Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics-Informed Neural Network Architecture for Real-Time Dynamic Security Assessment in Low-Inertia Transmission Networks

Venkatesh Daggupati

venkateshdaggupati2003@gmail.com

Texas Tech University

Research Article

Keywords:

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11637283/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Physics-Informed Neural Network Architecture for Real-Time Dynamic Security Assessment in Low-Inertia Transmission Networks

Venkatesh Daggupati[1*]

[1*] Department of Electrical Engineering, Texas Tech University, Lubbock, Texas, USA .

Corresponding author(s). E-mail(s): venkateshdaggupati2003@gmail.com;

Abstract

In power systems, the fast transition towards converter interfaced generations (CIGs) has reduced the inertia available to the transmission system, thus increasing its susceptibility to large disturbances. In low inertia conditions, fault and severe contingencies cause the system to experience sharp frequency fluctuations, rapid rotor angle divergence and poor voltage recovery. Dynamic security assessment (DSA) methods that rely on stiff non-linear differential algebraic equations (DAEs) are computationally expensive and challenging to apply in practice. Purely data driven alternatives often struggle to generalise to unseen operating scenarios. To address these challenges, we propose a physics informed neural network (PINN) that can be used for real time DSA in low inertia transmission networks. This approach includes multi-objective regularisation that incorporates centre of inertia (COI) swing equations, converter current tracking and semiconductor current limits into the loss function using Automatic Differentiation. The network classifies whether a scenario is secure (i.e. transiently unstable is defined as the positive class) and predicts the post-fault trajectory. The ensemble mean accuracy across an adapted IEEE 39 bus, 10 machine benchmark system with renewable penetration levels up to 70% was found to be 98.7% $\pm$ 0.2% (98.40% in median representative run). Accuracies in unseen topological $\boldsymbol{N-2}$ contingencies were found to be 94.5% $\pm$ 0.3%. The average false secure (false negative) rate was found to be 2.22% $\pm$ 0.18% (2.22% in representative run). This model can be implemented in real time with inference times of 7.2 ms on a dedicated GPU (21.4 ms estimated contingency screening time). This shows that it could be used for online security screening in the benchmark studied.

Keywords: Converter-interfaced generation, dynamic security assessment, low-inertia systems, physics-informed neural networks, real-time screening, transient stability.

1 Introduction

The integration of inverter-based resources (IBRs), such as large-scale photovoltaics and wind farms [1, 2], has increased in recent years due to the transition to carbon neutral energy systems. This has changed the mechanism by which the grid becomes stable during disturbances, compared to conventional systems where large synchronous generators provide inertial response as well as synchronising torque. IBRs are connected to the AC grid via a converter interface and do not inherently have inertia [3]. Due to the removal of synchronous units, H_{sys} will drop across the system. This makes it more susceptible to severe contingencies such as multi-phase faults, generator outages and transmission corridor outages [4]. Low inertia networks will exhibit large changes in rate of change of frequency (RoCoF), due to the large mismatch in the active power at the moment of the fault. This will cause deep frequency dips. Additionally, the current limit of converter devices (typically about 1.1–1.3 p.u. of rated current) will limit the amount of reactive support that can be provided during voltage depression. This can cause faster separation between the rotor angles of the remaining synchronous units [5]. In this situation, safety actions should be taken within a very short timeframe (often less than 500ms) [6].

Dynamic security assessment (DSA) is an essential part of proactive transient contingency screening [7]. Online DSA is typically performed using time domain analysis tools which can solve large non-linear DAEs [9, 10]. Whilst these methods provide great accuracy in capturing the dynamics of the controllers and electromagnetic devices included, the computation time can become very large when considering large networks and multiple contingencies. In highly stressed converter dominated networks, stiff dynamics of the converter inner control loops and phase locked loops (PLLs) means that very small time-steps are required ($\Delta t \leq 1$ ms). This causes delays in the real-time scanning of contingencies in wide area networks [11].

To overcome this, many data driven ML techniques have been implemented for fast DSA. These include support vector machines, random forests and deep feed forward networks [12–14]. These can assess contingencies within milliseconds by learning the relationship between PMUs and transient stability. However, data driven methods act as black-box estimators, meaning they have no constraints. They can therefore produce unrealistic results if the data is outside of the training set, for example unexpected $N-2$ line outages, topology changes or unknown levels of renewable generation [15, 16].

Physics informed neural networks (PINNs) overcome this issue by incorporating analytical differential equations into the loss function of the network using auto differentiation [17, 18]. Initial studies within the power industry show great potential for using PINNs for dynamic state estimation [19], power flow surrogate models [21] and estimating transient stability boundaries [8, 20]. However, current physics informed DSA methods have the following issues within modern low inertia transmission networks:

1. *Omission of Non-Smooth Converter Current Dynamics:* This work considers transient stability problems that have been recently tackled using PINNs (e.g., [20, 21]), where the dynamics of the system are governed by conventional second order swing equations for synchronous machines or algebraic equations for power balance ($\mathbf{g}(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{0}$). This assumes that only the inertial dynamics of the rotor and continuous changes in flux linkages govern the power system dynamics. This assumption is no longer valid in converter-dominant systems where CIG controller currents become non-linear saturated currents ($I_d^2 + I_q^2 \leq I_{\max}^2$) during large voltage drops. Since converter inner-loop current dynamics and hard limits on semiconductor overcurrents are not considered in existing PINN frameworks, the regularised loss landscapes are unable to capture the non-smooth structural changes in the system.
2. *Reliance on Uncoupled or Simplified Topology Assumptions:* Most recent physics informed models apply to simple single machine infinite bus (SMIB) systems or assume constant admittance matrices [19, 20]. They therefore cannot cope with large topology changes (e.g. double circuit $N-2$ transmission line outages), where transfer reactances and local short-circuit ratios vary rapidly.
3. *Absence of Multi-Attribute Dynamic Stability Criteria in Inverter-Dominated Grids:* Existing PINN stability classifiers rely on a single-attribute criterion, the rotor angle separation ($\max_{i,j} |\delta_i - \delta_j| > 180°$). However in modern converter dominant systems, dynamic security collapse often occurs due to large RoCoF and frequency nadir violations that cause UFLS within 200-500 ms.

