localized inertia changes shift those modes rather than from a purely heuristic spatial rule. This motivates the use of the eigenvalue sensitivity analysis derived in Chapter 2. The objective here is to quantify how a marginal change in inertia at bus i affects the damping of the critical oscillatory modes and to convert that information into a topology-aware allocation signal. The overall procedure is illustrated conceptually in Fig. 3.6 using a simple 5-bus example.

In the electromechanical model, virtual inertia enters through the effective inertia parameter at a bus. Equivalently, through the mapping between inertia constant and swing-equation coefficient established previously, adjustments in H_i modify the corresponding entry of the matrix M. Physically, this virtual inertia is realized by inverter-based active-power control acting on a fast response time scale relative to the dominant electromechanical oscillations. The role of the allocator is therefore to specify where fast synthetic inertial support should be concentrated so that the least-damped modes are improved most effectively under the current topology and operating point.

The sensitivity of an electromechanical eigenvalue λ_k with respect to the inertia parameter M_i is obtained from first-order perturbation theory applied to the linearized state matrix $A(M, D, K(t))$ defined in (3.2). Using the closed-form result derived from the Chapter 2 analysis, the modal sensitivity is written as

$$\frac{\partial \lambda_k}{\partial M_i} = \frac{1}{M_i^2} \frac{(v_{k,i}^{L,\omega})^* \left[(K v_k^{R,\theta})_i + D_i (v_{k,i}^{R,\omega}) \right]}{(v_k^L)^\dagger v_k^R}. \tag{3.21}$$

This expression quantifies the first-order movement of mode k under a localized inertia perturbation at bus i. In particular, its real part indicates whether an increase in inertia shifts the corresponding mode leftward in the complex plane, thereby improving modal damping, or rightward, thereby degrading it.

To identify the buses at which inertia placement most effectively improves system stability, the real parts of these sensitivities are aggregated over the set of critical oscillatory modes $\mathcal{K}$, yielding the bus-level sensitivity score

$$g_i = \sum_{k \in \mathcal{K}} \mathrm{Re} \left\{ \frac{\partial \lambda_k}{\partial M_i} \right\}. \tag{3.22}$$

Negative values of g_i indicate that increasing inertia at bus i shifts the critical eigenvalues further into the left half of the complex plane and therefore improves the damping of the dominant electromechanical modes. In this sense, g_i provides a direct physics-based measure of the stabilizing value of inertia at each bus. A simple illustration of these raw bus-level scores is shown in Fig. 3.6a.

Because the critical modes are shaped not only by individual buses but also by the network corridors that couple them, a purely pointwise use of g_i may produce spatially abrupt allocations. To incorporate topological coherence, the raw sensitivity field is regularized using the graph Laplacian:

$$\tilde{g} = g - \rho_s L_g g, \tag{3.23}$$

where L_g is the bus-connectivity Laplacian and ρ_s is a smoothing coefficient. This operation discourages sharp variations across electrically connected buses and therefore produces a topology-aware allocation signal that better reflects corridor-level electromechanical structure. The effect of this smoothing step is illustrated in Fig. 3.6b.

Using the regularized sensitivity scores, the reference virtual inertia allocation is obtained by solving the constrained optimization problem

$$
\begin{aligned}
\min_{H} \quad & \sum_{i \in \mathcal{G}_{inv}} \tilde{g}_i H_i \\
\text{subject to} \quad & \sum_{i \in \mathcal{G}_{inv}} H_i \leq H_{total}, \\
& \epsilon_M \leq H_i \leq H_{max}, \qquad i \in \mathcal{G}_{inv}, \\
& H_i = H_i^{real}, \qquad i \in \mathcal{G}_{gen}.
\end{aligned}
\tag{3.24}
$$

Here, $\mathcal{G}_{inv}$ denotes the set of inverter-interfaced buses eligible for virtual inertia support, $\mathcal{G}_{gen}$ denotes the set of synchronous-generator buses, H_{total} is the total virtual inertia budget, and H_{max} is the per-bus upper bound. The lower bound ϵ_M ensures a strictly positive admissible inertia level. Because the objective weights each candidate bus by its topology-regularized stabilizing sensitivity, the solution prioritizes allocating inertia to those locations that provide the strongest first-order improvement in modal damping. The resulting reference allocation for the illustrative example is shown in Fig. 3.6c.

It is important to distinguish the role of this allocator from that of the RL controller. The optimization in (3.24) produces a reference profile H_{ref} by solving a fast, quasi-static allocation problem at the current operating condition. It is therefore a model-based guidance mechanism defined on the electromechanical time scale of the linearized system. By contrast, the RL controller operates sequentially over time, coordinating inertia adjustments together with redispatch and load-control actions under uncertainty. The allocator thus provides the stability-oriented baseline, while the controller determines how that baseline should be tracked, blended, or adjusted sequentially.

The reference inertia allocation H_{ref} obtained from solving (3.24) is incorporated into the closed-loop framework in two ways. First, it is embedded in the graph-based state representation so that the predictive and control modules remain informed by the current physics-guided target allocation. Second, it serves as a reference signal within the coordination mechanism shown later in Fig. 3.7, where the reinforcement-learning policy's proposed inertia adjustments are blended with the spectral reference before projection onto the feasible set. In this way, the spectral allocator anchors the closed-loop policy to the small-signal stability structure derived from the electromechanical model.

Therefore the reference profile H_{ref} can act as a topology-aware, physics-guided baseline for control decisions aimed at maintaining system stability under evolving operating conditions.

3.5 Markov Decision Process Formulation

The closed-loop control problem is formulated as a discrete-time Markov Decision Process (MDP) in which the controller observes the current graph-structured grid state and selects coordinated control actions over successive decision intervals. This formulation is not introduced independently of the electromechanical model derived in Chapter 2; rather, it provides a sequential decision-making layer around the topology-dependent dynamics described by the linearized state-space model $\dot{x} = A(M, D, K)x + Bu$. In the present framework, the MDP state transition represents the sampled evolution of these dynamics under changing operating conditions, topology updates, and exogenous disturbances.

It is important to clarify the role of the physics model within the DRL framework. The system model derived in Chapter 2 is not used by the PPO policy at inference time; the policy maps directly from observed state s_t to action a_t without querying the swing equations. Instead, the physics model enters the learning loop in three ways: it defines the simulation environment in which the agent trains, it shapes the reward signal through the stability penalty r_t^{stab} which uses eigenvalue-based damping metrics, and it provides the spectral reference H_{ref} that regularizes the inertia component of the policy output through the blending-and-projection coordinator. The spectral guidance in Fig. 3.8 therefore enters the DRL loop as a shaped reference signal, not as a policy input or a differentiable model. This makes the framework model-informed rather than model-based in the strict sense.

Formally, the control problem is defined by the tuple $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{S}$ denotes the state space, $\mathcal{A}$ the action space, P the state-transition kernel, R the reward function, and $\gamma \in (0, 1)$ the discount factor.

The system state at decision interval t is given by

$$s_t = \{\mathcal{G}(t), x_t\}, \tag{3.25}$$

where $\mathcal{G}(t)$ captures the time-varying network topology and x_t collects the bus-level operating variables defined in Section 3.2. Through this definition, the MDP state retains both the electrical operating condition and the topology-dependent structure that influences the electromechanical spectrum through $K(t)$.

The control action is defined at the bus level, reflecting the distributed nature of grid support. For each bus $n \in \mathcal{V}$, the action vector is

$$a_t^n = \{P_{g,t+1}^n,\ Q_{g,t+1}^n,\ P_{t+1}^{sh,n},\ Q_{t+1}^{sh,n},\ H_{t+1}^n\}, \tag{3.26}$$

where $P_{g,t+1}^n$ and $Q_{g,t+1}^n$ denote active and reactive redispatch at the next operating step, $P_{t+1}^{sh,n}$ and $Q_{t+1}^{sh,n}$ denote controllable load-shedding actions, and H_{t+1}^n denotes the commanded effective virtual inertia level. Generator buses in $\mathcal{G}_{gen}$ retain their physical inertia, whereas only inverter-interfaced buses in $\mathcal{G}_{inv}$ are eligible for virtual inertia adjustment. In this sense, H^n_{t+1} should be interpreted not as an instantaneous change in physical inertia, but as a control setpoint to be realized by inverter-based fast frequency support over the subsequent control interval.

Demand response enters the framework as a projection only constraint rather than as a primary control objective. The load-shedding components P^{sh} and Q^{sh} in (3.26) allow the RL policy to modulate controllable loads within feasibility bounds, consistent with the demand-response channel described in Chapter 2. This treatment keeps demand response within the unified actuation framework without elevating it to a separate control layer, preserving the consistency of the overall architecture.

The system evolves according to the stochastic transition

$$s_{t+1} \sim P(\,\cdot \mid s_t, a_t), \tag{3.27}$$

which captures the sampled AC power-flow and electromechanical response of the grid together with stochastic disturbances arising from load variability, renewable fluctuations, and contingency events. The dependence of this transition on a_t reflects the fact that redispatch, load control, and virtual inertia support all influence the subsequent operating point and, through it, the effective damping and stability margins of the system.

The reward function is constructed to encode the two operational objectives that must be balanced in closed-loop operation: economic efficiency and dynamic security. It is written as

$$r_t = r_t^{\text{econ}} + r_t^{\text{stab}}, \tag{3.28}$$

where the economic component penalizes costly redispatch and disruptive curtailment,

$$r_t^{\text{econ}} = -\left[\alpha_g C_g(P_{g,t+1}) + \alpha_s \sum_{n\in\mathcal{V}} P_{t+1}^{sh,n}\right], \tag{3.29}$$

and the stability component penalizes violations of voltage-security and modal-damping requirements,

$$r_t^{\text{stab}} = -\left[\alpha_v \sum_{n\in\mathcal{V}} \max\bigl(0, |V_{t+1}^n - V_n^*| - \epsilon_v\bigr) + \alpha_\zeta \sum_{k\in\mathcal{K}} \max\bigl(0, \zeta_{\min} - \zeta_k\bigr)\right]. \tag{3.30}$$

This reward construction is consistent with the Chapter 2 analysis: voltage deviation is treated as an operational security indicator, while the damping-ratio penalty directly reflects the small-signal electromechanical objective of maintaining adequate modal damping in the critical mode set $\mathcal{K}$.

The reinforcement-learning agent seeks a policy that maximizes the expected discounted cumulative reward,

$$\pi^* = \arg\max_\pi \mathbb{E}_\pi\left[\sum_{\tau=0}^{\infty} \gamma^\tau r_\tau\right]. \tag{3.31}$$

The associated value function is defined as

$$V^\pi(s_t) = \mathbb{E}_\pi\left[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_\tau \,\middle|\, s_t\right], \tag{3.32}$$

and satisfies the Bellman recursion

$$V^\pi(s_t) = \mathbb{E}_\pi[r_t + \gamma V^\pi(s_{t+1}) \,|\, s_t]. \tag{3.33}$$

This representation makes explicit that the value of the current state depends not only on the immediate effect of the present action, but also on its discounted long-term influence on future stability and operating cost. In the proposed framework, this value function is approximated by the learned critic.

3.6 Actor-Critic Policy Parameterization

Having defined the grid-control problem as an MDP, the next component of the framework is the reinforcement-learning control policy that performs sequential decision-making under evolving grid conditions. While the spectral allocator provides a topology-aware inertia reference and the ST-GNN provides predictive state information, neither component alone determines the coordinated control actions required for sequential operation. The policy is therefore represented as

$$\pi(a_t \mid s_t),$$

which maps the current graph-structured system state to a continuous control action.

The policy receives the current state s_t, together with the learned spatio-temporal representations generated by the ST-GNN and the spectral reference inertia profile $H_{\mathrm{ref}}(t)$, and produces a distribution over the continuous action vector defined in (3.26). To support continuous adjustment of redispatch, demand response, and virtual inertia, the controller is implemented using an actor–critic architecture. The actor proposes the control action, while the critic estimates the corresponding long-term value.

Both the actor and critic are implemented as multilayer perceptrons. For continuous control, the actor outputs the mean and standard deviation of a Gaussian action distribution, allowing exploration during training. During evaluation, the mean action is applied directly to ensure deterministic operation. The critic approximates the value function in (3.32), thereby supplying the state-value estimates required for policy improvement.

The policy is optimized using the clipped Proximal Policy Optimization (PPO) objective with clipping parameter $\epsilon_{\mathrm{clip}} = 0.2$. The clipping mechanism restricts large policy updates between successive training iterations and thereby improves training stability. The raw policy output a_t specifies proposed redispatch, load-control, and virtual-inertia commands. However, these proposed actions are not executed directly; instead, they are passed to the supervisory coordination mechanism, where they are reconciled with the spectral reference inertia profile and projected onto the feasible action set.

3.6.1 Spectral–RL Coordination and Stability Analysis

The role of the coordinator is to couple the adaptive flexibility of reinforcement learning with the topology-aware stability guidance provided by the spectral allocator. The RL policy is responsible for sequential correction and adaptation under uncertainty, but it is not permitted to override the physics-guided stability logic arbitrarily. Instead, the inertia-related component of the RL action is interpreted as a residual adjustment around the reference allocation $H_{\mathrm{ref}}(t)$ derived from the Chapter 2 sensitivity analysis.

As illustrated in Fig. 3.7, the executed inertia command is computed through a blending-and-projection mechanism. The coordinator first blends the RL proposal with the physics-guided spectral reference and then projects the result onto the feasible inertia set. This guarantees that the executed inertia command remains both physically admissible and stability-oriented.

Formally, the executed inertia command is

$$H_{t+1} = \Pi_{\mathcal{H}}\Big[(1-\eta)H_{\mathrm{ref}}(t) + \eta\big(H_t + \Delta H_t\big)\Big], \tag{3.34}$$

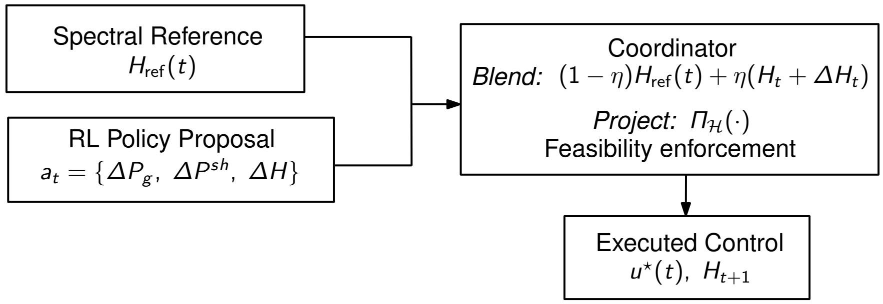

Figure 3.7: Blending-and-projection coordination mechanism used to reconcile the RL inertia proposal with the spectral reference inertia profile before execution.

where $\eta \in [0, 1]$ is the blending coefficient and $\Pi_{\mathcal{H}}[\cdot]$ denotes projection onto the feasible inertia set defined by (3.24). The parameter η governs the relative authority of the two signals: small values keep the executed command close to the spectral reference, whereas larger values allow stronger policy-driven adaptation.

The projection operator is defined by

$$\Pi_{\mathcal{H}}[z] = \arg\min_{H \in \mathcal{H}} \|H - z\|_2^2, \tag{3.35}$$

which returns the closest feasible inertia allocation to a candidate vector z in the Euclidean norm. Since the feasible set is determined by lower and upper bounds together with a total inertia budget, this projection can be computed efficiently while ensuring realizability of the final command. In physical terms, the coordinator ensures that the commanded virtual inertia remains compatible with the device-level and system-level limits under which inverter-based fast response can be delivered over the control interval.

The stability implications of this coordination mechanism can be analyzed using the Lyapunov energy function

$$V(\theta, \dot{\theta}, t) = \frac{1}{2}\dot{\theta}^\top M(t)\dot{\theta} + \frac{1}{2}\theta^\top L(t)\theta, \tag{3.36}$$

where $L(t)$ is the weighted network Laplacian capturing the synchronizing coupling among buses. Between switching events, the derivative of V satisfies

$$\dot{V} = -\dot{\theta}^\top D\dot{\theta} \leq 0,$$

So the energy is non-increasing whenever the damping matrix is positive definite.

Because both inertia allocations and network topology evolve in a piecewise-constant manner, the closed-loop system is naturally interpreted as a switched dynamical system. The purpose of the coordinator is therefore not only to enforce feasibility, but also to preserve the boundedness assumptions under which the Chapter 2 stability analysis remains valid. Under these conditions, the same Lyapunov function can be used across all admissible topologies and inertia updates, yielding the following result.

Theorem 3.1 (Uniform Asymptotic Stability). *Consider the grounded swing dynamics*

$$\dot{x}(t) = A(M(t), t)x(t)$$

with $u(t) \equiv 0$*, where the grounded Laplacian* $L(t)$ *is positive definite,* $D = \mathrm{diag}(D_i) \succ 0$ *is fixed, and each inertia satisfies*

$$0 < m_{\mathrm{lo}} \leq M_i(t) \leq m_{\mathrm{hi}} < \infty$$

for all t*. Suppose* $L(t)$ *and* $M(t)$ *vary in a piecewise-constant manner. Then the origin is a uniformly asymptotically stable equilibrium for any admissible evolution of* $L(t)$ *and* $M(t)$*. In particular, stability is preserved under any sequence of blended-and-projected updates* (3.34) *satisfying the feasibility constraints of* (3.24).

Proof (summary). Define the energy function

$$V(\theta, \dot{\theta}, t) = \frac{1}{2}\dot{\theta}^\top M(t)\dot{\theta} + \frac{1}{2}\theta^\top L(t)\theta.$$

Uniform bounds on $M(t)$ together with positive definiteness of the grounded Laplacian imply that V is uniformly positive definite and radially unbounded with respect to the grounded state. Between switching instants,

$$\dot{V} = -\dot{\theta}^\top D\dot{\theta} \leq 0.$$

At switching instants, the value of V may change, but the same quadratic bounds continue to hold because $M(t)$ and $L(t)$ remain in a uniformly bounded admissible family. The only invariant set for which $\dot{V} = 0$ is

$$\{\dot{\theta} = 0,\ L(t)\theta = 0\} = \{(0, 0)\},$$

since the grounded Laplacian is positive definite for every admissible topology. Uniform asymptotic stability then follows from the LaSalle–Yoshizawa invariance argument for piecewise-constant switched systems. □

Corollary 3.1 (Stability under Blended Allocations)**.** *Under the assumptions of Theorem 3.1, let H_{t+1} be generated by (3.34) for any $\eta \in [0, 1]$ and projected onto the feasible set $\mathcal{H}$. Then the corresponding inertia matrix satisfies the uniform bounds required by Theorem 3.1. Consequently, the grounded swing dynamics remain uniformly asymptotically stable under any admissible sequence of blended-and-projected inertia updates.*

It is important to note that the stability guarantee of Theorem 3.1 is conditional. It holds provided the network graph remains connected between switching events, the inertia

Algorithm 3.2 Closed-Loop Coordination of RL Actions with Spectral Inertia Guidance

Input: Time-varying grid $\mathcal{G}(t)$, feasible sets $\mathcal{U}, \mathcal{H}$
Initialize: node features x_t, inertia state $\mathbf{H}_t$
for each control interval t **do**
    Observe state $s_t = \{\mathcal{G}(t), x_t\}$
    Compute embeddings and predict $\hat{V}(t+1)$, $\hat{\delta}(t+1)$, and $\hat{H}(t+1)$
    Compute spectral reference $\mathbf{H}_{\text{ref}}(t)$ via (3.21)–(3.24)
    Compute control action $a_t = \pi(s_t)$ using (3.26)
    Extract inertia adjustment $\Delta\mathbf{H}_t$ from a_t
    Compute $\mathbf{H}_{t+1} = \Pi_{\mathcal{H}}[(1-\eta)\mathbf{H}_{\text{ref}}(t) + \eta(\mathbf{H}_t + \Delta\mathbf{H}_t)]$
    Apply control $(P_g, Q_g, P^{sh}, Q^{sh}, \mathbf{H}_{t+1})$ to the grid
    Observe s_{t+1} and reward r_t
    Store (s_t, a_t, r_t, s_{t+1})
end for

values stay within the bounds $[m_{\text{lo}}, m_{\text{hi}}]$ enforced by the feasibility projection, and the disturbance input satisfies $u(t) \equiv 0$ in the grounded swing model. Under severe multi-element contingencies that disconnect the network, or under large exogenous disturbances that push the system far from the linearization point, these conditions may not hold and the guarantee does not extend. The stability claim should therefore be understood as a conditional formal guarantee within the stated operating envelope, not an unconditional property of the closed-loop system.