To address these limitations, this study develops a physics-informed deep learning architecture for real-time DSA in low-inertia transmission systems. Unlike prior PINN formulations that embed only synchronous-machine swing dynamics, the proposed framework couples an aggregated center-of-inertia (COI) power-balance model with explicit converter inner-loop trajectory decoders and nonlinear semiconductor current-saturation penalties ($I_d^2 + I_q^2 \leq I_{\max}^2$) within a unified, dimensionally normalized multi-objective loss function. The main technical contributions of this paper are:

- **Physics-Informed Dual-Task Deep Architecture:** We design a neural surrogate that maps pre-fault PMU measurements and sequence-aware contingency descriptors to binary security decisions (where instability is explicitly designated as the positive class to evaluate critical false-secure risks) while simultaneously outputting continuous post-contingency COI frequency and converter current trajectories constrained by the underlying physical DAEs via automatic differentiation.
- **Semiconductor-Constrained Multi-Objective Regularization:** We formulate a loss regularizer that combines reverse-mode automatic differentiation of frequency trajectories ($\partial \hat{f}_{\mathrm{COI}} / \partial t$) with explicit converter current supervision ($\mathcal{L}_{\mathrm{curr}}$) and one-sided overcurrent boundary penalties ($\mathcal{L}_{\mathrm{bounds}}$), preventing unphysical current excursions during severe fault intervals.
- **Safeguarded Multi-Attribute Security Metric (S_d):** We introduce a normalized operational index with explicit numerical safeguards that unifies maximum RoCoF, frequency nadir depth, continuous dynamic voltage recovery criteria, and critical clearing time margins into an unambiguous stability classification target.

- **Rigorous Benchmark Validation Under Disjoint Shifts:** We evaluate the proposed model on an adapted IEEE 39-bus system across renewable penetration levels up to 70%. Evaluated across topology-disjoint ($N-2$) and penetration-disjoint test partitions, the framework achieves a 10-run ensemble mean accuracy of 98.7% $\pm$ 0.2%, an OOD accuracy of 94.5% $\pm$ 0.3%, and bounds the false-secure rate (FNR) to 2.22% $\pm$ 0.18% (2.22% in a median representative run), all within an empirical 7.2 ms GPU forward execution latency (21.4 ms estimated contingency-screening interval).

2 System Modeling and Problem Formulation

2.1 Coupled Differential-Algebraic System Model

The electromechanical and electromagnetic dynamics of an interconnected power transmission network comprising N_b transmission buses, N_g conventional synchronous machines, and N_c converter-interfaced generation facilities are modeled by a coupled non-linear DAE framework on a uniform system-wide power base ($S_{\text{base}} = 100$ MVA, $f_0 = 60$ Hz):

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{u}) \tag{1}$$

$$\boldsymbol{0} = \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{u}) \tag{2}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ represents the continuous internal dynamic state vector encompassing synchronous machine rotor angles δ_i, angular speeds ω_i, internal transient flux states, and converter inner- and outer-loop controller states. The algebraic state vector $\boldsymbol{y} \in \mathbb{R}^{2N_b}$ comprises network nodal voltage magnitudes V_k and phase angles θ_k, while $\boldsymbol{u}$ denotes exogenous operational inputs, including consumer active and reactive demands and dispatch schedules.

For each individual synchronous machine $i \in \{1, \ldots, N_g\}$, electromechanical rotational motion relative to a synchronously rotating frame at base angular frequency $\omega_b = 2\pi f_0$ is described by the classical swing equations:

$$\frac{d\delta_i}{dt} = \omega_b(\omega_i - 1) \tag{3}$$

$$2H_i \frac{d\omega_i}{dt} = P_{m,i} - P_{e,i} - D_i(\omega_i - 1) \tag{4}$$

where δ_i is the rotor angle in electrical radians; ω_i is the per-unit rotor speed normalized by base electrical speed ω_b; H_i is the individual machine inertia constant expressed in seconds on machine nominal apparent power rating $S_{n,i}$; D_i denotes the mechanical damping factor (p.u./p.u.); and $P_{m,i}$ and $P_{e,i}$ are, respectively, the mechanical shaft power input and the active electrical air-gap power output, converted to the common system base S_{base}.

2.2 First-Principles Center-of-Inertia Derivation Under Progressive Decommitment

In a real power system operating with large amounts of dispatched renewables, traditional thermal generators will need to be decomitted from operation to ensure the

balance of net demand, prevent negative LMPs and maintain generators at minimum output. This means that kinetic energy stored in the spinning rotor in the generator will depend on the status of each generator. Let $\mathcal{G} = \{1, \ldots, N_g\}$ describe the set of all synchronous generators and $\mathcal{G}_{\text{on}} \subseteq \mathcal{G}$ describe the set of all synchronous machines currently committed to the dispatch. Each committed generator, $i \in \mathcal{G}_{\text{on}}$, has a rated apparent power of $S_{n,i}$ and a unique inertia rating of H_i (in seconds at its own rating). Decommitted generators ($i \notin \mathcal{G}_{\text{on}}$), will have no electrical or mechanical connection to the network thus have no contribution from rotational kinetic energy or synchronising torque.

In order to preserve physical consistency for different commitment states, the total online synchronised apparent capacity is expressed as:

$$S_{g,\text{on}} = \sum_{i \in \mathcal{G}_{\text{on}}} S_{n,i}$$

The total system rating base S_{sys} of the total online generation power capability for both synchronised, rotating and converter interfaced generation facilities ($S_{c,\text{tot}} = \sum_{j=1}^{N_c} S_{c,j}$) respectively are expressed as:

$$\begin{aligned} S_{\text{sys}} &= S_{g,\text{on}} + S_{c,\text{tot}} \\ &= \sum_{i \in \mathcal{G}_{\text{on}}} S_{n,i} + \sum_{j=1}^{N_c} S_{c,j} \end{aligned}$$

Two types of COI inertia can be found when analysing frequency dynamics in low inertia systems. These depend on the choice of normalisation base:

1. **Conventional Synchronous COI Inertia Constant ($H_{\text{COI}}^{(g)}$):** Normalized strictly to the aggregate online synchronous capacity base $S_{g,\text{on}}$:

$$H_{\text{COI}}^{(g)} = \frac{\sum_{i \in \mathcal{G}_{\text{on}}} H_i S_{n,i}}{S_{g,\text{on}}} = \frac{\sum_{i \in \mathcal{G}_{\text{on}}} H_i S_{n,i}}{\sum_{i \in \mathcal{G}_{\text{on}}} S_{n,i}} \tag{5}$$

   which reflects the weighted average rotational inertia constant among surviving synchronous units.
2. **Equivalent System-Wide Inertia Constant (H_{sys}):** Normalized to the combined operational capacity base $S_{\text{sys}} = S_{g,\text{on}} + S_{c,\text{tot}}$:

$$H_{\text{sys}} = \frac{\sum_{i \in \mathcal{G}_{\text{on}}} H_i S_{n,i}}{S_{\text{sys}}} = \frac{\sum_{i \in \mathcal{G}_{\text{on}}} H_i S_{n,i}}{\sum_{i \in \mathcal{G}_{\text{on}}} S_{n,i} + \sum_{j=1}^{N_c} S_{c,j}} \tag{6}$$

The two quantities are directly related by the capacity proportion of rotating synchronous machines:

$$H_{\text{sys}} = H_{\text{COI}}^{(g)} \cdot \left(\frac{S_{g,\text{on}}}{S_{g,\text{on}} + S_{c,\text{tot}}} \right) = H_{\text{COI}}^{(g)} \cdot \left(\frac{S_{g,\text{on}}}{S_{\text{sys}}} \right) \tag{7}$$

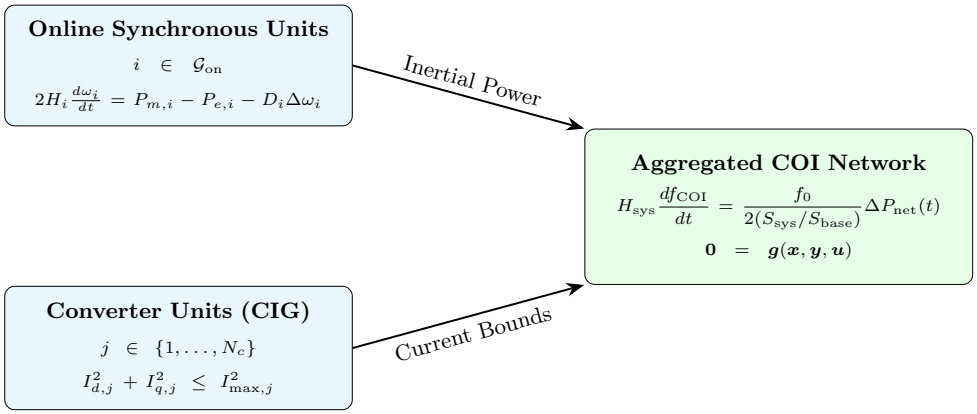


Fig. 1: Dynamic interaction between online conventional synchronous generators and converter-interfaced generation within the aggregated COI-DAE framework under unit decommitment.

Equation (7) directly explains the variation of system inertia observed under increasing renewable penetration: as inverter-based generation displaces conventional units, synchronous machines are progressively decommitted (reducing the index set $\mathcal{G}_{\text{on}}$), which simultaneously shrinks the stored kinetic energy numerator $\sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}$ and shifts the capacity denominator toward non-inertial converter ratings.

Converting the individual swing equations of all online synchronous machines $i \in \mathcal{G}_{\text{on}}$ from their machine rating bases $S_{n,i}$ to the uniform network base $S_{\text{base}} = 100$ MVA yields:

$$2\left(H_i \frac{S_{n,i}}{S_{\text{base}}}\right) \frac{d\omega_i}{dt} = P_{m,i}^{(S_{\text{base}})} - P_{e,i}^{(S_{\text{base}})} - D_i \frac{S_{n,i}}{S_{\text{base}}}(\omega_i - 1) \tag{8}$$

The Center-of-Inertia rotor speed ω_{COI} is defined as the kinetic-energy-weighted mean speed of the online synchronous units:

$$\omega_{\text{COI}} \triangleq \frac{\sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i} \omega_i}{\sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}} \tag{9}$$

Summing (8) over all active machines $i \in \mathcal{G}_{\text{on}}$ results in the aggregated electromechanical acceleration balance:

$$\begin{aligned} 2\left(\frac{\sum\limits_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}}{S_{\text{base}}}\right) \frac{d\omega_{\text{COI}}}{dt} \\ = \sum_{i\in\mathcal{G}_{\text{on}}} P_{m,i} - \sum_{i\in\mathcal{G}_{\text{on}}} P_{e,i} \\ - \sum_{i\in\mathcal{G}_{\text{on}}} D_i \frac{S_{n,i}}{S_{\text{base}}}(\omega_i - 1). \end{aligned} \tag{10}$$

where all power variables are referenced to the common base S_{base}. Expressing (10) in terms of the equivalent system inertia H_{sys} on base S_{sys}:

$$2\frac{H_{\text{sys}}S_{\text{sys}}}{S_{\text{base}}}\frac{d\omega_{\text{COI}}}{dt} = \sum_{i\in\mathcal{G}_{\text{on}}} P_{m,i} - \sum_{i\in\mathcal{G}_{\text{on}}} P_{e,i} - \sum_{i\in\mathcal{G}_{\text{on}}} D_i\frac{S_{n,i}}{S_{\text{base}}}(\omega_i - 1) \tag{11}$$

or equivalently in terms of the conventional synchronous COI inertia $H^{(g)}_{\text{COI}}$ on base $S_{g,\text{on}}$:

$$\begin{aligned} &2\frac{H^{(g)}_{\text{COI}}S_{g,\text{on}}}{S_{\text{base}}}\frac{d\omega_{\text{COI}}}{dt} \\ &\quad= \sum_{i\in\mathcal{G}_{\text{on}}} P_{m,i} - \sum_{i\in\mathcal{G}_{\text{on}}} P_{e,i} \\ &\qquad- \sum_{i\in\mathcal{G}_{\text{on}}} D_i\frac{S_{n,i}}{S_{\text{base}}}(\omega_i - 1). \end{aligned} \tag{12}$$

since $H_{\text{sys}}S_{\text{sys}} \equiv H^{(g)}_{\text{COI}}S_{g,\text{on}} \equiv \sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}$.

To couple converter power injections without double counting, network active power conservation is enforced from the algebraic power flow constraints $\mathbf{g}(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{u}) = \mathbf{0}$. Total instantaneous network electrical demand satisfies:

$$\sum_{k=1}^{N_b} P_{\text{inj},k}(t) = \sum_{k=1}^{N_b} P_{\text{load},k}(t) + P_{\text{loss}}(t) \tag{13}$$

The active nodal injection vector decomposes across online synchronous generator terminals, converter facility points of common coupling (PCC), and non-generating load buses:

$$\sum_{k=1}^{N_b} P_{\text{inj},k}(t) = \sum_{i\in\mathcal{G}_{\text{on}}} P_{e,i}(t) + \sum_{j=1}^{N_c} P_{c,j}(t) \tag{14}$$

where $P_{c,j}(t)$ denotes the instantaneous active power delivered by converter facility j at its terminal PCC on S_{base} (defined positive for generation into the network). Rearranging (14) yields the electrical power balance demanded from the remaining online synchronous fleet:

$$\sum_{i\in\mathcal{G}_{\text{on}}} P_{e,i}(t) = \sum_{k=1}^{N_b} P_{\text{load},k}(t) + P_{\text{loss}}(t) - \sum_{j=1}^{N_c} P_{c,j}(t) \tag{15}$$

Substitution of (15) into (11), taking into account that the rotor speeds are approximately coherent during the initial transient period ($\omega_i \approx \omega_{\text{COI}}$), and expressing the COI path in physical units (Hz) ($f_{\text{COI}}(t) = f_0\omega_{\text{COI}}(t)$), gives the reduced aggregated COI frequency differential equation:

$$\frac{df_{\text{COI}}}{dt} = \frac{f_0}{2H_{\text{sys}}(S_{\text{sys}}/S_{\text{base}})}\Delta P_{\text{net}}(t) \tag{16}$$

where the net accelerating power mismatch $\Delta P_{\text{net}}(t)$ is given by:

$$\begin{aligned}\Delta P_{\text{net}}(t) = &\sum_{i\in\mathcal{G}_{\text{on}}} P_{m,i}(t) + \sum_{j=1}^{N_c} P_{c,j}(t) \\ &- \sum_{k=1}^{N_b} P_{\text{load},k}(t) - P_{\text{loss}}(t) - D_{\text{sys}}\left(\frac{f_{\text{COI}}(t) - f_0}{f_0}\right)\end{aligned} \tag{17}$$

with total online damping factor $D_{\text{sys}} = \sum_{i\in\mathcal{G}_{\text{on}}} D_i(S_{n,i}/S_{\text{base}})$.