In practice, the spectral-sensitivity allocation $H_{\text{ref}}(t)$ is computed both during training and during closed-loop operation. During training, it provides supervision for the inertia prediction head of the ST-GNN through the loss term in (3.13), while during deployment it acts as a stabilizing reference that regularizes the RL policy's inertia adjustments.

The complete interaction among the spectral allocator, the ST-GNN predictor, and the RL controller at each decision interval is summarized in Algorithm 3.2.

Algorithm 3.2 highlights how the spectral allocator, the ST-GNN predictor, and the RL policy interact within the closed-loop supervisory framework. The spectral module provides a stability-oriented reference, the ST-GNN supplies short-horizon state information, and the RL controller proposes coordinated control actions. The coordinator then enforces feasibility and preserves alignment with the physics-guided stability objective before execution in the grid.

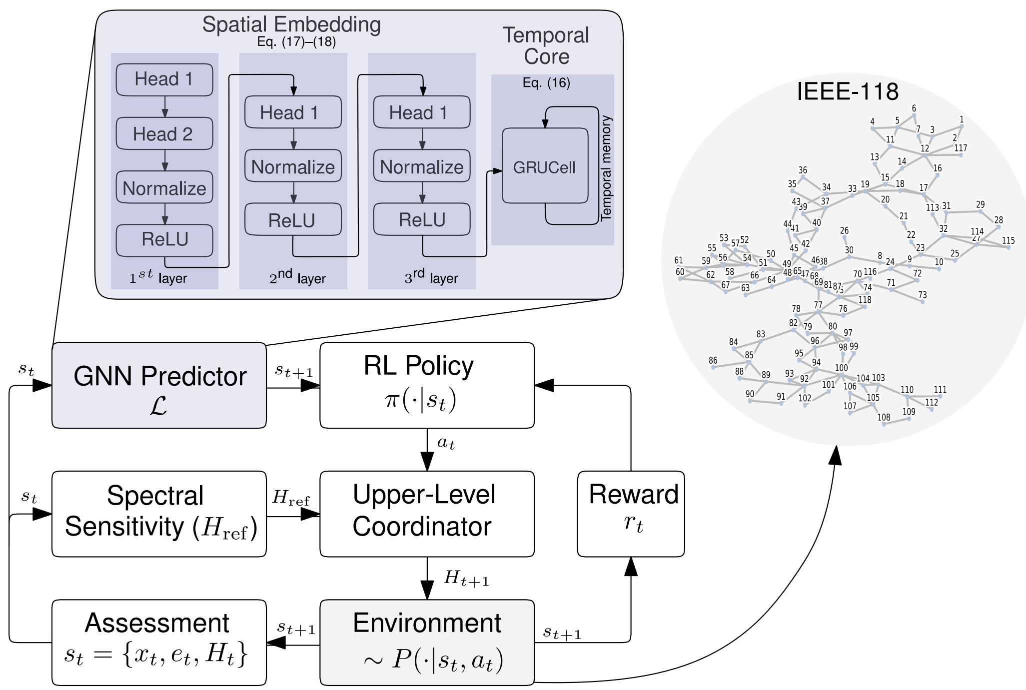

Figure 3.8: Integrated closed-loop framework combining spectral-sensitivity-based virtual inertia allocation, spatio-temporal graph neural network prediction, reinforcement-learning-based control, and supervisory coordination for stability-aware grid operation.

Figure 3.8 summarizes the complete closed-loop interaction among the ST-GNN predictor, the spectral-sensitivity allocator, the reinforcement-learning controller, the upper-level coordinator, and the grid environment. At each control interval, the environment provides the current graph-structured state s_t, which contains the network topology and bus-level operating variables. Based on this state, the ST-GNN produces short-horizon predictive features by processing the graph through its spatial embedding layers and temporal core, while the spectral-sensitivity module computes the physics-guided reference inertia profile H_{ref} from the current electromechanical condition. These two information streams are then supplied to the RL policy, which proposes the control action a_t. Before execution, the upper-level coordinator reconciles the policy output with the spectral reference and projects the resulting command onto the feasible set, yielding the executed inertia and control actions applied to the environment. The environment then evolves to the next state s_{t+1}, and the resulting operational performance generates the reward signal r_t, which closes the learning loop. In this way, the figure highlights that the proposed framework is not a sequence of independent modules, but a supervisory architecture in which physics-based guidance, graph-based prediction, and learning-based control interact continuously to support adaptive and stability-aware grid operation.

Chapter 4

Numerical Evaluation and Case Study

The integrated framework developed in Chapter 3 combines a physics-informed spatio-temporal graph neural network for state prediction, a spectral-sensitivity module for topology-aware virtual inertia allocation, and a reinforcement learning controller for coordinated closed-loop grid operation. While Chapter 3 establishes the theoretical foundations and algorithmic design of these components, their practical value depends on whether the learned representations remain accurate under realistic operating variability, whether the inertia allocation reproduces the physics-guided reference under contingency conditions, and whether the closed-loop controller maintains system feasibility and stability over extended operating horizons. These questions motivate the empirical evaluation presented in this chapter.

The assessment is conducted on the IEEE 118-bus transmission benchmark, a large meshed standard test system in which the buses are interconnected through multiple parallel transmission paths rather than a single radial chain. As a result, disturbances at one part of the network propagate through several routes, and system stability depends on the set of transmission lines in service at any given time. Experiments are designed to reflect the conditions under which the framework is intended to operate, including time-varying load demand, stochastic renewable generation, and N–2 transmission-line contingencies that modify the network Laplacian and, in doing so, alter the damping of the critical electromechanical modes identified in Chapter 2.

The evaluation proceeds in three stages. First, the dataset generation procedure, simulation environment, and training configuration are described. Second, the predictive accuracy of the ST-GNN is validated against the AC optimal power flow reference, which enforces the underlying network physics and therefore provides a physically consistent target for the learned model. This stage includes a comparative study of GNN backbones, a runtime characterization, and an explicit verification of the predictor under both single-line (N–1) and double-line (N–2) outages. Third, the closed-loop performance is assessed under a representative N–2 disturbance on the bus-68 corridor over a 24-hour horizon, examining in turn the structural interpretability of the trained model, the spatio-temporal evolution of the system state, the virtual inertia allocation produced by the controller, the demand-response behaviour, and the aggregate reinforcement-learning control metrics.

4.1 Environment and Experimental Setup

Experiments are conducted on the IEEE 118-bus benchmark network, comprising 118 buses, 54 synchronous generators, and 186 transmission lines. The network mixes single-path branches that feed peripheral loads, corridors with several parallel paths between regions, and heavily loaded lines that operate close to their limits, which makes it wellsuited for evaluating the topology-aware representations and spectral-sensitivity-driven inertia allocation developed in Chapter 3. The network also exhibits inter-area oscillations, in which clusters of generators in different regions of the grid swing collectively against each other after a disturbance. The damping of these oscillations depends on both the inertia at each bus and the set of transmission lines currently in service, as established in the linearized swing model of Chapter 2.

To generate a dataset that reflects the variability of realistic grid operation, 1,000 day-long scenario trajectories are produced by Monte Carlo sampling over the demand profiles, the renewable generation profiles, and the line-outage events that define each scenario. Each trajectory spans a 24-hour horizon at a five-minute resolution, yielding 288 time steps per trajectory and 288,000 system snapshots in total. The demand at each load bus is built in two steps. The first step takes the base IEEE 118-bus active and reactive demand at that bus and multiplies it by a daily load pattern that depends on the type of load served by the bus, with separate residential, commercial, and industrial patterns assigned to the corresponding buses. The second step adds a small random variation at each five-minute step, so that two trajectories that share the same daily pattern still differ in their minute-to-minute fluctuations.

Renewable generation variability is introduced through time-varying generation profiles that follow a daily solar or wind shape and that are themselves randomized across trajectories, so that the renewable contribution rises and falls through the day in a way that is recognizable but not identical from one trajectory to the next. The power factor at the generation buses is kept within practical operating limits throughout the 24-hour horizon, thereby keeping each unit's reactive output within the range against which real generators are dispatched in practice.

To model structural uncertainty, transmission-line contingencies are introduced into each trajectory by randomly removing a subset of lines, with higher selection probabilities assigned to lines that carry greater active power in the standard IEEE 118-bus reference solution. Any sampled contingency that would disconnect the network is discarded and resampled to ensure that the post-fault topology remains connected. A single fault onset time is placed uniformly over the 24-hour window, so that the network operates on one fixed topology before the fault and on a different fixed topology afterwards. Each topology directly modifies the susceptance-weighted graph Laplacian governing the swing dynamics in (2.22). Consequently, the modal properties of the state matrix $A(M, D, K(t))$ used for stability assessment are reshaped at the contingency time.

For each realized operating condition and topology, an AC optimal power flow problem is solved to obtain a feasible steady-state operating point. The resulting voltage magnitudes, voltage angles, and power injections serve as ground-truth supervision targets for the spatio-temporal predictor described in Section 3.3. In parallel, the spectral-sensitivity module introduced in Section 3.4 is applied to each realized topology to compute a reference virtual inertia allocation by evaluating the eigenvalue sensitivities in (3.21) and solving the budget-constrained allocation problem in (3.24). The resulting inertia values are incorporated into the node feature vector defined in (3.3), enabling the learning framework to capture stability-relevant structural information alongside electrical operating-point quantities.

The resulting dataset comprises node-level features defined as in (3.3), including the forecasted active and reactive power demands over the prediction horizon, the active and reactive generation injections, the bus voltage magnitude and phase angle, and the effective inertia constant applied at each bus, together with edge-level features that record the loading of each transmission line in service at each time step. The complete dataset is partitioned into training, validation, and test sets at the trajectory level. As a result, no time steps from the same operating scenario appear in more than one partition, and the model is never validated or tested on snapshots that share their underlying trajectory with any training sample. The complete dataset, the data-generation scripts, and the trained models are publicly available [100].

The ST-GNN is pre-trained in a supervised, one-step-ahead setting on the generated dataset, with the ground-truth state at each time step used as input to the next prediction, so that errors do not accumulate across the prediction window. Given a temporal window of recent system states, the model learns to predict the next-step voltage magnitudes, voltage angles, and virtual inertia allocations under evolving operating conditions, with training guided by the composite physics-informed loss in (3.13). The relative weights of the supervised and physics-informed terms in this loss are tuned on the validation set to balance state-prediction fidelity, physics-constraint satisfaction, and inertia-tracking accuracy. The training scripts and the resulting model checkpoints are provided in the accompanying repository [100].

The blending parameter that combines the RL-proposed inertia update with the spectral-sensitivity reference is set to $\eta = 0.5$, giving equal weight to the two signals as specified in (3.34).

Following pre-training, the reinforcement learning policy is trained through interaction with the simulated grid environment. At each decision interval, the ST-GNN provides the predicted next-step system state and the aggregated latent representation, which the policy network uses to generate coordinated control actions, including generator redispatch, demand response, and virtual inertia adjustments. The proposed action is then passed through the blending-and-projection coordination mechanism of Section 3.6.1. This mechanism reconciles the RL-proposed inertia adjustment with the spectral-sensitivity reference and projects the result onto the feasible action set defined by the generator limits, the demand-response bounds, and the inertia budget. The resulting feasible action is then applied to the environment. The environment advances to the next interval and returns a reward signal computed from the decomposed objective in (3.28), which penalizes load shedding, generation cost, and voltage-limit violations. The sections that follow assess the predictive accuracy of the trained ST-GNN and the policy's closed-loop control performance under realistic temporal variability and topology changes.

4.2 Model Evaluation and Validation

This section characterizes the predictive and computational performance of the framework components prior to the closed-loop study in Section 4.4. The metrics used to assess prediction fidelity, inertia allocation quality, and closed-loop control effectiveness are defined first. The choice of GNN backbone is then validated through an architectural comparison among the GCN, GAT, and GIN, and the per-step runtime of the proposed pipeline is reported to confirm that it fits inside the five-minute control interval. The section closes with an explicit verification of the predictor against the OPF reference under both single-line (N–1) and double-line (N–2) contingencies.

4.2.1 Evaluation Metrics

The predictive accuracy of the ST-GNN is assessed using the mean absolute error (MAE) and root-mean-square error (RMSE), computed for voltage magnitudes and angles across all buses and time steps. These metrics quantify the fidelity of the one-step-ahead predictions relative to the OPF-derived ground truth, providing a direct measure of how reliably the learned model tracks system evolution under both normal and post-contingency operating conditions.

The quality of virtual inertia allocation is evaluated by measuring the agreement between the GNN-predicted inertia $\hat{H}_i$ and the spectral-sensitivity-based reference H_i^{ref} computed by the allocator in (3.24). High agreement indicates that the learned allocation preserves the spatial structure dictated by the electromechanical sensitivity analysis of Chapter 2, confirming that the inertia prediction head captures topology-dependent stability priorities rather than learning a spatially uniform or topology-agnostic mapping.

Closed-loop control performance is assessed using system-level indicators that reflect both economic and security objectives, including cumulative reward, total load shedding as a fraction of system demand, aggregate generation cost, and the number of bus-voltage-limit violations over the operating horizon. These indicators correspond directly to the terms of the decomposed reward signal in (3.28) and together measure how effectively the RL policy balances feasibility, economy, and stability under contingency conditions.

4.2.2 Predictive Model Validation

To identify the most effective graph neural network backbone for the spatio-temporal prediction task, three architectures are compared: the Graph Convolutional Network (GCN), the Graph Attention Network (GAT), and the Graph Isomorphism Network (GIN). All three are trained with the same supervision targets, data splits, normalization settings, and loss weighting, so that the comparison isolates the effect of the architectural choice rather than any variation in the training procedure. For each backbone, two metrics are used. The first is the physics-informed loss in (3.13), evaluated on the validation set, which serves as a single summary of the training objective. The second is the coefficient of determination, R^2, defined as the fraction of variance in the target variables explained by

Table 4.1: Performance comparison of GNN architectures for OPF state prediction under faulted conditions.

GNN Method	$\mathcal{L}$	R^2
GCN	0.438	0.985
GAT	0.775	0.968
GIN	0.895	0.961

the model. This is computed on the state variables (V, δ, P_d, Q_d) after the training-time normalization has been reversed, so that the reported scores reflect prediction quality in physical units.

Table 4.1 shows that the GCN achieves the lowest validation loss and the highest R^2 score of the three architectures. The loss column reports the value of $\mathcal{L}$ in (3.13) averaged over the validation set, where lower values indicate a closer match to the training objective. The R^2 column reports how closely the predicted state variables track the OPF reference after the training-time normalization has been reversed, where values closer to one indicate a closer match in physical units. The GCN reaches $\mathcal{L} = 0.438$ and $R^2 = 0.985$, against $\mathcal{L} = 0.775$ and $R^2 = 0.968$ for the GAT and $\mathcal{L} = 0.895$ and $R^2 = 0.961$ for the GIN, so the gap between the GCN and the next-best architecture is visible on both columns.

The GCN updates each bus by combining the values at its neighbouring buses according to the network's connectivity. This operation has the same form as the graph Laplacian, which governs how disturbances propagate through the swing dynamics in (2.22), so the GCN matches the problem's physics from the start. The GAT and the GIN are designed for cases where the right way to combine neighbour information is not known in advance, so they learn it during training. In a power system where connectivity is fixed by the network topology, this learning step is unnecessary, and the GCN performs better. The GCN is therefore used in all subsequent experiments. The next subsection reports the per-step runtime of the pipeline, covering the spectral-sensitivity module, ST-GNN inference, RL policy, and coordinator, and confirms that it fits inside the five-minute control interval.

4.2.3 Computational Efficiency

A practical requirement of any online grid control framework is that its per-step inference time remains well within the five-minute decision interval at which the controller operates. To assess whether this requirement is satisfied, the wall-clock runtime of the proposed online pipeline is measured and compared against two optimization-based references: a single-step AC optimal power flow oracle solved independently at each time step, and a consecutive OPF roll-out in which the solution at time t is used to warm-start the solve at $t+1$ over the full operating horizon.

Table 4.2 shows that the proposed pipeline runs in 24.5 ms at the median and 54 ms at the 95th percentile. The single-step OPF and the sequential OPF take 250 ms and 350 ms at the median, so the pipeline is approximately ten times faster on the same hardware. Within the pipeline, the spectral-sensitivity block at 20 ms accounts for 82% of the median runtime, while the ST-GNN inference, RL policy inference, and coordinator contribute

Table 4.2: Median and 95th-percentile per-step runtimes for the reference methods and the proposed pipeline.

Computation	**Median (ms)**	**p95 (ms)**
Reference baselines (not part of the proposed pipeline)		
OPF oracle (per step)	250	600
Consecutive OPF roll-out	350	820
Proposed pipeline components		
Spectral sensitivity (H_{ref})	20	45
ST-GNN inference	3	6
RL policy inference	1	2
Coordinator	0.5	1
Proposed pipeline (total)	24.5	54

3, 1, and 0.5 ms, respectively. The full per-step runtime sits well below the five-minute control interval at both the median and the 95th percentile.

4.2.4 Predictor Robustness Under N–1 and N–2 Contingencies

The predictor is evaluated against the OPF reference under two transmission-line contingencies of increasing severity. The contingencies are drawn from a combined criticality ranking in which each line is scored on its base-case active-power flow, which captures how much load the line carries under nominal operation, and its edge-betweenness centrality, which captures how often the line lies on shortest paths connecting distant parts of the network. The two scores are weighted equally and summed to yield a single criticality value, and the contingencies tested below correspond to the highest-ranked single line and the highest-ranked pair of lines at this criticality value.

The single-line (N–1) case removes the highest-ranked line, $(68, 76)$, which carries the maximum centrality score on the IEEE 118-bus system and the largest combined criticality score overall. The double-line (N–2) case removes the top two lines from the same ranking, $(68, 76)$ and $(48, 68)$, both of which are incident on bus 68 and which together constitute a localized corridor outage in the same region of the network. The two contingencies therefore challenge the predictor on the most stressed corridor of the IEEE 118-bus system, and the close structural relationship between them prepares the ground for the closed-loop case study of Section 4.4, in which the same bus-68 corridor is exercised under the related N–2 disturbance defined by lines $(46, 68)$ and $(68, 76)$ over a full 24-hour horizon. For each severity level, the GNN-predicted voltage magnitude, voltage angle, and virtual inertia allocation are compared with the OPF-derived ground truth across all 118 buses.