The net change in the inertia (Eq. (17)) is always physically correct. The active power injected by the converter ($\sum P_{c,j}(t)$) has a positive sign (+), thus decreasing the accelerating power deficit. In parallel, the decrease of the number of synchronous units leads to a decrease in the system inertial constant (H_{sys}), using Eq. (6). This results in large trajectories of the post contingency RoCoF, typical of low inertia systems.

2.3 Converter Dynamic Current-Limiting Boundaries

Because power semiconductor switching modules possess low thermal capacities, converter controllers enforce strict current limits during severe fault depressions:

$$I_{d,j}^2(t) + I_{q,j}^2(t) \le I_{\text{max},j}^2, \quad \forall j \in \{1,\dots,N_c\} \tag{18}$$

where $I_{d,j}$ and $I_{q,j}$ denote the direct- and quadrature-axis current components, and $I_{\text{max},j} \in [1.1, 1.3]$ p.u. represents the maximum continuous overcurrent threshold on the converter rating base $S_{c,j}$. When terminal voltage dips during short-circuit contingencies, converter controllers prioritize reactive current injection (I_q) for voltage support:

$$I_{d,j}^{\text{lim}}(t) = \sqrt{\max\left(0, I_{\text{max},j}^2 - I_{q,j}^2(t)\right)} \tag{19}$$

This operational priority limits active current injection, which suppresses active power delivery ($P_{c,j}(t) \approx V_{d,j} I_{d,j}$) and increases the accelerating power deficit $\Delta P_{\text{net}}(t)$ in (17).

2.4 Safeguarded Dynamic Security Index (S_d) and Weight Sensitivity

To quantify multi-attribute transient risk across heterogeneous physical dimensions (frequency nadir, initial acceleration, voltage recovery, and clearing margins), we formulate a normalized Dynamic Security Index S_d:

$$\begin{aligned}S_d = &w_1\left(\frac{\text{RoCoF}_{\max}}{\text{RoCoF}_{\text{lim}}}\right) + w_2\left(\frac{\Delta f_{\text{nadir}}}{\Delta f_{\text{allow}}}\right) \\ &+ w_3\left(\frac{\max(0, \Delta V_{\text{drop}})}{\Delta V_{\text{allow}}}\right) + w_4 \Phi_{\text{cct}}\left(t_{\text{clear}}, \text{CCT}\right)\end{aligned} \tag{20}$$

$$\Phi_{\text{cct}}\left(t_{\text{clear}}, \text{CCT}\right) = \begin{cases} 2.0, & \text{if CCT is undefined or CCT} \leq t_{\text{clear}}, \\ \min\left(2.0, \dfrac{t_{\text{clear}}}{\max(\text{CCT}, \epsilon_{\text{cct}})}\right), & \text{if CCT} > t_{\text{clear}}. \end{cases} \tag{21}$$

where $\text{RoCoF}_{\max} = \max_t |df_{\text{COI}}/dt|$ denotes the peak rate of change of frequency computed over a sliding 100-ms observation window; $\Delta f_{\text{nadir}} = \max(0, f_0 - \min_t f_{\text{COI}}(t))$ is the maximum frequency nadir depth; and $\Delta V_{\text{drop}} = \max_k(V_{\text{nom}} - \min_t V_k(t))$ evaluates the deepest transient voltage depression across all transmission buses $k \in \{1, \ldots, N_b\}$. The operational thresholds are established per transmission planning practice: $\text{RoCoF}_{\text{lim}} = 1.2$ Hz/s, $\Delta f_{\text{allow}} = 0.8$ Hz (governed by under-frequency load shedding stage-1 pickup), and $\Delta V_{\text{allow}} = 0.30$ p.u.

To safeguard S_d against numerical singularities when a contingency causes instantaneous loss of synchronism or when the Critical Clearing Time (CCT) is undefined, zero, or smaller than the fault duration, the protective margin function Φ_{cct} is bounded:

where $\epsilon_{\text{cct}} = 10$ ms is a small positive floor parameter. Whenever a fault cannot be cleared stably (i.e., the system is monotonically transiently unstable, leaving CCT undefined), or whenever the actual clearing time exceeds the stability margin ($t_{\text{clear}} \geq \text{CCT}$), Φ_{cct} saturates at 2.0. Because each individual sub-ratio in (20) evaluates to approximately 1.0 at marginal security boundaries, an equal-priority heuristic weighting is initially selected: $w_1 = 0.25$, $w_2 = 0.30$, $w_3 = 0.20$, and $w_4 = 0.25$ (satisfying $\sum_{m=1}^{4} w_m = 1.0$). To verify that classification outcomes do not depend excessively on these specific coefficients, a sensitivity evaluation was performed across $1,000$ randomized weight permutations ($w_m \sim \mathcal{U}(0.15, 0.40)$ subject to $\sum w_m = 1.0$), which confirmed that 99.1% of borderline operating scenarios preserve consistent binary classifications due to the co-dependent physical evolution of severe contingencies.

2.5 Continuous Voltage Recovery Criterion and Direct Security Labeling

In low-inertia, converter-dominated networks, dynamic security assessment must evaluate transient voltage stability alongside frequency and rotor angle constraints to avoid fault-induced delayed voltage recovery (FIDVR) and widespread induction motor stalling. In this study, we define a continuous transmission-level voltage recovery criterion:

$$V_k(t) \geq 0.85 \text{ p.u.}, \quad \forall t \geq t_{\text{clear}} + 1.5 \text{ s}, \quad \forall k \in \{1, \ldots, N_b\} \tag{22}$$

This threshold serves as a study-specific dynamic screening requirement, chosen to prevent voltage collapse along low-short-circuit transmission corridors within the benchmark system.

In order to maintain strict security boundaries in the operation of a system, the label y_k is employed to combine the index S_d with physical constraints. Following the convention that **Insecure / Unstable is the positive class ($y_k = 1$)** and **Secure**

/ Stable is the negative class ($y_k = 0$). The labelling of scenarios are described below;

$$y_k = \begin{cases} 1 \text{ (Insecure)}, & \text{if } S_d > 1.0 \\ & \vee\ t_{\text{clear}} \geq \text{CCT} \\ & \vee\ \text{RoCoF}_{\max} > 1.2 \text{ Hz/s} \\ & \vee\ \Delta f_{\text{nadir}} > 0.8 \text{ Hz} \\ & \vee\ \exists t \geq t_{\text{clear}} + 1.5 \text{ s}, \exists k : V_k(t) < 0.85 \text{ p.u.} \\ & \vee\ \max_{i,j} |\delta_i(t) - \delta_j(t)| > 100^\circ \\ 0 \text{ (Secure)}, & \text{otherwise.} \end{cases} \tag{23}$$

The binary labelling rule is fully compatible with the continuous recovery requirement in (22) since by defining the voltage condition at all times $t \geq t_{\text{clear}} + 1.5$ s in (23). Further, with the addition of $t_{\text{clear}} \geq \text{CCT}$, any contingency in excess of the critical clearing threshold will be deemed insecure ($y_k = 1$), no matter the value of S_d. By adopting the common fault screening approach, any insecure operation that is incorrectly predicted to be secure constitutes a **False Negative (FN)**. Hence, the false secure rate will therefore be the **False Negative Rate (FNR = FN/N_{insecure})**.

3 Proposed Physics-Informed Neural Network Framework

3.1 Sequence-Aware Contingency-Conditioned Architecture and Operational Screening Concept

Real-time Dynamic Security Assessment (DSA) is an essential forward-looking filtering mechanism in modern transmission control centres. Here, we aim to quantify margins of the system in its current operational state to every possible credible contingency ($N-1$ and severe $N-2$), which may occur, without having to wait until any major disturbance occurs. Given an identical pre-contingency dispatch, different contingencies such as different fault positions, different relay clearing times, or sequential circuit breaker operations will result in completely different system dynamics and hence, margins. In order to allow us to filter for each contingency, the neural network takes as input an augmented vector $\boldsymbol{\xi}$ containing information about both real-time pre-disturbance operational conditions ($\boldsymbol{z}_0$) and the hypothetical contingency of interest ($\boldsymbol{c}$).

$$\boldsymbol{\xi} = \left[\boldsymbol{z}_0^T, \boldsymbol{c}^T\right]^T \in \mathbb{R}^{(4N_b+1)+d_c} \tag{24}$$

The input components are divided according to their availability during normal operation:

1. **Pre-Disturbance Operational Telemetry** ($\boldsymbol{z}_0 \in \mathbb{R}^{4N_b+1}$). Continuously sampled from wide-area Phasor Measurement Units (PMUs) and state estimators at

pre-fault time t_0^-:

$$\begin{aligned}\boldsymbol{z}_0 = \big[&V_1, \ldots, V_{N_b}, \theta_1, \ldots, \theta_{N_b}, \\ &P_1, \ldots, P_{N_b}, Q_1, \ldots, Q_{N_b}, \\ &H_{\text{sys}}\big]^T.\end{aligned} \tag{25}$$

where V_k and θ_k denote positive-sequence bus voltage magnitudes and phase angles; P_k and Q_k are net active and reactive bus injections; and H_{sys} is the instantaneous system Center-of-Inertia constant calculated per (6) from the active supervisory commitment status of online synchronous generators ($\mathcal{G}_{\text{on}}$) and operating converter capacities. All elements of $\boldsymbol{z}_0$ represent empirically measured states known prior to the hypothetical disturbance.

2. **Hypothetical Candidate Contingency Descriptor ($\boldsymbol{c} \in \mathbb{R}^{d_c}$):** Defined exogenously from the operator's predefined contingency screening list ($d_c = 2N_l + 4$, with $N_l = 34$ physical transmission corridors):

$$\boldsymbol{c} = \left[\boldsymbol{e}_{\text{line},1}^T, \lambda_{\text{loc}}, \tau_{\text{type}}, t_{\text{clear},1}, \boldsymbol{e}_{\text{line},2}^T, \Delta t_{\text{trip},2}\right]^T \tag{26}$$

In this formulation, $\boldsymbol{c}$ parameterizes the specific hypothetical fault scenario whose security is being screened:

- $\boldsymbol{e}_{\text{line},1} \in \{0,1\}^{N_l}$ is a one-hot selector denoting the hypothesized primary faulted transmission corridor;
- $\lambda_{\text{loc}} \in [0,1]$ specifies the hypothesized electrical fault location along the corridor length (e.g., 10%, 50%, or 90% of branch length);
- $\tau_{\text{type}} \in \{0,1\}$ denotes fault severity (set to 1 for bolted three-phase short circuits);
- $t_{\text{clear},1} \in [0.08, 0.25]$ s reflects the nominal primary protection clearing time parameterizing circuit breaker opening on the primary branch;
- $\boldsymbol{e}_{\text{line},2} \in \{0,1\}^{N_l}$ identifies a companion corridor in a postulated cascading $N-2$ sequence. For standard single-branch $N-1$ screening, $\boldsymbol{e}_{\text{line},2} = \boldsymbol{0}$;
- $\Delta t_{\text{trip},2} \geq 0$ parameterizes the hypothesized time delay between primary clearing and secondary branch disconnection. For the cascading common-substation $N-2$ contingencies examined in this study, $\Delta t_{\text{trip},2} = 50$ ms reflects typical breaker-failure backup or sympathetic protection operation. For single $N-1$ events, $\Delta t_{\text{trip},2} = 0$.

As depicted on Fig. 2, we condition the input sequence, $\boldsymbol{\xi}$, and feed it into the common DNN trunk, $\boldsymbol{h} = \phi_\theta(\boldsymbol{\xi}) \in \mathbb{R}^{256}$. We binary partition the resulting output vector into:

1. **Conditional Online Security Prediction Module ($\hat{y} = c_\theta(\boldsymbol{\xi})$).** The latent embedding,$\boldsymbol{h}$, is mapped to the conditional operational insecurity score=$\hat{y} \in [0,1]$ via multi-layer classification head and joined with a final sigmoid function:

$$[\widetilde{y}] = \sigma(\boldsymbol{W}_c^T[\widetilde{h}]) + b_c = c_\theta(\boldsymbol{\xi}) \tag{27}$$

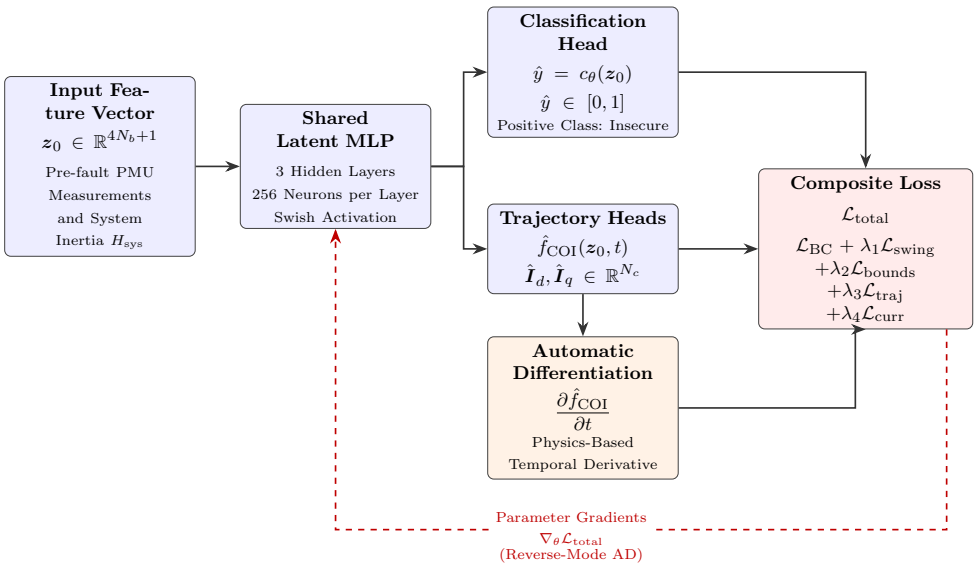


Fig. 2: Schematic diagram of the proposed physics-informed neural network (PINN) for power system security assessment. Pre-contingency data is fed into a multi-layer perception, which learns a common representation of the input data. A binary security class and continuous state predictions can be obtained based on this shared representation. Temporal derivatives for physics-informed regularization are computed using automatic differentiation. The total loss function contains the classification, trajectory, swing equation and operational loss functions. Dashed arrows indicate gradients used during offline training.