Figure 4.1 shows the state-prediction comparison. The voltage magnitude profiles in

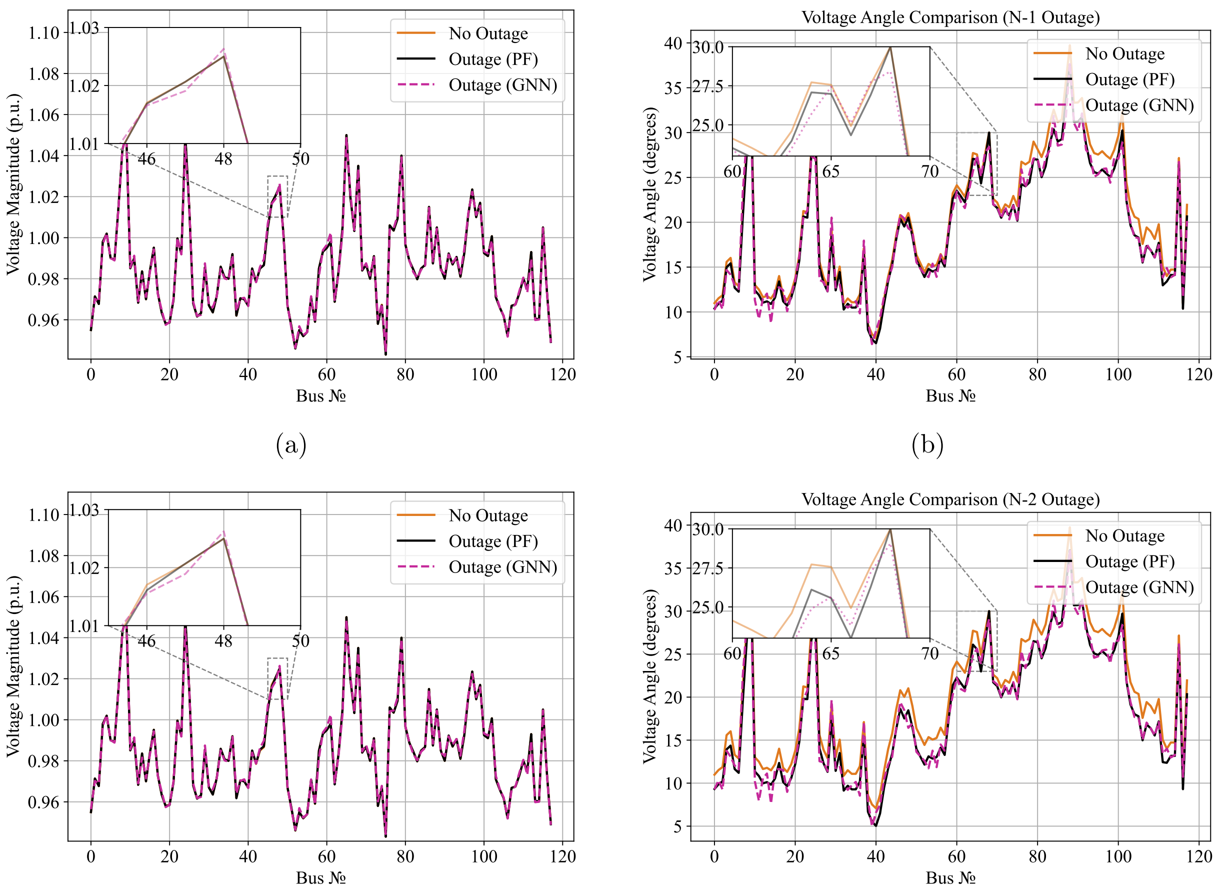

(a) (b)

(c) (d)

Figure 4.1: State-prediction verification of the GNN on the IEEE 118-bus system: (a) voltage magnitude under the N–1 outage on line $(68, 76)$; (b) voltage angle under the same N–1 outage; (c) voltage magnitude under the N–2 outage on lines $(68, 76)$ and $(48, 68)$, both incident on bus 68; (d) voltage angle under the same N–2 outage. In each panel, the no-outage baseline is shown in orange, the OPF-derived post-fault solution in black, and the GNN prediction as a magenta dashed line, with insets resolving the corridor most affected by each disturbance. The N–1 case removes the highest-ranked line based on the combined flow and centrality criticality score; the N–2 case removes the top two lines using the same ranking.

Fig. 4.1a and Fig. 4.1c show that the three traces overlap across the full set of buses, with the no-outage baseline, the OPF post-fault reference, and the GNN prediction remaining within a band that is too narrow to be resolved at the figure scale. The inset zooms in on a representative cluster around bus indices 45–50, where the three traces coincide in both shape and amplitude. The voltage angle profiles in Fig. 4.1b and Fig. 4.1d reveal a more visible response to the disturbance.

In both contingencies, the OPF post-fault trace sits systematically below the no-outage baseline across the high-angle buses in the upper half of the network, and the offset is larger under N–2 than under N–1. This downward shift is consistent with the redistribution of active-power transfer that follows a topology change. The removed lines no longer carry their share of the flow, so the surviving routes must absorb the rerouted power. To do so, they operate at smaller phase differences relative to the slack bus, and the effect compounds as more lines are removed. The GNN dashed trace tracks the OPF post-fault solution at every bus rather than the no-outage baseline, recovering both the global shift and the local angle features visible in the inset around buses 60–70. The mean absolute angle error remains below 0.13° under both contingencies, with the maximum absolute error staying under 0.60°, so the predictor's accuracy on the angle field is essentially the same in both cases. This sub-degree accuracy is comparable to the phase-angle uncertainty reported for synchrophasor measurements in practical deployments [12], and is small relative to the angle excursions that drive protection and stability margins, so the predictor's angle accuracy is operationally acceptable for stability assessment and control purposes. Table 4.3 shows that the GNN-predicted allocation reproduces the sensitivity-based reference produced by the allocator in (3.21)–(3.24) on both contingencies. The reference allocation places inertia on a compact set of stability-critical buses selected by the eigenvalue sensitivities, several of which saturate at the per-bus budget cap $H_{\max} = 2.5$; buses 70, 95, 81, 82, 83, and 85 reach this cap under both contingencies, followed by a short tail of lower allocations at buses such as 94, 2, 16, and 4. The predicted values match the reference to within a few thousandths of a per-unit inertia constant across the saturated set and the tail, giving a virtual-inertia mean absolute error of approximately 0.004 over all 118 buses under both contingencies. The agreement is essentially identical in the two cases, reflecting their close structural relationship: both share line $(68, 76)$, both concentrate stress on bus 68, and both leave the network's dominant electromechanical modes largely unchanged. As a result, the eigenvalue sensitivities in (3.21) return the same critical buses as the largest contributors to modal damping under either contingency.

Together, the state-prediction comparison in Fig. 4.1 and the inertia comparison in Table 4.3 establish that the predictor reproduces both the post-fault electrical state and the sparse spatial inertia pattern targeted by the allocator. The errors observed on the bus-68 corridor are at the same level as those measured on a separate hold-out set drawn from outages not used in training, indicating that predictive fidelity is preserved when the predictor is evaluated on the most stressed corridor of the network rather than on the average sample. Section 4.3 next examines whether this predictive accuracy reflects the same structural dependencies exploited by the spectral allocator.

Section 4.4 then extends the analysis to the full closed-loop framework on the same bus-68 corridor, under a different N–2 disturbance defined by lines $(46, 68)$ and $(68, 76)$ over

Table 4.3: Top-10 buses by reference virtual inertia, under the N–1 outage on line $(68, 76)$ and the N–2 outage on lines $(68, 76)$ and $(48, 68)$. Values are in per-unit inertia constants (H), capped at $H_{\max} = 2.5$ by the allocation budget.

Bus	**N–1**		**N–2**	
	$\mathbf{H}^{\text{ref}}$	$\hat{\mathbf{H}}$	$\mathbf{H}^{\text{ref}}$	$\hat{\mathbf{H}}$
70	2.500	2.486	2.500	2.486
95	2.500	2.511	2.500	2.510
81	2.500	2.508	2.500	2.508
82	2.500	2.510	2.500	2.511
83	2.500	2.510	2.500	2.512
85	2.500	2.503	2.500	2.502
94	2.100	2.098	1.950	1.948
2	1.950	1.950	1.850	1.849
16	1.850	1.836	2.050	2.037
4	1.650	1.649	1.300	1.297

a 24-hour horizon. The bus-68 corridor was selected as the evaluation scenario because it carries the highest combined flow-and-centrality criticality score on the IEEE 118-bus system and therefore represents the most stressed operating condition for the framework. Generalization to other corridors and to N–x ($x > 2$) contingencies remains an open direction for future work.

4.3 Structural Awareness and Interpretability

Building on the predictive validation of Section 4.2, this section examines how the trained ST-GNN uses network structure to generate its virtual inertia predictions. The spectral-sensitivity allocator introduced in Chapter 2 provides a physics-grounded reference for inertia placement, derived from the eigenstructure of the linearized state matrix in (2.22). The trained ST-GNN, by contrast, learns the same association directly from data, without being given the spectral calculation, and the question is whether the learned model relies on the same structural features that the spectral allocator exploits. To answer this question, the gradient-based saliency framework defined in (3.20) is applied to the trained model. For each edge in the network, the saliency score measures how strongly a small change in the graph-theoretic descriptors at its two endpoint buses would alter the model's predicted bus-wise inertia. The descriptors used are the degree, closeness, and betweenness centrality of each endpoint bus. Edges with large saliency, therefore, mark corridors whose structural role most influences the learned allocation, and the spatial distribution of these values reveals which regions of the network the model treats as electromechanically critical.

Figures 4.2–4.4 show the saliency distribution for the N–2 contingency on lines $(46, 68)$ and $(68, 76)$ used as the reference scenario throughout this chapter. In the pre-fault state shown in Fig. 4.2, the saliency field is concentrated on a small number of corridors and

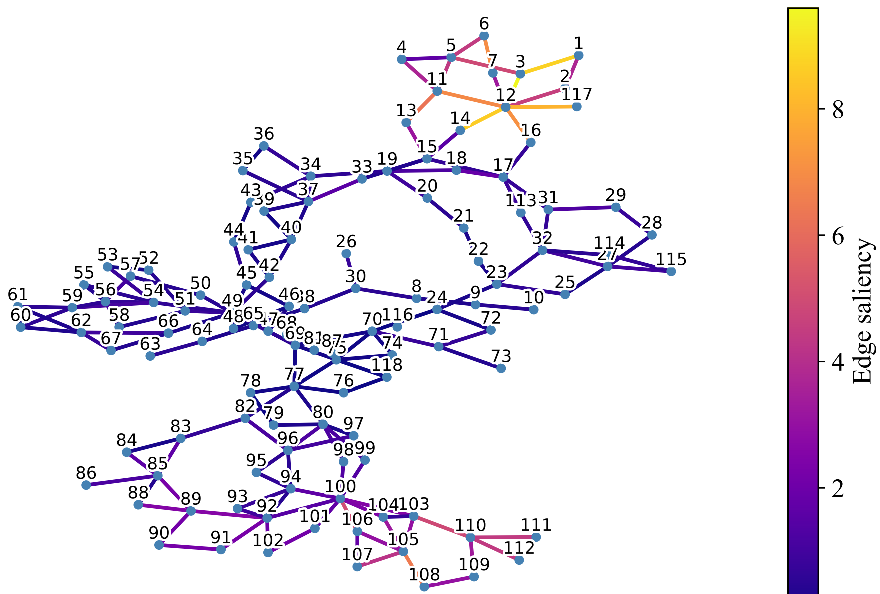

Figure 4.2: Pre-fault edge saliency for the inertia head. Higher values indicate edges whose structural role most strongly influences the predicted virtual inertia H.

remains close to zero across most of the densely meshed interior. The brightest edges sit on corridors that connect subgraphs of high betweenness and high degree to the rest of the network, with the largest cluster located around bus 12 in the upper region of the system. The reason for this pattern follows from the construction of the saliency score in (3.20), which measures the difference in structural descriptors at the two endpoints of each line. Edges that join a high-degree hub to a lower-degree neighbour, or that lie on the only short path between two otherwise distant clusters, have endpoints whose descriptors differ markedly, so their saliency is correspondingly large. Edges in well-meshed regions, where the descriptor values at neighbouring buses are similar, carry near-zero saliency.

Following the simultaneous removal of lines $(46, 68)$ and $(68, 76)$, illustrated in Fig. 4.3, the saliency distribution undergoes a structural reorganization rather than a uniform reduction.

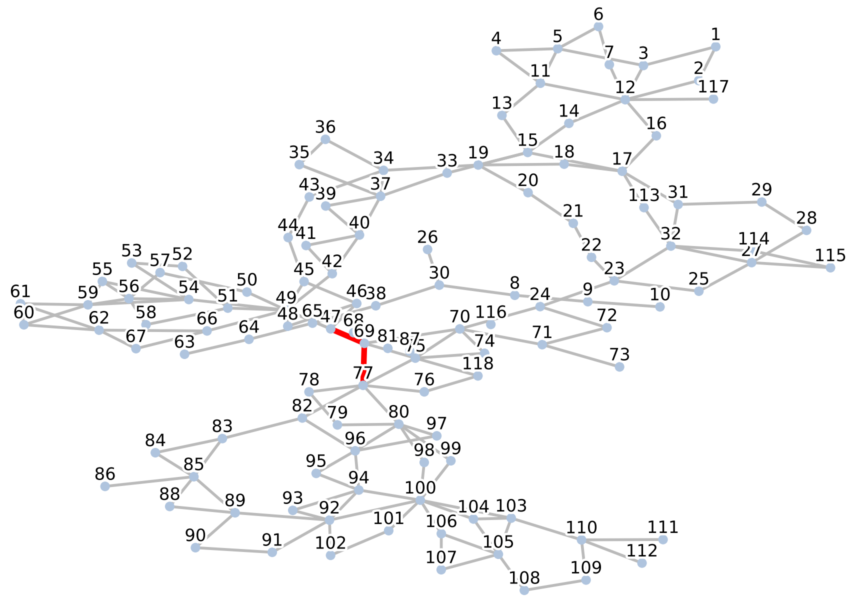

Figure 4.3: Post-fault topology of the IEEE 118-bus system. The two transmission lines $(46, 68)$ and $(68, 76)$ removed by the N–2 contingency are highlighted in red, both incident on bus 68.

The post-fault map in Fig. 4.4 shows two visible effects. First, the southern corridor connecting buses 89 through 101 exhibits markedly higher saliency than before the outage. This corridor lies on the alternative transfer paths that must absorb the flow previously routed through bus 68, and as those paths become more central in the post-fault graph, their endpoint descriptors begin to differ more sharply, raising the saliency of every edge along the corridor. Second, the cluster around bus 12 in the upper region softens because the loss of two lines incident to bus 68 does not alter the structural role of the northern corridor itself. Regions of the network that continue to provide multiple substitute routes after the outage show stable, low saliency, consistent with a diffusion of structural influence across many parallel paths.

The pre-fault and post-fault saliency maps therefore confirm that the learned model uses

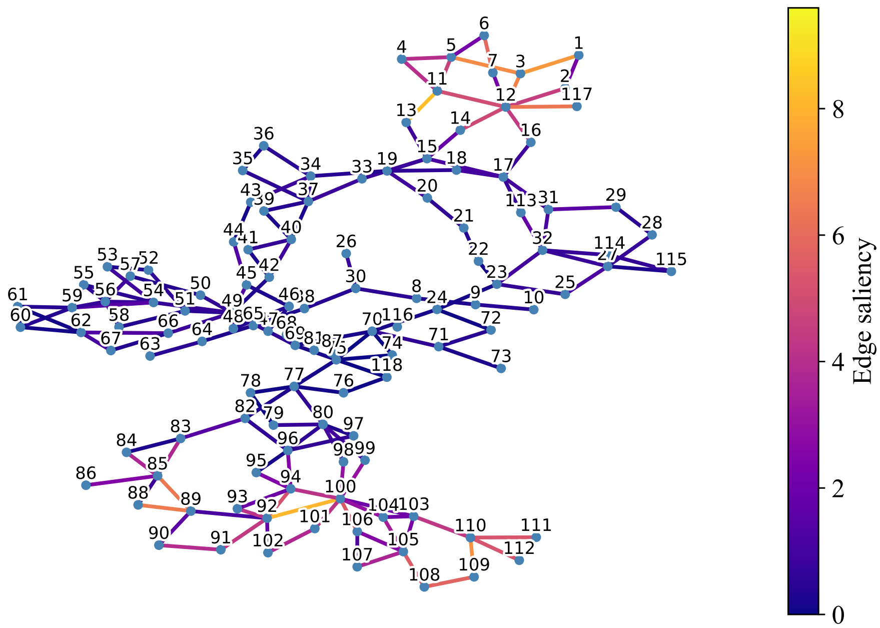

Figure 4.4: Post-fault edge saliency computed on the updated topology. The southern corridor near buses 89–101 intensifies while the upper cluster around bus 12 softens, indicating a redistribution of the model's structural sensitivity toward the alternative routing paths that absorb the rerouted flow.

network structure rather than only bus-level signals to generate its inertia predictions. The redistribution of saliency mirrors the redistribution of post-fault routing capacity that motivates the spectral-sensitivity allocator of Section 3.4, providing data-driven evidence that the predictor and the physics-based allocator agree on which regions of the network become electromechanically pivotal under topology change. This analysis is diagnostic rather than prescriptive. It does not directly modify the inertia allocation, but it identifies the corridors whose structural sensitivity shapes the closed-loop control behaviour examined in the following sections.

4.4 System Behaviour Under an N–2 Contingency

Having established the structural interpretability of the trained ST-GNN in Section 4.3, this section evaluates the predictor's ability to reproduce system states under an explicit topology disturbance. A representative N–2 contingency is considered in which transmission lines $(46, 68)$ and $(68, 76)$ are removed simultaneously. The contingency forces the network to redistribute its power flows along the surviving corridors and updates the susceptance-weighted Laplacian that governs the swing dynamics in (2.22). The simultaneous loss of two lines incident on bus 68 concentrates electromechanical stress on the surrounding corridors. Because the modal properties of the state matrix $A(M, D, K(t))$ depend on the post-fault topology itself, the post-disturbance dynamics must be re-evaluated on the updated network rather than inferred from the pre-fault operating point. Building on the static prediction validation of Section 4.2.4, the following subsections examine the spatio-temporal evolution of the system state over the full 24-hour horizon and the spatial sensitivity of the network to the topology change.

4.4.1 Spatio-Temporal State Evolution

The static profile analysis of Section 4.2.4 establishes that the ST-GNN reproduces bus-level voltage and angle distributions under the faulted topology at a single time instant. To assess whether this fidelity extends to the full temporal evolution of the system, the spatio-temporal behaviour of the predicted state variables is examined over a complete 24-hour operating horizon under the same N–2 contingency. The outage is introduced at $t = 600$ min, after which the system evolves on the post-fault topology for the remainder of the trajectory. Throughout this analysis, predicted refers to the GNN output and actual refers to the OPF-based reference obtained on the corresponding topology and load realization.