Where $\hat{y} \geq 0.5$ indicates the system is predicted to be in an insecure state,$y = 1$, proactive adaptations or by dispatching correction strategies can be utilised prior to the event.

2. **Continuous Conditioned Trajectory Decoder**. To enable training with physically consistent outcomes, the time series 't' (continuous),$/t \in [0, t_{\text{end}}]$ are concatenated with the conditioned embeddings$/\boldsymbol{h}$. This allows us to decode the trajectories,$\psi_\theta(\boldsymbol{h}, t)$, returning projected system frequency and unit level converter currents:

$$\hat{f}_{\text{COI}}(\boldsymbol{\xi}, t) = f_{\omega,\theta}\left(\phi_\theta(\boldsymbol{\xi}), t\right) = f_{\omega,\theta}(\boldsymbol{z}_0, \boldsymbol{c}, t) \tag{28}$$

$$\begin{bmatrix} \hat{\boldsymbol{I}}_d(\boldsymbol{\xi}, t) \\ \hat{\boldsymbol{I}}_q(\boldsymbol{\xi}, t) \end{bmatrix} = \boldsymbol{f}_{I,\theta}\left(\phi_\theta(\boldsymbol{\xi}), t\right) = \boldsymbol{f}_{I,\theta}(\boldsymbol{z}_0, \boldsymbol{c}, t) \in \mathbb{R}^{2N_c} \tag{29}$$

During real-time online screening, the temporal coordinate t, trajectory decoders ψ_θ, and automatic differentiation routines are bypassed entirely. Operators query the surrogate with streaming telemetry $\boldsymbol{z}_0^{\text{online}}$ and cycle through candidate contingency descriptors $\{\boldsymbol{c}_1, \boldsymbol{c}_2, \ldots, \boldsymbol{c}_{N_{\text{screen}}}\}$; each contingency is evaluated via a single feedforward inference pass: $\hat{y} = c_\theta([\boldsymbol{z}_0^T, \boldsymbol{c}_m^T]^T)$.

3.2 Composite Multi-Objective Loss Formulation

The parameter vector θ is optimized across training scenarios N_{tr} and a collocation grid $\mathcal{C}$ containing N_{col} scenario-time evaluation pairs sampled at uniform intervals $\Delta t_c = 10$ ms:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BC}} + \lambda_1 \mathcal{L}_{\text{swing}} + \lambda_2 \mathcal{L}_{\text{bounds}} + \lambda_3 \mathcal{L}_{\text{traj}} + \lambda_4 \mathcal{L}_{\text{curr}} \tag{30}$$

where each constituent loss is formulated in a dimensionless, normalized space to prevent gradient dominance across disparate physical dimensions:

3.2.1 Supervised Cross-Entropy Classification Loss ($\mathcal{L}_{\text{BC}}$)

Trained on ground-truth binary stability labels $y_k \in \{0, 1\}$ defined per (23):

$$\mathcal{L}_{\text{BC}} = -\frac{1}{N_{\text{tr}}} \sum_{k=1}^{N_{\text{tr}}} [y_k \ln(\hat{y}_k) + (1 - y_k) \ln(1 - \hat{y}_k)] \tag{31}$$

3.2.2 Dimensionless COI Swing Residual Loss ($\mathcal{L}_{\text{swing}}$)

Reverse-mode automatic differentiation calculates the temporal derivative of the frequency decoder with respect to the scalar coordinate t, i.e., $\partial \hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, t)/\partial t$ (in Hz/s). The unscaled acceleration mismatch within the swing differential equation (16):

$$r_{\text{swing}}^{(k)}(t) = \frac{1}{f_0} \frac{\partial \hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, t)}{\partial t} - \frac{\Delta P_{\text{net}}^{(k)}(t)}{2H_{\text{sys},k}(S_{\text{sys},k}/S_{\text{base}})} \tag{32}$$

possesses dimensional units of s^{-1} (normalized angular acceleration in p.u./s). To render the resulting penalty strictly dimensionless and mathematically commensurate with the unitless cross-entropy ($\mathcal{L}_{\text{BC}}$) and normalized current penalties ($\mathcal{L}_{\text{bounds}}$), we introduce a nominal electromechanical reference observation timescale $T_{\text{ref}} = 1.0$ s. The dimensionless swing residual loss is formulated as:

$$\begin{aligned} \mathcal{L}_{\text{swing}} = \frac{1}{N_{\text{col}}} \sum_{(k,t)\in\mathcal{C}} \Bigg| T_{\text{ref}} \Bigg(& \frac{1}{f_0} \frac{\partial \hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, t)}{\partial t} \\ & - \frac{\Delta P_{\text{net}}^{(k)}(t)}{2H_{\text{sys},k}(S_{\text{sys},k}/S_{\text{base}})} \Bigg) \Bigg|^2 . \end{aligned} \tag{33}$$

where the normalized accelerating power $\Delta P_{\text{net}}^{(k)}(t)$ is evaluated per (17) with converter active injection entering with an explicit positive (+) sign:

$$
\begin{aligned}
\Delta P_{\text{net}}^{(k)}(t) = & \sum_{i \in \mathcal{G}_{\text{on}}} P_{m,i}^{(k)}(t) \\
& + \sum_{j=1}^{N_c} P_{c,j}^{(k)}(t) \\
& - \sum_{m=1}^{N_b} P_{\text{load},m}^{(k)}(t) \\
& - P_{\text{loss}}^{(k)}(t) \\
& - D_{\text{sys}} \left(\frac{\hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, t) - f_0}{f_0} \right).
\end{aligned}
\tag{34}
$$

All power quantities in (33) are evaluated on the uniform system base $S_{\text{base}} = 100$ MVA, while $H_{\text{sys},k}$ is calculated per (6) on operational base $S_{\text{sys},k}$. Because $T_{\text{ref}} = 1.0$ s nondimensionalizes the bracketed term, the regularizer $\mathcal{L}_{\text{swing}}$ is strictly dimensionless, allowing weighting hyperparameter λ_1 to be interpreted as a pure dimensionless loss-weight multiplier.