Figure 4.5 shows the predicted active power evolution at ten representative PQ buses selected by peak demand. The trajectories follow a pronounced diurnal pattern, with low values during the early morning hours, a coordinated rise toward midday peak demand, and a gradual decline through the evening. After the contingency at $t = 600$ min, a localized rise in active power is observed at buses 100 and 101, indicating that these buses absorb part of the active-power redistribution induced by the loss of the $(46, 68)$ and $(68, 76)$ corridor. Conversely, buses 51 and 80 exhibit reduced post-fault loading, consistent with the rerouting of flows away from the affected region. The model therefore captures both the underlying daily load cycle and the spatial reorganization triggered by the topology change without explicit information about which buses are downstream of the failure.

To quantify prediction fidelity directly, Fig. 4.6 shows the difference between the GNN-predicted and OPF-derived active-power trajectories at the same buses. The map is dominated by near-zero values across most of the horizon, indicating that the predictor tracks the OPF reference closely under both pre-fault and post-fault conditions. The largest residuals are localized in time and space, appearing around rapid load transitions, such as the late-morning ramp toward peak demand at bus 101, and immediately following the disturbance, where short-lived errors of order 20 MW are observed at buses 100 and 101, where the post-fault active-power redistribution is largest. These residuals correspond to the most challenging conditions for a one-step predictor, in which the input state has not yet incorporated the post-fault topology change. Outside these intervals, the error remains small and does not accumulate over the horizon. This behaviour reflects the action of the physics-informed penalty terms in (3.13), which constrain the predicted trajectories to remain close to physically feasible operating points.

The corresponding voltage angle evolution at the ten PV buses with the largest peak volt-

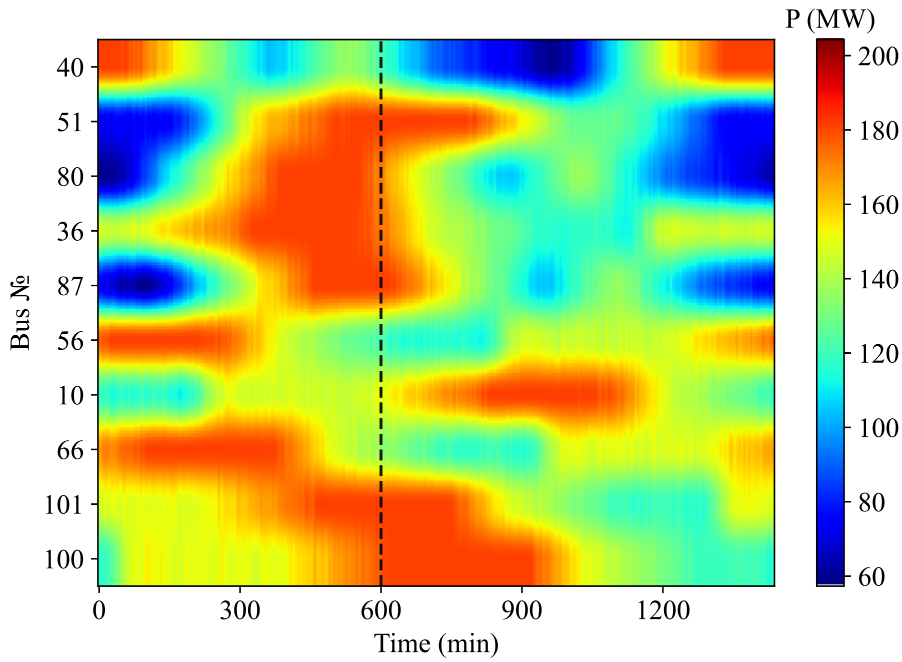

Figure 4.5: Predicted spatio-temporal evolution of active power at the ten highest-demand PQ buses (MW) over a 24-hour horizon with 5-minute resolution. The vertical dashed line marks the N–2 outage at $t = 600$ min on lines $(46, 68)$ and $(68, 76)$.

ages is shown in Fig. 4.7. Two angle groups are visible. The upper group, comprising buses such as 65, 3, 5, and 61, sits at relatively low angle values throughout the horizon, while the lower group, containing buses 76, 79, and 48, sits at substantially higher angle values. After the outage at $t = 600$ min, the lower group intensifies further, with bus 79 rising into the same band as bus 76, both of which sit near the upper end of the colour scale during the post-fault hours. This sustained increase in angle at buses adjacent to the removed corridor is consistent with the redistribution of active-power transfer toward alternative paths, which forces the surviving routes to operate at larger angle differences in order to carry the rerouted flow. The upper group experiences a complementary slight cooling, indicating that the post-fault angle field is reshaped along the network rather than uni-

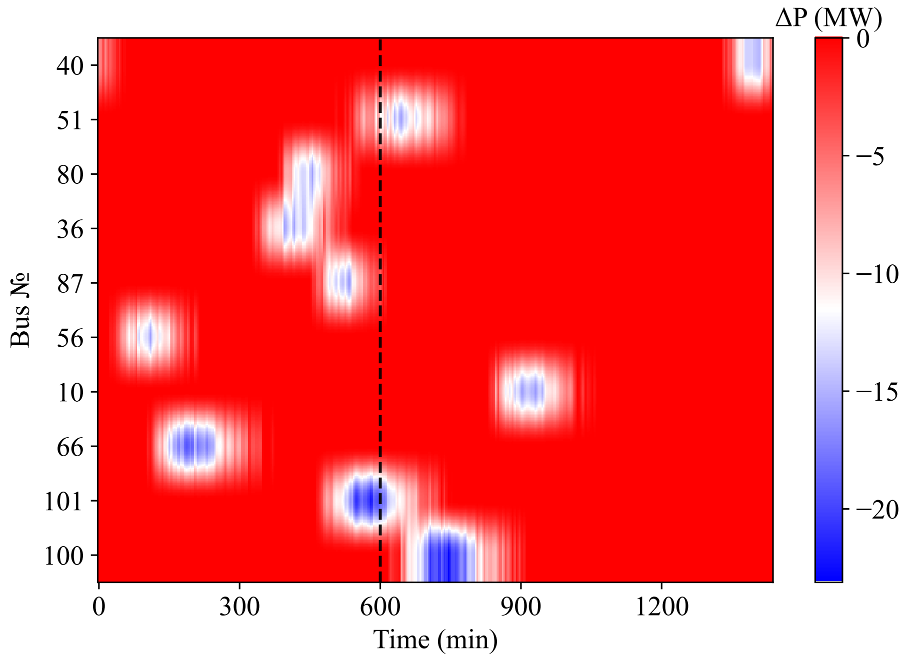

Figure 4.6: Difference map between GNN-predicted and OPF-derived active-power trajectories at ten PQ buses over a 24-hour horizon with 5-minute resolution under the N–2 contingency. The vertical dashed line marks the disturbance time. Negative values (blue) indicate buses for which the GNN underpredicts the OPF reference.

formly shifted. This spatial pattern follows from the susceptance-weighted Laplacian $K(t)$ in (2.22), which couples bus angles through active-power transfer and therefore propagates the local effect of a topology change along the surviving transmission corridors.

The voltage magnitude evolution at the same PV buses is shown in Fig. 4.8. Magnitudes remain bounded across the horizon, varying within a narrow band consistent with generator voltage regulation. Buses near the lower end of the y-axis ordering, such as 48 and 11, sit close to the upper end of the displayed range, while buses near the top, such as 65, 3, and 5, sit close to the lower end of the displayed range. After the contingency, the distribution of voltage magnitudes changes only modestly, with short-lived excursions at a small subset of buses including bus 79 between roughly 300 and 800 min, but no progressive drift develops

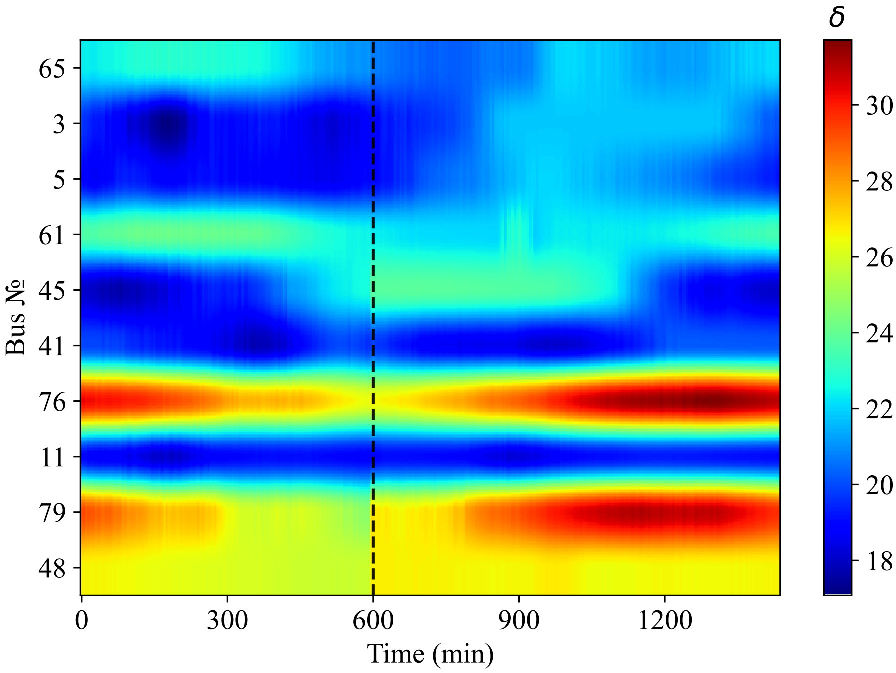

Figure 4.7: Predicted spatio-temporal evolution of voltage angle δ at ten PV buses over a 24-hour horizon with 5-minute resolution. The vertical dashed line marks the N–2 outage at $t = 600$ min.

over the post-fault horizon. The bounded response reflects the corrective re-dispatch and local reactive support applied by the controller, which restores voltage magnitudes toward their set points after the topology change. The GNN reproduces this behaviour by propagating neighbourhood information through its graph convolutional layers and aggregating temporal context via the attention-weighted encoder defined in (3.9), thereby enabling it to track coordinated voltage recovery without explicit knowledge of the applied control actions.

Across the four spatio-temporal views, the ST-GNN reproduces the system's daily operational pattern, captures the spatial reorganization of active power and voltage angles induced by the contingency, and maintains stable tracking of the OPF reference over the full 24-hour horizon. The bounded prediction errors and the absence of progressive drift

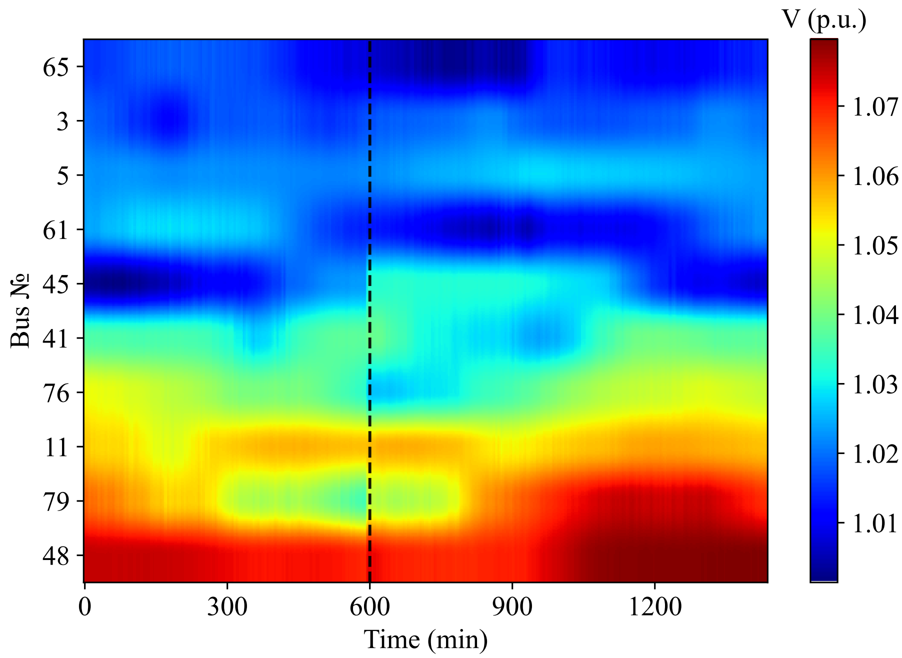

Figure 4.8: Predicted spatio-temporal evolution of voltage magnitude V (p.u.) at ten PV buses over a 24-hour horizon with 5-minute resolution. The vertical dashed line marks the N–2 outage at $t = 600$ min.

establish a reliable predictive foundation for the spatial sensitivity analysis and closed-loop control studies presented in the following sections.

4.4.2 Spatial Sensitivity to Topology Disturbances

To identify the buses most affected by the contingency, the aggregated deviations of the four bus-state and injection variables $(\Delta V, \Delta\theta, \Delta P, \Delta Q)$ are examined. For each variable, the bars in Figs. 4.9–4.12 show the sum of absolute pre- and post-contingency differences accumulated across the recorded fault realization, so that the ordering reflects how strongly each bus responds to the disturbance.

The bus-state response is shown in Figs. 4.9 and 4.10. The voltage magnitude ranking in

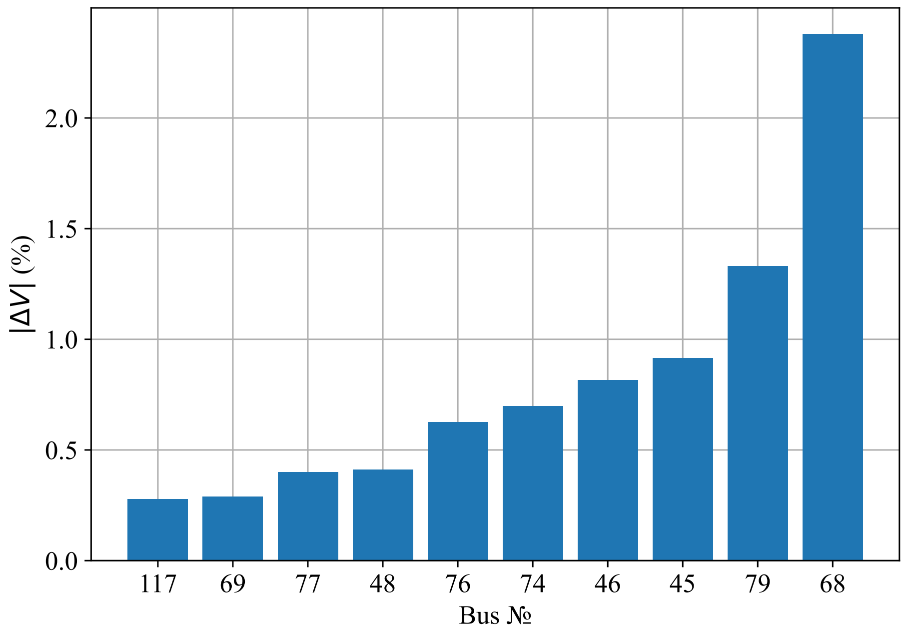

Figure 4.9: Top ten buses ranked by aggregated absolute change in voltage magnitude, $|\Delta V|$ (%). Larger totals indicate greater sensitivity to changes in topology.

Fig. 4.9 is dominated by bus 68, whose $|\Delta V|$ is markedly larger than that of any other bus in the ranking. The next largest entry is bus 79, with the remaining buses arranged in a gradual descent toward values an order of magnitude below the peak. The dominance of bus 68 reflects its position as the common endpoint of both removed lines $(46, 68)$ and $(68, 76)$. The simultaneous loss of two incident branches concentrates a voltage deviation at this bus that cannot be fully absorbed by the reactive support available at neighbouring buses. Away from bus 68, the magnitude deviations remain comparatively small, which is consistent with the action of generator voltage regulation that limits the propagation of voltage deviations beyond the immediate neighbourhood of the removed lines.

The voltage angle ranking in Fig. 4.10 shows a different shape. Bus 79 stands well above the rest of the ranking, while a tight cluster of buses comprising 46, 77, 68, and 45 takes

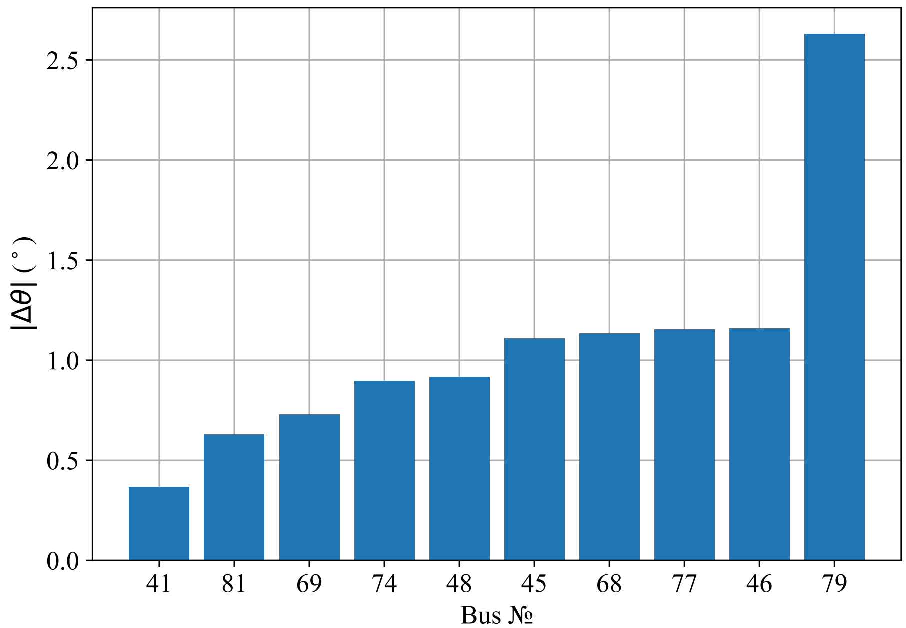

Figure 4.10: Top ten buses ranked by aggregated absolute change in voltage angle, $|\Delta\theta|$ ($^\circ$).

intermediate values within a narrow range just above 1°. The values then taper down through buses 48 and 74 to the lowest entries at buses 69, 81, and 41. The gradual rise across most of the ranking reflects the longer-range influence of θ through the active-power balance equations and the coupling between buses imposed by the susceptance-weighted Laplacian $K(t)$ in (2.22). An angle change introduced at the faulted corridor is therefore propagated across a wider set of buses than a magnitude change. The prominence of bus 79 at the top of the ranking stems from its position in the corridor that absorbs a large share of the post-fault active-power transfer. As shown in Fig. 4.7, this bus also exhibits the most pronounced post-fault increase in voltage angle in the spatio-temporal heatmap, which is consistent with its position at the head of the angle-impact ranking.

The injection response is shown in Figs. 4.11 and 4.12. The active power ranking in Fig. 4.11 places buses 46 and 68 at the top with very similar values. Bus 69 follows in third position,

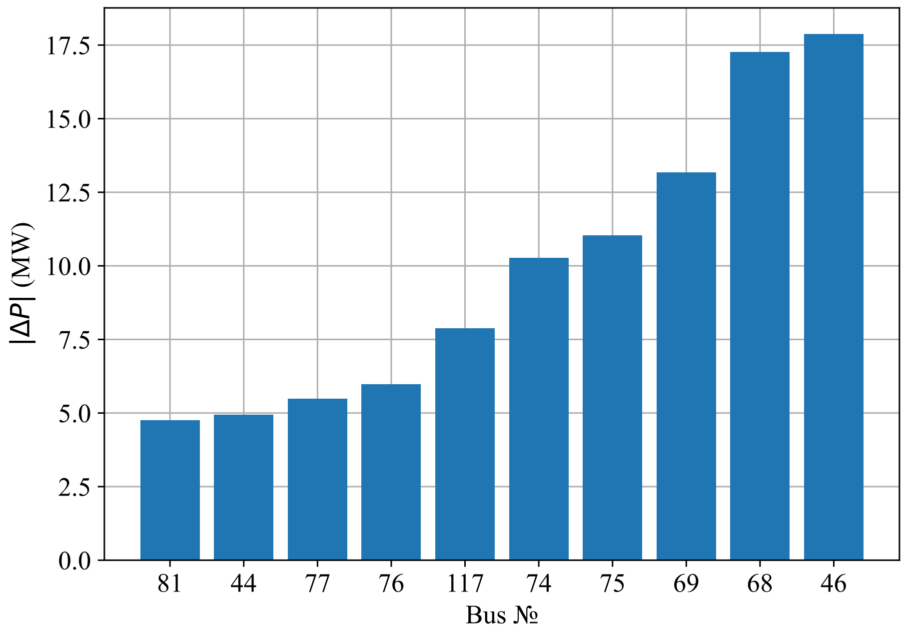

Figure 4.11: Top ten buses ranked by aggregated absolute change in active power demand, $|\Delta P|$ (MW).

with buses 75 and 74 in close support. Buses 46 and 68 are the two endpoints of the removed line $(46, 68)$, and their joint appearance at the top is a direct consequence of the redistribution of active power along the lines adjacent to the outage. Bus 69 also appears among the top entries because its incident lines carry a larger share of the rerouted flow once the $(46, 68)$ and $(68, 76)$ branches are removed. The progressive decline through the rest of the ranking reflects the diminishing share of the rerouted flow carried by buses further from the affected corridor.

The reactive power ranking in Fig. 4.12 places bus 68 at the top, but the supporting cast differs from the active-power panel. Buses 75, 69, and 76 appear in the upper ranks, while bus 73, which is not prominent in the $|\Delta P|$ ordering, also appears in the $|\Delta Q|$ ranking. These buses provide reactive support to the region around the removed lines, so they appear in the $|\Delta Q|$ ranking even when they do not carry a significant share of the redistributed

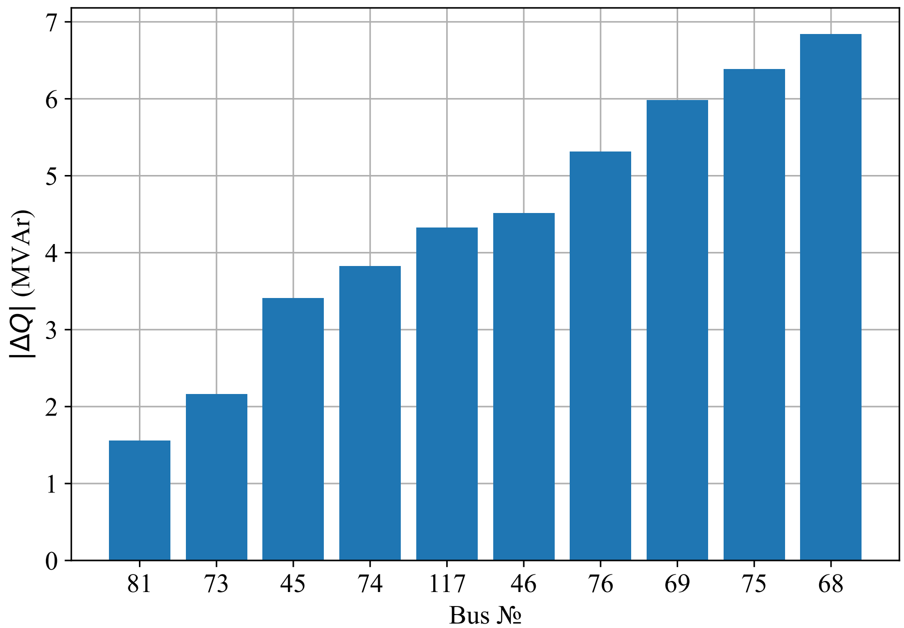

Figure 4.12: Top ten buses ranked by aggregated absolute change in reactive power demand, $|\Delta Q|$ (MVAr).

active flow. The overall magnitude of the reactive response is smaller than that of the active response, reflecting the smaller role of Q in redistributing power at the post-fault operating point.

Across the four panels, buses 46, 68, and 76, all of which are directly incident on the removed lines, appear in the upper ranks of multiple rankings. The voltage-magnitude ranking is dominated by bus 68, while the voltage-angle ranking places bus 79 at the head as the bus that carries the largest share of the rerouted active power transfer. The active-power ranking is concentrated at the line endpoints, whereas the reactive-power ranking includes buses with local reactive support. The effect of the topology disturbance is therefore concentrated on a small set of buses located on or directly adjacent to the removed lines, consistent with the eigenvalue sensitivity analysis in (3.21), which identifies these same locations as those with the strongest influence on the critical modes. As shown in the next section, these are also the buses where virtual inertia is preferentially allocated.

4.5 Stability Support and Control

The preceding section established that the ST-GNN reproduces the faulted operating point with high fidelity and that the response to the contingency is concentrated at a small set of buses located on or directly adjacent to the removed lines. The remaining question is whether the closed-loop framework acts on this information in a manner that is consistent with the small-signal stability objective derived in Chapter 2 and with the coordination mechanism introduced in Chapter 3. Three aspects of the closed-loop response are examined in turn. The virtual inertia allocation is evaluated against the spectral-sensitivity reference to verify that the learned allocation reproduces the physics-guided spatial structure under contingency conditions. The demand-response behaviour is examined to determine whether curtailment decisions remain targeted in both time and space. The aggregate closed-loop indicators are then analyzed to assess how effectively the RL policy coordinates these actions over the full operating horizon.

4.5.1 Virtual Inertia Allocation

The static verification shown in Section 4.2.4 established that the GNN-predicted allocation reproduces the spatial structure of the spectral-sensitivity reference under both single-line and double-line outages on the bus-68 corridor. The remaining question for the closed-loop case study is how that allocation evolves through the 24-hour horizon when the predictor is exercised inside the full control loop under the focal N–2 disturbance on lines $(46, 68)$ and $(68, 76)$. The inertia signal applied to the grid at each step combines the spectral-

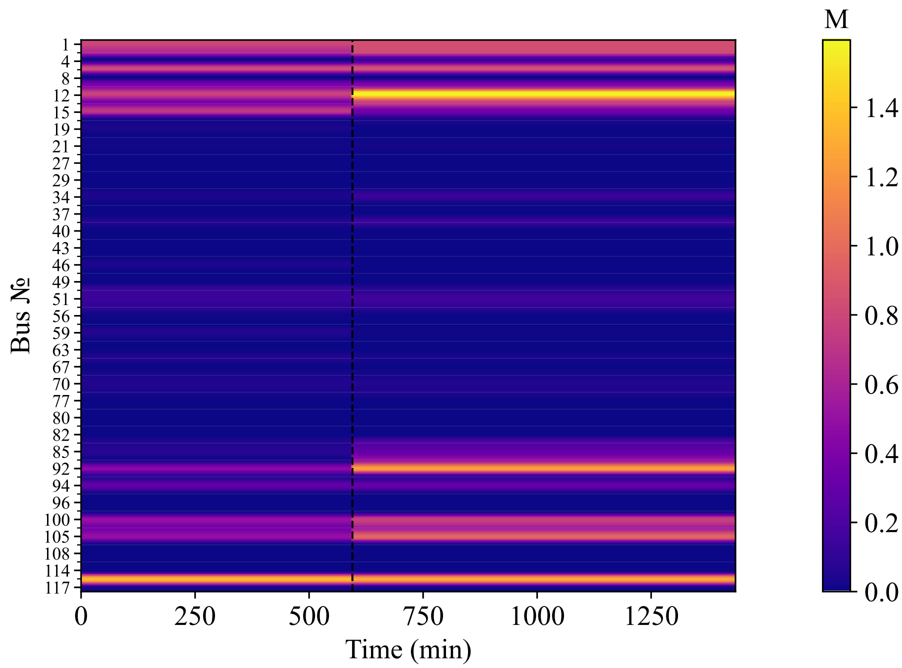

Figure 4.13: Temporal evolution of the virtual inertia allocation M across all non-generator buses over the 24-hour horizon. The vertical dashed line marks the contingency at $t = 600$ min on lines $(46, 68)$ and $(68, 76)$.

sensitivity reference of (3.24) with the RL-proposed update through the blending rule of (3.34). The resulting allocation is therefore expected to remain close to the topology-aware reference between switching events and adjust at the contingency time to track the post-fault redistribution of modal participation.

The temporal behaviour of the allocation is shown in Figs. 4.13–4.16. Figure 4.13 shows the full heatmap across all non-generator buses. The map is piecewise constant in time. Each row holds an essentially fixed value before the fault and a different fixed value afterwards, with the transition occurring sharply at $t = 600$ min. Bus 117 sits at the upper end of the scale throughout the horizon. Buses 12 and 92 are visible at elevated pre-fault values and brighten further after the fault, joined post-fault by a smaller set of additional bright rows,

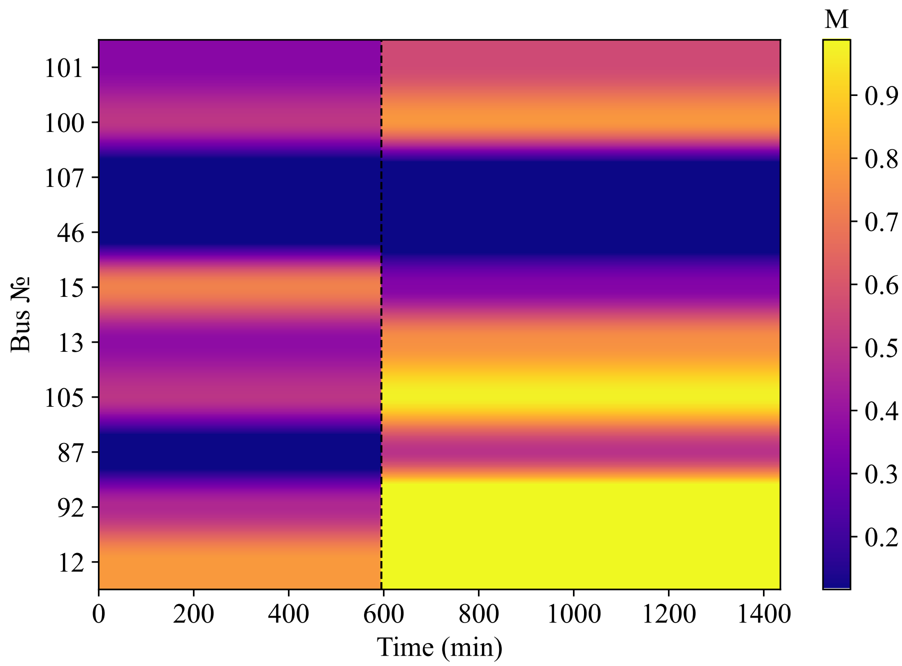

Figure 4.14: Temporal evolution of the virtual inertia allocation M at the ten most sensitive buses, on a narrower colour range that resolves the post-fault change. The vertical dashed line marks the contingency time.

including buses 100, 105, and 87. Most of the network sits at the lower end of the colour scale at all times. This piecewise-constant structure with a step at the contingency time is the expected behaviour of H_{ref} as a topology-aware reference signal, whose spatial structure updates only when network connectivity changes, rather than in response to instantaneous operating-point fluctuations driven by the daily load profile.

Figure 4.14 restricts attention to the ten most sensitive buses and uses a narrower colour range, which brings the post-fault change into sharper relief. Buses 12 and 92 both move from intermediate pre-fault values toward the upper end of the colour scale after the fault, exhibiting the largest sustained changes in the figure. Bus 87 rises from a low pre-fault value to a mid-range post-fault value, and buses 105 and 13 each strengthen by a clearly visible

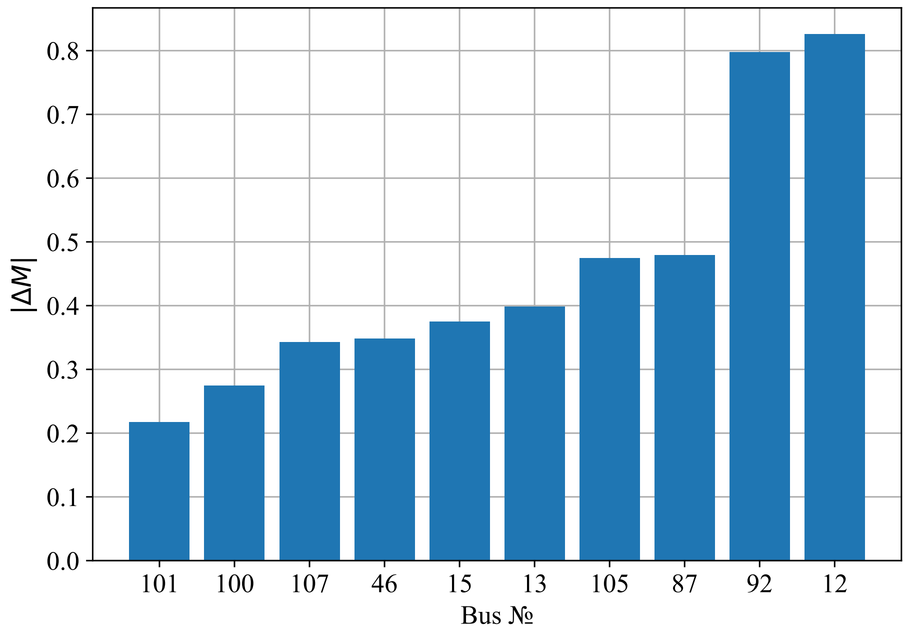

Figure 4.15: Top ten buses ranked by aggregated absolute change in virtual inertia, $|\Delta M|$. Larger values indicate buses whose allocation is most strongly modified by the contingency.

margin. Bus 15 moves in the opposite direction, dropping from a moderate pre-fault level to a dimmer post-fault level. Buses 46, 100, 101, and 107 exhibit smaller adjustments, with bus 107 showing little visible change on either side of the fault. The mixture of increases and decreases follows from the budget constraint in (3.24), which requires that any inertia added at one bus be compensated by a reduction elsewhere, so that the redistribution reflects a reallocation of a fixed resource rather than an uncompensated increase.

The magnitude of these changes is summarised in Fig. 4.15, which ranks the same ten buses by the absolute change in virtual inertia from the pre-fault baseline. Buses 12 and 92 stand at the head of the ranking with comparable values, well above the next group. A second tier comprising buses 87 and 105 takes intermediate values that are roughly half the peak. Buses 13, 15, and 46 occupy a third tier, and the remaining buses descend further toward

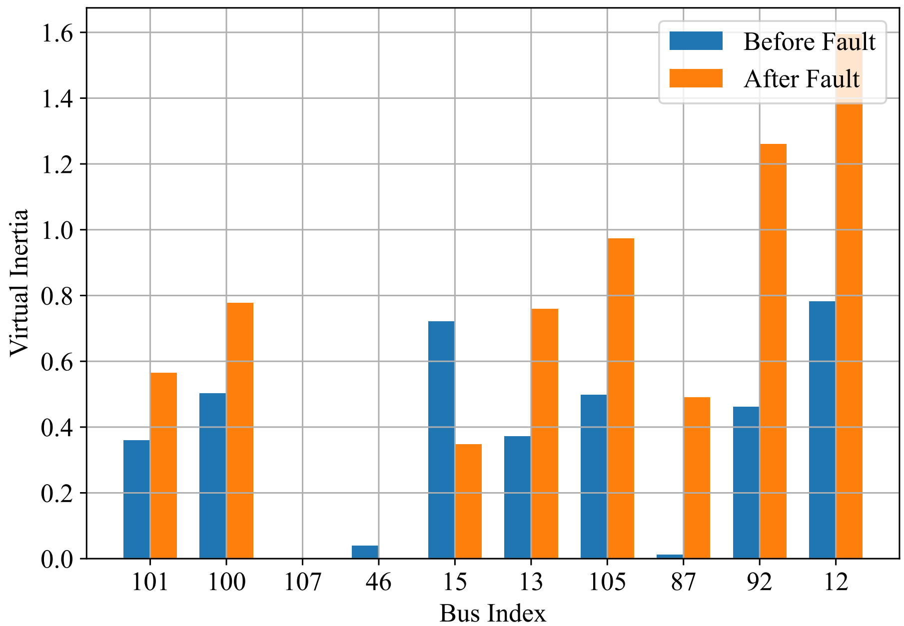

Figure 4.16: Before-and-after comparison of virtual inertia at the ten most sensitive buses for the contingency. The blue bars report the pre-fault allocation, and the orange bars report the post-fault allocation.

smaller individual changes. The fact that two buses dominate the ranking by a wide margin indicates that the post-fault reallocation is highly localized. The compensating reductions required by the budget constraint are spread across the rest of the network in small per-bus amounts, with bus 15 providing the most visible single contribution to that compensation in the top-ten subset.

Figure 4.16 resolves the magnitude of each change directly by showing the pre- and post-fault allocations as paired bars. Bus 12 exhibits the largest absolute increase, with its allocation roughly doubling across the disturbance. Bus 92 undergoes an even larger relative increase, moving from a low pre-fault value to a level comparable to that of bus 12 after the fault. Buses 87, 105, and 13 all receive additional inertia after the fault, while buses 100 and 101 rise by more modest amounts. Bus 15 is the only bus in the top-ten subset whose allocation is clearly drawn down, consistent with its role in providing part of the compensating reduction required by the budget constraint. The pattern confirms the observation from the heatmap. The redistribution is not a uniform increase but a structured reallocation, in which buses whose modal participation grows after the topology change are reinforced at the expense of buses whose role in the post-fault modes is reduced. The pattern aligns with the blending-and-projection rule of (3.34), in which the spectral reference and the RL-proposed adjustment are combined before projection onto the feasible set. As a result, the executed allocation tracks the updated physics-guided target while remaining within the prescribed bounds.

The four views of the allocation therefore give the same picture. The framework reproduces the spatial structure of the spectral-sensitivity reference under contingency conditions, responds to topology changes by reallocating inertia toward the buses whose post-fault participation in critical modes is largest, and does so under a fixed total budget that induces compensating reductions elsewhere in the network. This alignment provides a reliable basis for integrating virtual inertia into the coordinated control actions evaluated in the remaining subsections.

4.5.2 Demand Response Behaviour

Attention now turns to the demand-response actions selected by the controller. The question addressed here is whether curtailment decisions track the evolution of system stress over time and across space, or follow a less discriminating pattern. Demand response enters the reward function through the economic term in (3.29) and is expected to operate as a corrective action that the RL policy invokes selectively rather than continuously.