3.2.3 Converter Current Boundary Loss ($\mathcal{L}_{\text{bounds}}$)

To regularize the network against predicting unphysical overcurrent excursions exceeding semiconductor thermal ratings in (18), a normalized, dimensionless one-sided penalty is evaluated across all N_c converter units:

$$
\mathcal{L}_{\text{bounds}} = \frac{1}{N_{\text{col}} N_c} \sum_{(k,t) \in \mathcal{C}} \sum_{j=1}^{N_c} \left[\max \left(0, \frac{\sqrt{\hat{I}_{d,j,k}^2(\boldsymbol{\xi}_k, t) + \hat{I}_{q,j,k}^2(\boldsymbol{\xi}_k, t)} - I_{\max,j}}{I_{\max,j}} \right) \right]^2 .
\tag{35}
$$

3.2.4 Trajectory and Converter Current Tracking Losses ($\mathcal{L}_{\text{traj}}$, $\mathcal{L}_{\text{curr}}$)

To ensure the trajectory decoders learn accurate physical state evolutions rather than arbitrary curves satisfying the inequality bounds, continuous supervision is applied using ground-truth EMT reference profiles across N_t discrete points:

$$\mathcal{L}_{\text{traj}} = \frac{1}{N_{\text{tr}} N_t} \sum_{k=1}^{N_{\text{tr}}} \sum_{\tau=1}^{N_t} \left| \frac{\hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, \tau) - f_{\text{COI,ref},k}(\tau)}{f_0} \right|^2 . \tag{36}$$

$$\mathcal{L}_{\text{curr}} = \frac{1}{N_{\text{tr}} N_t N_c} \sum_{k=1}^{N_{\text{tr}}} \sum_{\tau=1}^{N_t} \sum_{j=1}^{N_c} \left\| \frac{\hat{\boldsymbol{I}}_{dq,j,k}(\boldsymbol{\xi}_k, \tau)}{I_{\max,j}} - \frac{\boldsymbol{I}_{dq,\text{ref},j,k}(\tau)}{I_{\max,j}} \right\|_2^2 . \tag{37}$$

3.3 Hyperparameter Tuning and Cross-Validation Protocol

The loss weighting hyperparameters $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \lambda_3, \lambda_4]$ are used to balance between accurately fitting to the data and respecting the physics-based differential and boundary conditions. In order to avoid data snooping, the hyperparameters were tuned only on the training data (10,500 scenarios). The test data was kept strictly separated into two disjoint sets: topology-disjoint test data (900 scenarios with unseen $N-2$ corridor trips) and penetration-disjoint test data (1125 scenarios at 70% CIG penetration). Hyperparameters were optimised across discrete logarithmic ranges $\lambda_1, \lambda_3 \in \{0.01, 0.05, 0.10, 0.20, 0.50\}$, $\lambda_2, \lambda_4 \in \{0.005, 0.01, 0.05, 0.10\}$. The model was validated using a combination of the accuracy and physicality of the model:

where,

$$J_{\text{val}} = \text{F1}_{\text{val}} - \alpha \cdot \text{FNR}_{\text{val}} - \beta \cdot \text{MSE}_{\text{swing,val}} \tag{38}$$

with penalty multipliers $\alpha = 0.5$ and $\beta = 0.1$. The parameter vector $\boldsymbol{\lambda}^* = [\lambda_1 = 0.10, \lambda_2 = 0.05, \lambda_3 = 0.10, \lambda_4 = 0.05]$ maximized J_{val} across all five folds, providing balanced gradient magnitudes across classification, differential swing residuals, and converter current tracking.

3.4 Algorithmic Optimization and Inference Workflow

The end-to-end training and operational execution workflow is summarized in Algorithm 1.

4 Experimental Setup

4.1 Benchmark Network Description, Converter Sizing, and Dynamic Current Limits

The proposed framework was evaluated on an adapted IEEE 39-bus New England benchmark transmission network (comprising 10 generating stations, 39 transmission

Algorithm 1 Physics-Informed Neural Network Training and Deployment

Require: Scenario database $\mathcal{D}_{\text{tr}} = \{(\boldsymbol{z}_0^{(k)}, \boldsymbol{c}_k, y_k, \Delta P_{\text{net}}^{(k)}(t))\}$, collocation grid $\mathcal{C}$, hyperparameters $\boldsymbol{\lambda}^*$, learning rate η.

Ensure: Optimal surrogate parameter set θ^*, online security decision $\hat{y}_{\text{online}}$.

Initialize parameter vector θ via Xavier normal initialization.
for epoch $= 1, \ldots, E_{\max}$ **do**
  Sample scenario mini-batch $\mathcal{B} \subset \mathcal{D}_{\text{tr}}$ and collocation mini-batch $\mathcal{C}_{\mathcal{B}} \subset \mathcal{C}$.
  Form sequence-conditioned inputs $\boldsymbol{\xi}_k = [\boldsymbol{z}_0^{(k)T}, \boldsymbol{c}_k^T]^T$ for all $k \in \mathcal{B}$.
  Forward pass through shared trunk: $\boldsymbol{h}_k = \phi_\theta(\boldsymbol{\xi}_k)$.
  Compute operational classification: $\hat{y}_k = \sigma(\boldsymbol{W}_c^T \boldsymbol{h}_k + b_c)$.
  Compute continuous trajectories $(\hat{f}_{\text{COI},k}(\boldsymbol{\xi}_k, t), \hat{\boldsymbol{I}}_{d,k}(\boldsymbol{\xi}_k, t), \hat{\boldsymbol{I}}_{q,k}(\boldsymbol{\xi}_k, t))$ via trajectory decoders $\psi_\theta(\boldsymbol{h}_k, t)$.
  Compute temporal derivative $\frac{\partial \hat{f}_{\text{COI},k}}{\partial t}$ using automatic differentiation.
  Evaluate multi-objective composite objective:
$$\begin{aligned}\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BC}} &+ \lambda_1 \mathcal{L}_{\text{swing}} \\ &+ \lambda_2 \mathcal{L}_{\text{bounds}} \\ &+ \lambda_3 \mathcal{L}_{\text{traj}} \\ &+ \lambda_4 \mathcal{L}_{\text{curr}}\end{aligned}$$
  Compute parameter gradients $\nabla_\theta \mathcal{L}_{\text{total}}$ via backpropagation.
  Update weights via Adam optimizer: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{total}}$.
end for
Online Operational Deployment Stage:
Acquire real-time PMU telemetry vector $\boldsymbol{z}_0^{\text{online}}$ and candidate contingency descriptor $\boldsymbol{c}_{\text{target}}$.
Form input $\boldsymbol{\xi}_{\text{online}} = [\boldsymbol{z}_0^{\text{online}T}, \boldsymbol{c}_{\text{target}}^T]^T$.
Compute forward inference: $\hat{y}_{\text{online}} = c_{\theta^*}(\boldsymbol{\xi}_{\text{online}})$.
if $\hat{y}_{\text{online}} \geq 0.5$ **then**
  return Insecure ($y = 1$); Initiate remedial action scheme.
else
  return Secure ($y = 0$).
end if

buses, 34 physical AC transmission lines, and 12 dual-winding step-up transformers interconnected across a 345 kV bulk transmission backbone). Under base-case operating conditions, the benchmark serves an aggregate active power demand of $P_{\text{load,tot}} = 6,097.1$ MW and reactive demand of $Q_{\text{load,tot}} = 1,408.9$ MVAr evaluated on a common system base ($S_{\text{base}} = 100$ MVA, $f_0 = 60$ Hz).