Figure 4.17 shows the DR activation across PQ buses over the full 24-hour horizon, ex-

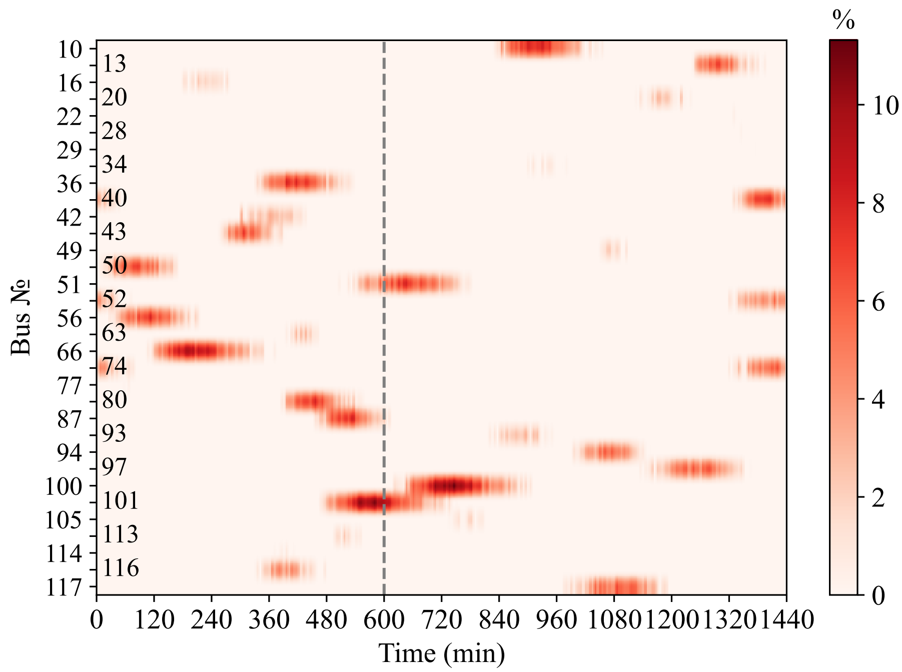

Figure 4.17: P

redicted demand-response activation at PQ buses over a 24-hour horizon under the N–2 contingency, expressed as a percentage of baseline demand. Each row corresponds to a distinct PQ bus; the colour encodes the curtailment level as a percentage of that bus's baseline demand, ranging from white (no curtailment) to dark (maximum activation). The vertical dashed line marks the fault at $t = 600$ min.

pressed as a percentage of baseline demand. The map is predominantly white, with nonzero activation confined to two distinct windows and to a small subset of buses within each window. The first window is located before the fault and covers the morning demand ramp. Within this window, activations appear at buses 50 and 56 in the early morning, at buses 36 and 43 during the late-morning rise, and at buses 80, 87, 101, and 105 as demand approaches its midday peak. The second window follows the fault and is concentrated in two separate bands. A short-lived band on bus 100 between approximately $t = 600$ and $t = 840$ min reflects the immediate post-fault redistribution, and a longer evening band

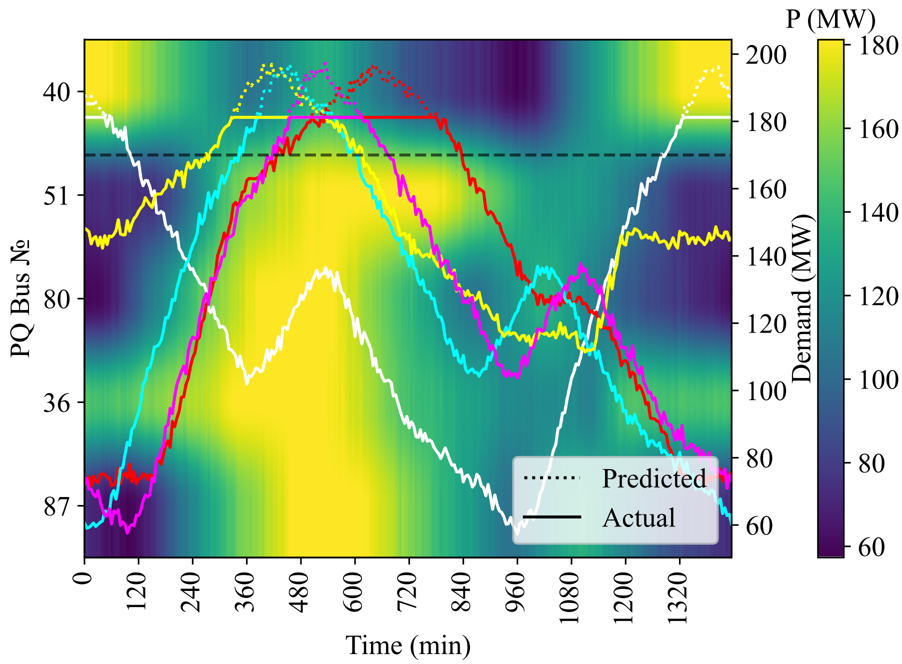

Figure 4.18: Demand-response profiles for the five representative PQ buses listed on the vertical axis (buses 40, 51, 80, 36, and 87). The background shows the demand field across the selected buses, and the overlaid traces compare the baseline predicted demand (dotted) with the served demand after DR (solid). Each coloured trace corresponds to the bus at the matching position on the vertical axis. The horizontal dashed line marks the curtailment threshold.

between roughly $t = 1080$ and $t = 1440$ min activates buses such as 10, 13, 52, 56, and 117 as the second daily demand rise interacts with the post-fault topology. Peak activation values reach approximately 8–10% of baseline demand at the most stressed buses, while the majority of PQ buses exhibit zero curtailment throughout the horizon. The sparsity of the map reflects the construction of the reward in (3.29), in which the shedding penalty discourages curtailment whenever operating limits can be maintained through generator redispatch alone.

Figure 4.18 traces the baseline and served demand for five representative PQ buses selected from the activation map. The predicted and served trajectories coincide across most of

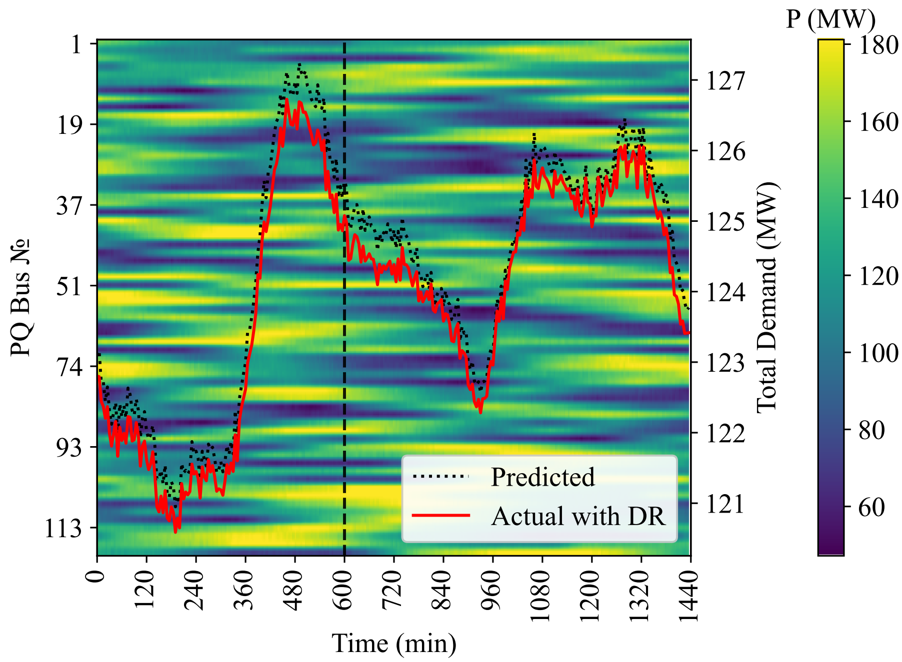

Figure 4.19: System-level demand response under the N–2 contingency. The system-average baseline demand (dotted) is compared with the served demand after DR (solid). The vertical dashed line marks the fault at $t = 600$ min.

the horizon, with separation confined to short intervals during which the baseline demand exceeds the curtailment threshold marked by the horizontal dashed line. Within those intervals, the served trajectory is held flat at the threshold while the predicted trajectory continues to rise, so that the controller curtails only the portion of demand that would otherwise exceed the operating limit. Once the baseline demand falls back below the threshold, the two trajectories rejoin without any residual curtailment. The same pattern is repeated across all buses shown, reproducing, at the per-bus level, the concentration of activation already visible at the system level and confirming that curtailment is triggered by the system's evolving state rather than by a preset schedule.

The aggregate view of the DR response is shown in Fig. 4.19, which compares the system-average baseline demand with the served demand after DR across all PQ buses. The two curves follow a common multi-peak profile across the horizon. They rise from approximately 124 MW at $t = 0$ to a morning maximum of roughly 127 MW near $t = 480$ min, then descend through a midday valley of approximately 123 MW near $t = 960$ min. A second evening peak of roughly 126 MW occurs near $t = 1320$ min, after which the curves fall back toward 124 MW at the end of the horizon. The dotted and solid curves are essentially coincident throughout the 24-hour window, with small but visible separations during the morning and evening peak intervals, where the activation map shows the most concentrated curtailment. The aggregate curtailment, therefore, remains a small fraction of system demand even under the faulted topology, indicating that DR is applied conservatively and serves as a feasibility-preserving action rather than a primary dispatch lever. The fault at $t = 600$ min coincides with the descent from the morning peak rather than with a moment of stress, which explains why the activation map in Fig. 4.17 does not exhibit a sharp post-fault spike and why the post-fault activation is instead distributed across the late-morning and evening periods that follow.

Figure 4.20 examines the per-bus allocation of curtailment more closely by plotting normalized demand trajectories for the ten highest-demand PQ buses over a window centred on the system-level peak at approximately 08:10. Each trajectory shows the baseline demand and the served demand after DR, with the shaded region between them representing the curtailment applied by the RL agent. The figure shows that curtailment is not uniformly distributed among the high-demand buses. Buses 102, 37, 114, and 52 exhibit a pronounced shaded region that extends from before the peak to well after it, indicating that the policy begins to curtail these buses as demand approaches its maximum and continues to do so until demand has receded. In contrast, buses 88, 81, 64, 43, and 82 show essentially no separation between the baseline and served trajectories throughout the

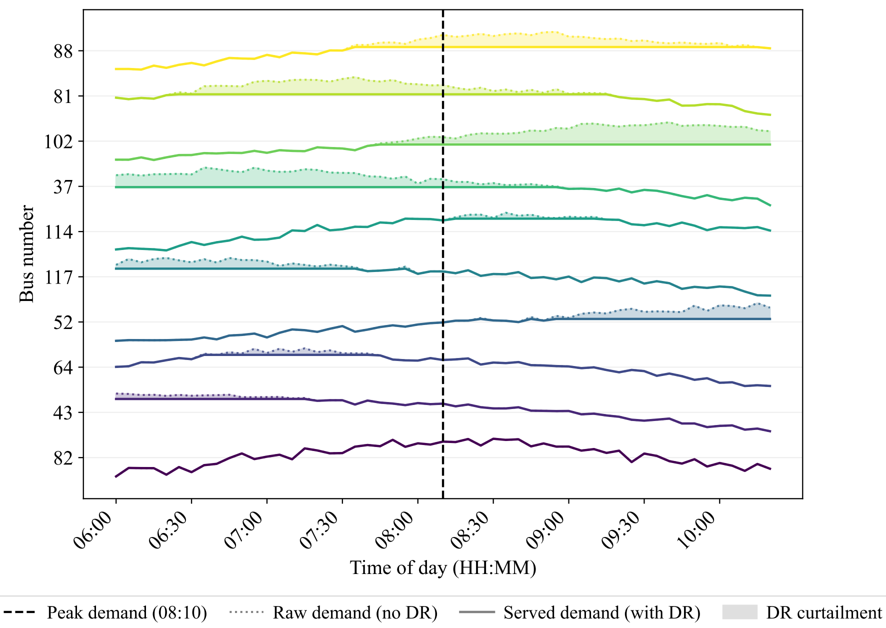

Figure 4.20: Top ten PQ bus demand trajectories around peak demand under the $N{-}2$ fault.

window, even at peak demand. The selective pattern confirms that DR is directed toward the most-stressed buses rather than spread uniformly across the high-demand subset, and that the policy's curtailment decisions are guided by the underlying spatial structure of demand stress rather than by a uniform load-reduction rule.

The four figures collectively indicate targeted, physically consistent demand-response behaviour under topology disturbances. The control actions are localized in both time and space, align with the intervals of elevated system stress identified in the activation map, and avoid unnecessary intervention during normal operating periods. This behaviour supports the reliability of the learned representation as an input to the coordinated control framework, whose aggregate performance is evaluated next.

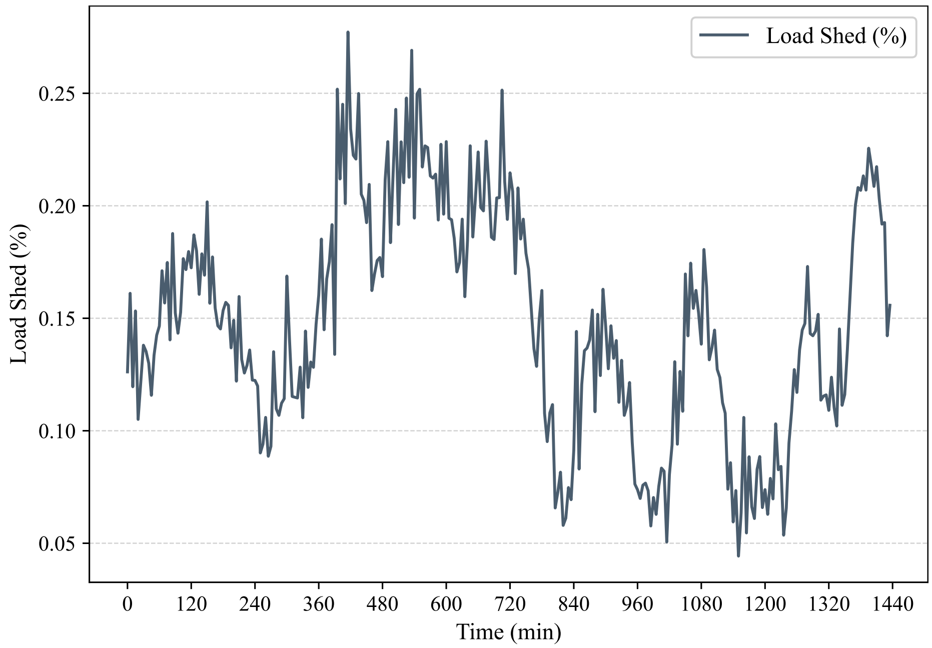

Figure 4.21: System-level load shedding over a 24-hour horizon under the N–2 contingency, expressed as a percentage of total system demand.

4.5.3 Closed-Loop Control Performance

The overall performance of the closed-loop controller is the final aspect considered in this section. Four system-level indicators are shown in sequence, corresponding to the components of the reward function in (3.28), and together they characterize how the integrated framework coordinates generator redispatch, demand response, and virtual inertia adjustments across the day. The controller operates at a five-minute decision interval, and at each step it receives the state representation assembled by the ST-GNN, comprising the predicted voltages, angles, and injections together with the topology-aware edge features, before projecting the proposed action onto the feasible operating region defined by the generator limits, DR bounds, and inertia budget as specified in the coordination mechanism of Section 3.6.

Figure 4.21 shows the system-level load-shed fraction, demand-weighted across PQ buses, over the 24-hour horizon. The trajectory tracks the daily load profile shown in Section 4.5.2, with the largest shedding event coinciding with the morning demand peak near $t = 480$ min, when aggregate load narrows the operating margin most severely, and voltage limits are progressively approached. During this period, DR is triggered more frequently to maintain feasibility. After the morning peak, the shed fraction declines as margins reopen and served demand is restored toward baseline levels. The fault at $t = 600$ min arrives during this descending phase, so the immediate post-fault interval does not coincide with a moment of peak operating stress. Post-fault shedding persists at a lower level throughout the afternoon and reactivates in the evening as the second daily demand rise interacts with the post-fault topology. Mechanistically, the shedding action $P_{t+1}^{\mathrm{sh},n}$ in (3.26) is applied when the predicted next-step evolution under (3.27) indicates that, absent curtailment, operating-limit violations would likely occur, and small curtailments are then applied at the most effective PQ buses to relieve stress and keep voltages and flows within the prescribed bounds.

Figure 4.22 shows the aggregated generation cost, which varies between roughly \$590 and \$640 across the day. The trajectory follows the load profile through the quadratic form of the unit cost curves $C_g(\cdot)$, with elevated values during the morning and evening peak windows where the dispatch leans more heavily on the higher-cost units and lower values during the early-morning and midday valleys where the system operates farther from unit limits. The fault at $t = 600$ min sits on the descending arm of the morning peak rather than at a separate inflection in the trajectory, and the cost remains within the same daily band afterwards rather than rising to a new operating regime. Generation cost and reward do not vary together at every instant. When demand response increases, less energy is served, and the evaluated cost from $C_g(\cdot)$ can decrease even though the reward declines

Figure 4.22: Aggregated generation cost over the 24-hour horizon under the N–2 contingency.

due to the shedding penalty. This trade-off is an immediate consequence of the multi-term reward construction in (3.28), which balances economic efficiency against resilience through the relative weights α_g and α_s in (3.29).

Figure 4.23 shows the number of bus-voltage violations recorded at each decision step, fluctuating between approximately 37 and 47 across the full horizon. The violation count tends to rise during high-load periods and around topology changes. This behaviour matches the state-transition structure in (3.27), since once the operating point sits close to the voltage-security limits, even small changes in injections or connectivity are amplified into visible deviations. When violations rise, the policy increases DR and redispatches to relax the local reactive demand and restore voltage margins, so that a decline in the violation count typically follows with a short lag. Although the violation term is not explicit in the reward decomposition in (3.28), it is treated as an auxiliary penalty aligned with the stability

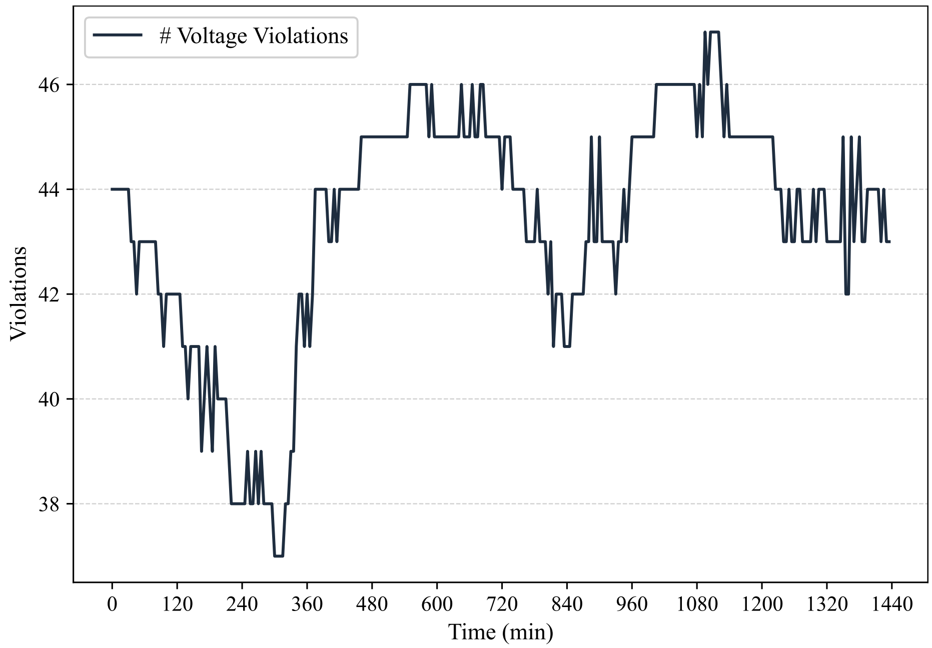

Figure 4.23: Number of bus-voltage violations over the 24-hour horizon under the N–2 contingency.

component in (3.30), and its temporal behaviour tracks the corrective action taken by the controller rather than sustained deterioration.

Figure 4.24 summarizes the total reward trajectory. The reward is driven primarily by the shedding signal, since the shedding term carries the larger weight in (3.29) to prioritize service continuity, with cost variations playing a secondary role under the chosen weights. The trajectory begins near -925 at $t = 0$ and rises to a local maximum of roughly -855 around $t = 240$ min as the system operates within a margin during the early-morning hours. It then drops to its lowest value of approximately -1070 near $t = 535$ min, when the morning demand peak coincides with the largest shedding event of the horizon. The reward then recovers through the fault, reaching a post-fault local maximum near -870 around $t = 840$ min as the system stabilizes on the new topology. For the remainder of the day, the reward oscillates between approximately -880 and -1000, with a brief evening

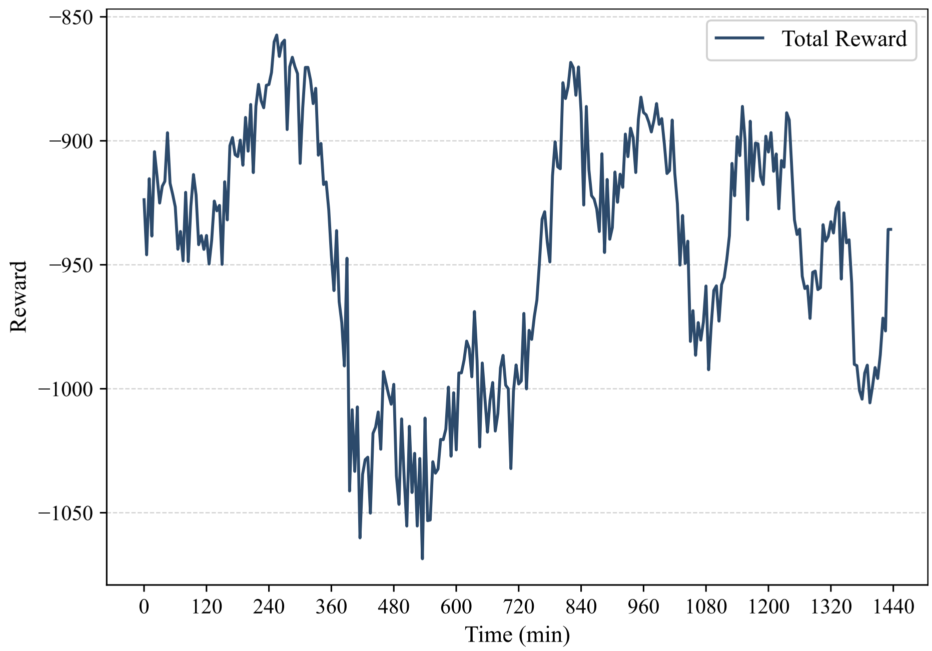

Figure 4.24: Total reward trajectory over the 24-hour horizon under the N–2 contingency.

dip toward -1010 near $t = 1320$ min that aligns with the second demand peak in Fig. 4.19. The fact that the global reward minimum occurs before the fault, rather than at or after the contingency time, indicates that the coordinated action of DR, redispatch, and virtual inertia absorbs the topology change with a moderate and bounded penalty rather than a renewed performance collapse. The temporal behaviour aligns with the Bellman structure in (3.33), in which action sequences that minimize shedding yield higher expected returns, and the reward tracks the weighted combination of the shedding, cost, and violation signals examined above.

The four indicators together describe a coherent closed-loop response under the RL formulation, in which rewards track the weighted combination of shedding and cost, violations modulate the need for corrective action, and the temporal patterns remain compatible with the state-transition structure and the optimization objective.

Taken together, the results presented in this chapter show that the proposed physics-informed framework achieves both predictive accuracy and stable closed-loop control under the N–2 contingency on the IEEE 118-bus system. The ST-GNN predictor reproduces voltage and angle profiles with small errors under the faulted topology; the virtual-inertia allocator delivers a topology-aware reference signal that remains temporally stable except at fault times; and the reinforcement-learning controller achieves feasible operation through a combination of targeted demand response and inertia reallocation. The greatest demand on the controller arises from the morning load peak rather than from the topology change itself, which indicates that the N–2 fault is absorbed with a smaller excursion in shedding, cost, and reward than the daily demand cycle imposes under normal operation. All source code, datasets, and trained models required to reproduce the results presented in this chapter are openly available [100].

Chapter 5

Conclusion and Future Work

This thesis addressed a central challenge in modern transmission systems: maintaining stability as inverter-based generation displaces the rotational inertia on which the grid has historically relied. The difficulty was not the absence of tools, but their fragmentation. Eigenvalue analysis, optimization-based allocation, model predictive control, and machine-learning predictors had each matured independently, yet none alone could provide online stability support that was simultaneously topology-aware, formally safe under stated conditions, and operationally interpretable. The contribution of this work is a single closed-loop framework in which physics-based modelling and learning-based decision making operate as parts of the same system rather than as competing alternatives.

The framework was built layer by layer. Chapter 2 began at the physical layer, formulating the swing-equation dynamics of the network and the linearized model used throughout the thesis. From this model, a sensitivity calculation quantified how much each bus's inertia affected the damping of the most fragile electromechanical modes, which in turn yielded a small optimization problem whose solution was a reference inertia profile H_{ref} adapted to the current operating point and topology. The result was a transparent, physics-based answer to where virtual inertia should be placed at any given moment.

A reference profile alone does not run a grid. Chapter 3 therefore embedded the allocation within a closed-loop architecture in which a graph neural network predicted where the system was heading next and a reinforcement-learning policy decided what to do about it. A supervisory layer reconciled the policy's proposed action with the physics-based reference and adjusted it as needed to ensure compliance with generator limits, demand bounds, and the feasible inertia range before anything was sent to the grid. A Lyapunov analysis built around this architecture established that, as long as inertia stayed within its design bounds and the network remained connected between switching events, the closed-loop response remained stable. The framework, therefore, retained the formal safety guarantees of classical control while gaining the adaptability of learning-based methods.

Chapter 4 tested the framework on the IEEE 118-bus system under N–2 contingencies not seen during training. The predictor accurately tracked the post-disturbance state of the system, the allocator concentrated virtual inertia at the buses where it was most effective for damping, and the controller kept the system within safe operating limits across the tested scenarios. A saliency analysis of the trained predictor highlighted the same transmission lines identified by the physics-based sensitivity analysis, providing two independent explanations for each control decision. In addition to reproducing system trajectories with low error, the closed-loop policy produced sparse demand-response interventions over limited time windows, and the realized virtual inertia allocations remained concentrated near stressed corridors while maintaining low per-step computational cost relative to OPF-based baselines.

What this work delivered, then, is not a single algorithm but a way of organizing stability support in low-inertia, topology-varying grids. Physics carried the load where transparency and safety guarantees were needed, learning carried it where adaptability and prediction were needed, and a coordination layer allowed the two to operate as parts of the same loop.

5.1 Contributions

This thesis closed a specific gap in the support for stability in low-inertia, topology-varying transmission systems by building a single, mathematically coherent pipeline rather than stitching together separate components. The pipeline began from the swing-equation dynamics of the network. From those dynamics emerged a closed-form sensitivity expression quantifying how each bus's inertia influenced the damping of the most fragile electromechanical modes. Those sensitivities were aggregated into bus-level criticality scores and formulated as a small, budget- and bound-constrained optimization problem whose solution was a reference inertia profile H_{ref} adapted to the current operating point and network topology. This sensitivity-guided allocation was the thesis's first contribution and the foundation on which everything else rested. It showed where virtual inertia mattered most in a way the operator could inspect.

A reference allocation alone does not stabilize a grid; a control policy has to act on it safely. The second contribution was the coordination architecture developed in Chapter 3, in which a supervisory layer blended the reinforcement-learning action with H_{ref} and adjusted the result to ensure compliance with generator limits, demand bounds, and the feasible inertia range before execution. The Lyapunov analysis built around this architecture showed that, as long as inertia stayed within design bounds and the network remained connected between switching events, the closed-loop electromechanical response preserved uniform asymptotic stability. The framework was therefore not only physics-aware but also provably stable in the small-signal sense, which is the property an operator most needs from any learning-augmented controller. For the control loop to be useful sequentially, the predictions feeding it had to remain physically consistent under conditions the model never explicitly saw. That requirement led to the third contribution, the spatio-temporal physics-informed graph neural network. By treating each time step as a graph-structured observation and embedding Kirchhoff's laws and operational limits as soft penalties during training, the predictor produced forecasts of voltages, angles, generation, demand, and virtual inertia that respected the underlying power-flow physics. The result was a predictor whose outputs were physically consistent and directly usable within the closed-loop control pipeline.

These three pieces, the allocator, the coordinator, and the predictor, were unified by the closed-loop pipeline of Section 3.1, summarized in Algorithm 3.1. This unification was itself the fourth contribution. Conventional pipelines treat prediction, stability assessment, and control as separate sequential processes. The architecture here tied them into a single decision cycle, in which the spectral module supplied the reference, the predictor supplied the near-term outlook, the policy supplied the action, and the coordinator enforced feasibility before anything reached the grid. The benefit was more than convenience: it was a closed-loop system whose physics, learning, and safety guarantees co-existed by construction rather than by negotiation.

A unified pipeline of this kind is only operationally meaningful if its decisions can be understood, and to that end, the fifth contribution was a set of interpretability tools based on edge-level saliency analysis of the trained predictor. The saliency maps surfaced the transmission lines that most strongly influenced each prediction, and they aligned with the eigenvalue-sensitivity rankings derived analytically in Chapter 2. Two independent diagnostics, one learning-based and one physics-based, identified the same critical structure. That convergence justified operator trust in the framework's outputs and made the predictor a transparent component rather than a black box.

The final contribution was the contingency-aware training methodology supporting all of the above. By generating training data across a representative set of single- and double-line contingencies on the IEEE 118-bus system, the predictor and policy were exposed to topology variations during learning rather than only at deployment. Chapter 4 showed that the resulting models retained accuracy and feasibility under N–2 contingencies not seen during training, addressing a known generalization gap in graph-based learning for power systems and showing that the framework held up under the kind of stress operational reliability planning is concerned with.

Together, these contributions form a single closed-loop framework for stability-aware operation of low-inertia transmission systems, in which physics-based modelling and learning-based decision making are not alternatives but layers of the same architecture.

5.2 Limitations and Future Work

5.2.1 Feasibility, Assumptions, and Validity of the Proposed Approach

The proposed framework rests on several assumptions that determine the conditions under which its results are valid. First, the spectral-sensitivity analysis and Lyapunov stability guarantee are derived from the small-signal linearization of the swing dynamics around a known operating point. This approximation holds for small perturbations but does not extend to large transient excursions, electromagnetic transients, or post-fault dynamics that drive the system far from the linearization point. Second, the Lyapunov stability guarantee is conditional: it holds provided the network graph remains connected between switching events and the inertia values remain within the design bounds enforced by the coordinator. Under severe multi-element contingencies that disconnect the network, or under large exogenous disturbances, these conditions may not hold and the guarantee does not extend. Third, the centralized training procedure assumes that bus-level electrical states, spectral sensitivities, and OPF solutions are available to a single operator. In practice, multi-utility grids and systems with significant prosumer participation may not satisfy this observability assumption. Fourth, the framework assumes that all inverter-interfaced buses in $\mathcal{G}_{inv}$ are fully controllable by the operator; in systems where some buses belong to different companies or operate under distributed ownership, this assumption may require a federated or hierarchical extension. Fifth, the validation is conducted on the IEEE 118-bus benchmark, which is a standard meshed transmission test system but does not represent a real utility network; generalization to larger systems, real PMU data streams, and networks with different inertia profiles remains to be demonstrated. Within these conditions, the framework produces physically consistent predictions, topology-aware inertia allocations, and formally stable closed-loop behaviour as demonstrated in Chapter 4.

5.2.2 Limitations

The framework rests on the small-signal linearization of the swing dynamics around an operating point. Within that range, the eigenvalue-sensitivity machinery, the allocation problem, and the Lyapunov result all hold cleanly, but the model does not capture large transient excursions, electromagnetic transients, or protection-system behaviour, and a different modelling layer would be needed to extend the framework to those regimes. The present implementation also adopts one-step prediction and a greedy control update, so model mismatch may accumulate near regime shifts such as fault onsets, potentially causing the closed-loop trajectory to drift from the training distribution. In addition, the current environment does not explicitly model large disturbances that split the network into multiple synchronous islands, nor does it capture the practical sparsity, certification constraints, and actuation uncertainty associated with real grid-forming resources that may be asked to supply virtual inertia. The action space is likewise restricted to re-dispatch, demand response, and virtual inertia adjustments, leaving out other stabilizing devices such as power system stabilizers and FACTS.

The validation in Chapter 4 was carried out on the IEEE 118-bus system under a sample of single- and double-line contingencies, which is enough to demonstrate the closed-loop behaviour but does not exhaust the space of credible disturbances or generalize automatically to larger or multi-area networks. Scaling the framework to larger systems, such as the IEEE 300-bus network or real utility-scale networks, would test the graph aggregation depth of the ST-GNN, the computational cost of the spectral-sensitivity module, and the feasibility of the OPF-based training data generation. Larger graphs may require deeper GNN architectures or hierarchical graph partitioning to maintain expressiveness, and the eigenvalue computation cost grows with system size. Assessing scalability on at least one larger benchmark is identified as a priority direction for future work. The training pipeline also assumed access to clean simulated data and did not address the noise, missing samples, or measurement delays that arise with real PMU streams. Finally, the contingency-aware training procedure incurs a nontrivial offline computational cost, which is acceptable for periodic retraining but would need to be re-examined for systems with frequent topology evolution.

Several extensions follow naturally from the limitations above. The most direct approach is to relax the linearization assumption by integrating nonlinear models of synchronous machines and inverter-based resources, thereby enabling evaluation of the framework under transient excursions where the small-signal approximation no longer holds. A second direction is to strengthen the closed-loop predictor and controller by moving beyond one-step greedy updates toward multi-step rollout or value targets, receding-horizon model-predictive control, and training procedures such as scheduled sampling or data aggregation that expose the predictor to its own closed-loop state distribution. A third direction is to explicitly incorporate uncertainty into the predictor and the policy, replacing point forecasts of demand and renewable generation with distributional ones and allowing the controller to anticipate variability rather than react to it. Extending the transition and control logic to detect islanding, support per-island objectives, and restrict virtual inertia actuation to certified eligible devices under uncertainty would further improve operational realism.

Broadening the action space to include additional stabilizing mechanisms, such as power system stabilizers and FACTS, would move the framework toward a more unified multi-policy controller. Validation on real PMU data, on larger test systems, and across multi-area networks is the natural next step toward operational readiness, and hardware-in-the-loop testing would help close the remaining gap between simulation and deployment. Together, these directions would extend the framework from a closed-loop demonstration on a benchmark grid into a tool capable of supporting stability decisions in the systems it is ultimately meant to serve.

References

[1] M. S. S. Danish and T. Senjyu. Shaping the future of sustainable energy through ai-enabled circular economy policies. *Circular Economy*, 2(2):100040, 2023.

[2] Yashen Lin, Joseph H. Eto, Brian B. Johnson, Jack D. Flicker, Robert H. Lasseter, Hugo N. Villegas Pico, Gab-Su Seo, Brian J. Pierre, Abraham Ellis, Jeremiah Miller, and Guohui Yuan. Pathways to the next-generation power system with inverter-based resources: Challenges and recommendations. *IEEE Electrification Magazine*, 10(1):10–21, 2022.

[3] J. Yang, Y. H. Song, and K. S. Kook. Critical inertia calculation method of generators using energy balance condition in power system. *Energies*, 17(5):1097, 2024.

[4] Philipp Andreas Gunkel, Henrik Klinge Jacobsen, Claire-Marie Bergaentzlé, Fabian Scheller, and Frits Møller Andersen. Variability in electricity consumption by category of consumer: The impact on electricity load profiles. *International Journal of Electrical Power & Energy Systems*, 147:108852, 2023.

[5] Tanate Kerdphol, F. S. Rahman, and Yasunori Mitani. Virtual inertia control application to enhance frequency stability of interconnected power systems with high renewable energy penetration. *Energies*, 11(4):981, 2018.

[6] 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*, 159:110067, 2024.

[7] Joshua W Busby, Kyri Baker, Morgan D Bazilian, Alex Q Gilbert, Emily Grubert, Varun Rai, Joshua D Rhodes, Sarang Shidore, Caitlin A Smith, and Michael E Webber. Cascading risks: Understanding the 2021 winter blackout in texas. *Energy Research & Social Science*, 77:102106, 2021.

[8] S. J. Yeboah, S. Nunoo, J. C. Attachie, W. D. Asihene, and E. Y. Asuamah. Impact and integration techniques of renewable energy sources on smart grid operations: A systematic review. *Scientific African*, 29:e02845, 2025.

[9] M. I. Saleem, S. Saha, and T. K. Roy. Inertia sensitivity analysis of power grids with high penetration of renewable energy sources. *IEEE Transactions on Industry Applications*, 61(4):5362–5380, 2025.

[10] NASA Earth Observatory. Smoke across north america. https://science.nasa.gov/earth/earth-observatory/smoke-across-north-america-148610/, 2023. Accessed: 2026-04-30.

[11] Environment and Climate Change Canada. Minimum temperature surface observations, 20 january 2025. ECCC Surface Observations, Government of Canada, 2025. Accessed: 2026-04-30.

[12] Nikolaos P. Theodorakatos. Application of synchronized phasor measurements units in power systems. *The International Journal of Engineering and Science (IJES)*, 6(3):25–39, 2017.

[13] Wadim Strielkowski, Andrey Vlasov, Kirill Selivanov, Konstantin Muraviev, and Vadim Shakhnov. Prospects and challenges of the machine learning and data-driven methods for the predictive analysis of power systems: A review. *Energies*, 16(10):4025, 2023.

[14] Saleha Tabassum, Attuluri Babu, and Dharmendra Dheer. A comprehensive exploration of iot-enabled smart grid systems: power quality issues, solutions, and challenges. *Science and Technology for Energy Transition*, 79, 2024.

[15] M. A. Mahmoud, N. R. Md Nasir, M. Gurunathan, P. Raj, and S. A. Mostafa. The current state of the art in research on predictive maintenance in smart grid distribution network: Fault's types, causes, and prediction methods—a systematic review. *Energies*, 14(16):5078, 2021.

[16] V. Hassija, V. Chamola, A. Mahapatra, and S. K. Ghosh. Interpreting black-box models: A review on explainable artificial intelligence. *Cognitive Computation*, 16:45–74, 2024.

[17] Reinhard Günther. How data analytics and predictive maintenance are boosting efficiency in smart grids, May 2023. Website article, last updated October 5, 2024.

[18] J. G. Hoffer, A. B. Ofner, F. M. Rohrhofer, et al. Theory-inspired machine learning—towards a synergy between knowledge and data. *Welding in the World*, 66:1291–1304, 2022.

[19] J. Bae. Cost-effective placement of phasor measurement units to defend against false data injection attacks on power grid. *Energies*, 13(15):3862, 2020.

[20] L. Oneto, S. Ridella, and D. Anguita. Informed machine learning: Excess risk and generalization. *Neurocomputing*, 646:130521, 2025.

[21] Z. Zhou, R. Basker, and D. Yeung. Graph neural networks for multivariate time-series forecasting via learning hierarchical spatiotemporal dependencies. *Engineering Applications of Artificial Intelligence*, 147:110304, 2025.

[22] Anna Vettoruzzo, Mohamed-Rafik Bouguelia, Joaquin Vanschoren, Thorsteinn Rögnvaldsson, and KC Santosh. Advances and challenges in meta-learning: A technical review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(7):4763–4779, 2024.

[23] Mohammad Heidari, Li Ding, Mohammad Kheshti, Wei Bao, Xiangyu Zhao, Metodi Popov, and Vladimir Terzija. A review on application of machine learning-based methods for power system inertia monitoring. *International Journal of Electrical Power & Energy Systems*, 162:110279, 2024.

[24] Y. Zhang, M. Han, and M. Zhan. The concept and understanding of synchronous stability in power electronic-based power systems. *Energies*, 16(6):2923, 2023.

[25] Qi Qiu, Rui Ma, Jürgen Kurths, and Meng Zhan. Swing equation in power systems: Approximate analytical solution and bifurcation curve estimate. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(1), 2020.

[26] Mathaios Panteli and Pierluigi Mancarella. Modeling and evaluating the resilience of critical electrical power infrastructure to extreme weather events. *IEEE Systems Journal*, 11(3):1733–1742, 2015.

[27] Ali Tayyebi, Adolfo Anta, and Florian Dörfler. Almost globally stable grid-forming hybrid angle control. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 830–835, 2020.

[28] Cong He, Hua Geng, Kaushik Rajashekara, and Ambrish Chandra. Analysis and control of frequency stability in low-inertia power systems: A review. *IEEE/CAA Journal of Automatica Sinica*, 11(12):2363–2383, December 2024.

[29] Tomislav Baškarad, Ninoslav Holjevac, and Igor Kuzle. A new perspective on frequency control in conventional and future interconnected power systems. *International Journal of Electrical Power & Energy Systems*, 156:109731, 2024.

[30] Anh-Tuan Tran and Van Van Huynh. Load frequency control design for complex power systems implementing integral single-phase sliding mode control. *Engineering, Technology & Applied Science Research*, 15(1):19802–19808, 2025.

[31] Xinrui Lin, Jimmy Chih-Hsien Peng, David Macii, Dario Petri, Junwei Yu, and He Wen. Frequency-domain small-signal stability analysis methods for grid-following converters systems—an overview. *Renewable and Sustainable Energy Reviews*, 211:115283, 2025.

[32] Nikos Hatziargyriou, Jovica Milanovic, Claudia Rahmann, Venkataramana Ajjarapu, Claudio Canizares, Istvan Erlich, David Hill, Ian Hiskens, Innocent Kamwa, Bikash Pal, Pouyan Pourbeik, Juan Sanchez-Gasca, Aleksandar Stankovic, Thierry Van Cutsem, Vijay Vittal, and Costas Vournas. Definition and classification of power system stability – revisited & extended. *IEEE Transactions on Power Systems*, 36(4):3271–3281, 2021.

[33] Chenxi Liu, Huanhai Xin, Di Wu, Huisheng Gao, Hui Yuan, and Yuhan Zhou. Generalized operational short-circuit ratio for grid strength assessment in power systems with high renewable penetration. *IEEE Transactions on Power Systems*, 39(4):5479–5494, 2024.

[34] Robert H. Lasseter, Zhe Chen, and Dinesh Pattabiraman. Grid-forming inverters: A critical asset for the power grid. *IEEE Journal of Emerging and Selected Topics in Power Electronics*, 8(2):925–935, 2020.

[35] Roberto Rosso, Xiongfei Wang, Marco Liserre, Xiaonan Lu, and Soenke Engelken. Grid-forming converters: Control approaches, grid-synchronization, and future trends—a review. *IEEE Open Journal of Industry Applications*, 2:93–109, 2021.

[36] Xiuqiang He and Hua Geng. Pll synchronization stability of grid-connected multi-converter systems with a reference choice of stationary point. *IEEE Transactions on Power Systems*, 37(3):2087–2099, 2022.

[37] Ali Arjomandi-Nezhad, Yifei Guo, Bikash C. Pal, and Guangya Yang. Modeling fault recovery and transient stability of grid-forming converters equipped with current reference limitation. *IEEE Transactions on Power Systems*, 40(2):1735–1747, 2025.

[38] Yu. V. Makarov, Zhao Yang Dong, and D. J. Hill. A general method for small signal stability analysis. *IEEE Transactions on Power Systems*, 13(3):979–985, 1998.

[39] Hong Li, Yuanye Lu, Jinchang Pan, and Mingbo Wei. An eigenvalue sensitivity-based stability improvement control method of vsg. In *Proceedings of the IECON 2023–49th Annual Conference of the IEEE Industrial Electronics Society*, pages 1–6. IEEE, 2023.

[40] Otavio Bertozzi, Harold R Chamorro, Edgar O Gomez-Diaz, Michelle S Chong, and Shehab Ahmed. Application of data-driven methods in power systems analysis and control. *IET Energy Systems Integration*, 6(3):197–212, 2023.

[41] Anas Altawallbeh, Ahmed Alassi, Nader Meskin, Mohammed A. Al-Hitmi, and Ahmed M. Massoud. Small-signal stability analysis and parameters optimization of

virtual synchronous generator for low-inertia power system. *IEEE Access*, 13:107227–107243, 2025.

[42] Federico Milano, Florian Dörfler, Gabriela Hug, David J. Hill, and Gregor Verbič. Foundations and challenges of low-inertia systems (invited paper). In *2018 Power Systems Computation Conference (PSCC)*, pages 1–25, 2018.

[43] H. Bulat, D. Franković, and S. Vlahinić. Enhanced contingency analysis—a power system operator tool. *Energies*, 14(4):923, 2021.

[44] Hassan K Khalil and Jessy W Grizzle. *Nonlinear systems*, volume 3. Prentice hall Upper Saddle River, NJ, 2002.

[45] Sara Deilami, Amir S Masoum, Paul S Moses, and Mohammad AS Masoum. Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile. *IEEE Transactions on smart grid*, 2(3):456–467, 2011.

[46] O. Bertozzi, H. R. Chamorro, E. O. Gomez-Diaz, M. S. Chong, and S. Ahmed. Application of data-driven methods in power systems analysis and control. *IET Energy Systems Integration*, 6(3):197–212, 2024.

[47] J. Horak, J. Vrbka, and P. Suler. Support vector machine methods and artificial neural networks used for the development of bankruptcy prediction models and their comparison. *Journal of Risk and Financial Management*, 13(3):60, 2020.

[48] Kayode Adetunji, Ivan Hofsajer, and Ling Cheng. A coordinated charging model for electric vehicles in a smart grid using whale optimization algorithm. In *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, pages 1–7. IEEE, 2020.

[49] A. Zeghina, A. Leborgne, F. Le Ber, and A. Vacavant. Deep learning on spatiotemporal graphs: A systematic review, methodological landscape, and research opportunities. *Neurocomputing*, 594:127861, 2024.

[50] M. Gautam. Deep reinforcement learning for resilient power and energy systems: Progress, prospects, and future avenues. *Electricity*, 4(4):336–380, 2023.

[51] Farrokh Aminifar, Mahmud Fotuhi-Firuzabad, Amir Safdarian, Ali Davoudi, and Mohammad Shahidehpour. Synchrophasor measurement technology in power systems: Panorama and state-of-the-art. *IEEE Transactions on Smart Grid*, 2:1607–1628, 2014.

[52] T. Ahmad and H. Chen. A review on machine learning forecasting growth trends and their real-time applications in different energy systems. *Sustainable Cities and Society*, 54:102010, 2020.

[53] Robert I. Hamilton and Panagiotis N. Papadopoulos. Using shap values and machine learning to understand trends in the transient stability limit. *IEEE Transactions on Power Systems*, 39(1):1384–1397, 2024.

[54] Tingjian Zhang, Mingyang Sun, Jochen L. Cremer, Ning Zhang, Goran Strbac, and Chongqing Kang. A confidence-aware machine learning framework for dynamic security assessment. *IEEE Transactions on Power Systems*, 36(5):3907–3920, 2021.

[55] Ziming Yan and Yan Xu. Real-time optimal power flow: A lagrangian based deep reinforcement learning approach. *IEEE Transactions on Power Systems*, 35(4):3270–3273, 2020.

[56] Damian Owerko, Fernando Gama, and Alejandro Ribeiro. Optimal power flow using graph neural networks. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5930–5934, 2020.

[57] Ganesh K. Venayagamoorthy, Dulip Madurasinghe, and Rajan Ratnakumar. Graph models of electric power transmission networks using synchrophasor data. In *SoutheastCon 2024*, pages 592–597. IEEE, 2024.

[58] Wenlong Liao, Birgitte Bak-Jensen, Jayakrishnan Radhakrishna Pillai, Yuelong Wang, and Yusen Wang. A review of graph neural networks and their applications in power systems. *Journal of Modern Power Systems and Clean Energy*, 10(2):345–360, 2021.

[59] B. Donon, R. Clément, B. Donnot, A. Marot, I. Guyon, and M. Schoenauer. Neural networks for power flow: Graph neural solver. *Electric Power Systems Research*, 189:106547, 2020.

[60] Jinxian Zhang, Junbo Zhao, Fei Ding, Jing Yang, and Junhui Zhao. A graph convolutional network for active distribution system anomaly detection considering measurement spatial-temporal correlations. In *2023 North American Power Symposium (NAPS)*, pages 1–6. IEEE, 2023.

[61] E. Lodhi, N. Dahmani, S. M. S. Bukhari, S. Gyawali, S. Thapa, L. Qiu, M. H. Zafar, and N. Akhtar. Enhancing microgrid forecasting accuracy with saq-mtclstm: A self-adjusting quantized multi-task convlstm for optimized solar power and load demand predictions. *Energy Conversion and Management: X*, 24:100767, 2024.

[62] Yexiang Chen, Subhash Lakshminarayana, Carsten Maple, and H. Vincent Poor. A meta-learning approach to the optimal power flow problem under topology reconfigurations. *IEEE Open Access Journal of Power and Energy*, 9:109–120, 2022.

[63] Christian Nauck, Michael Lindner, Konstantin Schürholt, Haoming Zhang, Paul Schultz, Jürgen Kurths, Ingrid Isenhardt, and Frank Hellmann. Predicting basin stability of power grids using graph neural networks. *New Journal of Physics*, 24(4):043041, 2022.

[64] Bang LH Nguyen, Tuyen V Vu, Thai-Thanh Nguyen, Mayank Panwar, and Rob Hovsapian. Spatial-temporal recurrent graph neural networks for fault diagnostics in power distribution systems. *IEEE Access*, 11:46039–46050, 2023.

[65] Shuaibo Wang, Xinyuan Xiang, Jie Zhang, Zhuohang Liang, Shufang Li, Peilin Zhong, Jie Zeng, and Chenguang Wang. A multi-task spatiotemporal graph neural network for transient stability and state prediction in power systems. *Energies*, 18(6):1531, 2025.

[66] H. Sun, S. Li, J. Huang, H. Li, G. Jing, Y. Tao, and X. Tian. Dynamic spatial–temporal graph neural network for cooling capacity prediction in hvdc systems. *Energies*, 18(2):313, 2025.

[67] Caleb Kojo Eshun, Shahab Bahrami, and Javad Fattahi. Robust local contingency planning in power grids using spatio-temporal graph neural networks. In *2025 IEEE Electrical Power and Energy Conference (EPEC)*, pages 152–156, 2025.

[68] Di Cao, Junbo Zhao, Jiaxiang Hu, Yansong Pei, Qi Huang, Zhe Chen, and Weihao Hu. Physics-informed graphical representation-enabled deep reinforcement learning

for robust distribution system voltage control. *IEEE Transactions on Smart Grid*, 15(1):233–246, 2023.

[69] Jules Authier, Rabab Haider, Anuradha Annaswamy, and Florian Dörfler. Physics-informed graph neural network for dynamic reconfiguration of power systems. *Electric Power Systems Research*, 235:110817, 2024.

[70] Y. Wu, B. Sicard, and S. A. Gadsden. Physics-informed machine learning: A comprehensive review on applications in anomaly detection and condition monitoring. *Expert Systems with Applications*, 255:124678, 2024.

[71] A.K. Khamis and M. Agamy. Circuit topology aware gnn-based multi-variable model for dc-dc converters dynamics prediction in ccm and dcm. *Neural Computing and Applications*, 36:20807–20822, 2024.

[72] Quang-Ha Ngo, Bang LH Nguyen, Tuyen V Vu, Jianhua Zhang, and Tuan Ngo. Physics-informed graphical neural network for power system state estimation. *Applied Energy*, 358:122602, 2024.

[73] Huayi Wu, Zhao Xu, Minghao Wang, and Xue Lyu. Meta-learning enhanced physics-informed graph attention convolutional network for distribution power system state estimation. *IEEE Transactions on Network Science and Engineering*, 2025.

[74] Xiang Pan, Minghua Chen, Tianyu Zhao, and Steven H. Low. Deepopf: A feasibility-optimized deep neural network approach for ac optimal power flow problems. *IEEE Systems Journal*, 17(1):673–683, 2023.

[75] Thanh Long Vu, Sayak Mukherjee, Tim Yin, Renke Huang, Jie Tan, and Qiuhua Huang. Safe reinforcement learning for emergency load shedding of power systems. In *2021 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5, 2021.

[76] Xinyang Chen, Guangyi Qu, Yong Tang, Steven Low, and Neng Li. Reinforcement learning for selective key applications in power systems: Recent advances and future challenges. *arXiv preprint arXiv:2105.00770*, 2021.

[77] Qiuhua Huang, Renke Huang, Weituo Hao, Jie Tan, Rui Fan, and Zhenyu Huang. Adaptive power system emergency control using deep reinforcement learning. *IEEE Transactions on Smart Grid*, 11(2):1171–1182, 2020.

[78] Jean Toubeau, Babak Bakhshideh Zad, Mathieu Hupez, Zeger De Grève, and François Vallée. Deep reinforcement learning-based voltage control to deal with model uncertainties in distribution networks. *Energies*, 13(15):3928, 2020.

[79] J. Huang, H. Zhang, D. Tian, Z. Zhang, C. Yu, and G. P. Hancke. Multi-agent deep reinforcement learning with enhanced collaboration for distribution network voltage control. *Engineering Applications of Artificial Intelligence*, 134:108677, 2024.

[80] Jian Xie and Wei Sun. Distributional deep reinforcement learning-based emergency frequency control. *IEEE Transactions on Power Systems*, 37(4):2720–2730, 2022.

[81] Simon Stock, Davood Babazadeh, Philipp Hund, and Christian Becker. Reinforcement learning based coordination of virtual inertia provision from inverter-dominated distribution grids. In *2023 IEEE 32nd International Symposium on Industrial Electronics (ISIE)*, pages 1–6. IEEE, 2023.

[82] C. Lozada, W. Vargas, N. Granda, and M. Chamba. Optimal load shedding scheme considering the dynamic frequency response. *Engineering Proceedings*, 77(1):25, 2024.

[83] Renke Huang, Yujiao Chen, Tianzhixi Yin, Qiuhua Huang, Jie Tan, Wenhao Yu, Xinya Li, Ang Li, and Yan Du. Learning and fast adaptation for grid emergency

control via deep meta reinforcement learning. *IEEE Transactions on Power Systems*, 37(6):4168–4178, 2022.

[84] W. Zhu, S. Wen, Q. Zhao, B. Zhang, Y. Huang, and M. Zhu. Deep reinforcement learning based optimal operation of low-carbon island microgrid with high renewables and hybrid hydrogen–energy storage system. *Journal of Marine Science and Engineering*, 13(2):225, 2025.

[85] N. S. Ungureanu and A. Petrovan. Analytical research regarding the implementation of a prescriptive maintenance system applied to it&c equipment. *Machines*, 13(5):412, 2025.

[86] Ruisheng Diao, Zhiwei Wang, Di Shi, Qianyun Chang, Jiajun Duan, and Xiaohu Zhang. Autonomous voltage control for grid operation using deep reinforcement learning. In *2019 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5, 2019.

[87] Clayton Barrows, Aaron Bloom, Ali Ehlen, Jussi Ikäheimo, Jennie Jorgenson, Dheepak Krishnamurthy, Jessica Lau, Brendan McBennett, Matthew O'Connell, Eugene Preston, Andrea Staid, Gord Stephen, and Jean-Paul Watson. The ieee reliability test system: A proposed 2019 update. *IEEE Transactions on Power Systems*, 35(1):119–127, 2020.

[88] Mikhak Samadi and Javad Fattahi. Low-inertia microgrid synchronization using data-driven digital twins. *IEEE Access*, 12:78534–78548, 2024.

[89] A. Saleh, H. M. Hasanien, B. Turdybek, M. Alharbi, F. Jurado, and W. A. Omran. Optimal model predictive control for virtual inertia control of autonomous microgrids. *Sustainability*, 15(6):5009, 2023.

[90] Lexuan Meng, Jawwad Zafar, Shafiuzzaman K. Khadem, Alan Collinson, Kyle C. Murchie, Federico Coffele, and Graeme M. Burt. Fast frequency response from energy storage systems—a review of grid standards, projects and technical issues. *IEEE Transactions on Smart Grid*, 11(2):1566–1581, 2020.

[91] Wei Du, Robert H. Lasseter, and Amrit S. Khalsa. Survivability of autonomous microgrid during overload events. *IEEE Transactions on Smart Grid*, 10(4):3515–3524, 2019.

[92] Dinesh Pattabiraman, Robert H. Lasseter, and Thomas M. Jahns. Comparison of grid following and grid forming control for a high inverter penetration power system. In *2018 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5, 2018.

[93] Mahmoud Sarwar, Deepak Ramasubramanian, Wenzong Wang, and Venkataramana Ajjarapu. A sensitivity-based approach for grid-forming inverter allocation in power system to improve voltage performance. In *2024 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2024.

[94] M. Ghofrani. Sensitivity analysis of distribution network reconfiguration optimization for electric vehicle and renewable distributed generator integration. *Energies*, 18(8):1903, 2025.

[95] Qifan Chen, Siqi Bu, Huaiyuan Wang, and Chao Lei. Real-time multi-stability risk assessment and visualization of power systems: A graph neural network-based method. *IEEE Transactions on Power Systems*, 2024.

[96] Jochen L. Cremer, Ioannis Konstantelos, and Goran Strbac. From optimization-based machine learning to interpretable security rules for operation. *IEEE Transactions on Power Systems*, 34(5):3826–3836, 2019.

[97] Yassine Zahraoui, Triin Korõtko, Argo Rosin, Saad Mekhilef, Mehdi Seyedmahmoudian, Alex Stojcevski, and Ibrahim Alhamrouni. Ai applications to enhance resilience in power systems and microgrids—a review. *Sustainability*, 16(12):4959, 2024.

[98] Caleb Kojo Eshun, Somayeh Bahrami, and Javad Fattahi. Robust local contingency planning in power grids using spatio-temporal graph neural networks. In *2025 IEEE Electrical Power and Energy Conference (EPEC)*, pages 152–156, 2025.

[99] Prabha Kundur. *Power System Stability and Control*. McGraw-Hill, New York, 1994.

[100] Caleb Kojo Eshun. Physics-Informed Graph Neural Networks for Power-Grid Stability: Source code and datasets. `https://github.com/cal-esh/Physics-informed-GNN-Grid-Stability`, 2026. GitHub repository. Accessed: 2026-07-28.