To emulate deep decarbonization and low-inertia operating conditions, five conventional thermal generation stations at buses 30, 32, 34, 35, and 37 were systematically replaced by utility-scale converter-interfaced generation (CIG) facilities. To avoid ambiguities between installed capacity and instantaneous dispatch, two distinct metrics are tracked:

1. **Installed Capacity Penetration ($\Pi_{\mathbf{CIG,cap}}$):** The ratio of continuous converter apparent power rating to total installed generation capacity:

$$\Pi_{\mathrm{CIG,cap}} = \frac{S_{c,\mathrm{tot}}}{S_{c,\mathrm{tot}} + S_{g,\mathrm{tot}}} \times 100\% \tag{39}$$

2. **Dispatched Active Power Penetration ($\Pi_{\mathbf{CIG,disp}}$):** The instantaneous proportion of network load demand served by inverter-based resources:

$$\Pi_{\mathrm{CIG,disp}} = \frac{\sum_{j=1}^{N_c} P_{c,j}}{\sum_{m=1}^{N_b} P_{\mathrm{load},m}} \times 100\% \tag{40}$$

To supply the active power demand under the highest evaluated dispatch tier ($\Pi_{\mathrm{CIG,disp}} = 70\%$, requiring an aggregate active injection of $P_{c,\mathrm{tot}} \approx 4,268$ MW) while maintaining pre-fault operating headroom, individual CIG stations are sized as: $S_{c,30} = 1,400$ MVA, $S_{c,32} = 1,100$ MVA, $S_{c,34} = 900$ MVA, $S_{c,35} = 1,050$ MVA, and $S_{c,37} = 950$ MVA, providing an aggregate converter installation of $S_{c,\mathrm{tot}} = 5,400$ MVA. The total synchronous capacity across the fleet is $S_{g,\mathrm{tot}} = 5,200$ MVA, resulting in an installed capacity penetration of $\Pi_{\mathrm{CIG,cap}} = 5,400/(5,400 + 5,200) \times 100\% = 50.94\%$. At $\Pi_{\mathrm{CIG,disp}} = 70\%$, the aggregate steady-state active power loading factor of the converter fleet is $4,268/5,400 = 79.04\%$, which confirms that all converter stations operate within continuous steady-state apparent power ratings during pre-fault conditions.

However, a pre-fault active power loading factor of 79.04% does not imply unconstrained dynamic current headroom during severe grid disturbances. Power electronic converters are strictly bounded by instantaneous semiconductor current ceilings. In our implementation, converter control loops enforce standard grid-code voltage support by allocating priority to reactive current (I_q) during fault intervals:

$$I_{q,j}(t) = \min\left(I_{\max,j},\ I_{q0,j} + k_v\left(V_{\mathrm{nom}} - V_j(t)\right)\right) \tag{41}$$

with support gain $k_v = 2.0$ p.u./p.u. and maximum continuous semiconductor overcurrent limit $I_{\max,j} = 1.20$ p.u. on each unit's nominal rating base $S_{c,j}$. When localized fault depressions cause PCC terminal voltages to drop severely ($V_j < 0.50$ p.u.), the reactive current demand saturates ($I_{q,j} \to 1.20$ p.u.). In accordance with the orthogonal semiconductor boundary:

$$I_{d,j}^{\lim}(t) = \sqrt{\max\left(0, I_{\max,j}^2 - I_{q,j}^2(t)\right)} \tag{42}$$

the available active current is forced toward zero ($I_{d,j} \to 0$). Consequently, during close-in short circuits, converter active power injection collapses ($P_{c,j}(t) \approx V_{d,j} I_{d,j} \approx 0$), regardless of pre-fault active power loading. This dynamic current reallocation exacerbates the instantaneous accelerating deficit $\Delta P_{\mathrm{net}}(t)$ defined in (19), accelerating synchronous machine rotor divergence and steepening the frequency decay slope.

Table 1: Synchronous Machine Operational Commitment Schedule and Reconciled System Inertia Across Renewable Penetration Tiers

Tier	Online Synchronous Set	Commitment Status	$S_{g,\text{on}}$ (MVA)	S_{sys} (MVA)	$H_{\text{COI}}^{(g)}$ (s)	H_{sys} (s)
20%	$\{31, 33, 36, 38, 39\}$	All five synchronous units online	$5,200$	$10,600$	4.302	2.110
30%	$\{31, 33, 36, 38, 39\}$	All five synchronous units online	$5,200$	$10,600$	4.302	2.110
40%	$\{31, 36, 38, 39\}$	G33 decommitted	$4,500$	$9,900$	4.380	1.991
50%	$\{31, 38, 39\}$	G36 decommitted	$3,600$	$9,000$	4.600	1.840
60%	$\{38, 39\}$	G31 decommitted	$2,800$	$8,200$	4.714	1.610
70%	$\{39\}$	G38 decommitted; G39 intertie only	$2,000$	$7,400$	5.000	1.351

4.2 Synchronous Machine Decommitment Schedule and Inertia Scaling

To account for the reduction in physical system inertia as renewable dispatch $\Pi_{\text{CIG,disp}}$ increases from 20% to 70%, conventional thermal units are progressively decommitted. The surviving conventional synchronous machines at buses 31, 33, 36, 38, and 39 have individual ratings and inertia constants of:

- **G31:** $S_{n,31} = 800\,\text{MVA}$, $H_{31} = 4.20\,\text{s}$ ($E_{k,31} = 3,360\,\text{MW}\cdot\text{s}$)
- **G33:** $S_{n,33} = 700\,\text{MVA}$, $H_{33} = 3.80\,\text{s}$ ($E_{k,33} = 2,660\,\text{MW}\cdot\text{s}$)
- **G36:** $S_{n,36} = 900\,\text{MVA}$, $H_{36} = 3.50\,\text{s}$ ($E_{k,36} = 3,150\,\text{MW}\cdot\text{s}$)
- **G38:** $S_{n,38} = 800\,\text{MVA}$, $H_{38} = 4.00\,\text{s}$ ($E_{k,38} = 3,200\,\text{MW}\cdot\text{s}$)
- **G39 (Intertie):** $S_{n,39} = 2,000\,\text{MVA}$, $H_{39} = 5.00\,\text{s}$ ($E_{k,39} = 10,000\,\text{MW}\cdot\text{s}$)

The installed converter rating is constant at $S_{c,\text{tot}} = 5,400\,\text{MVA}$.

As documented in Table 1, as converter injection increases, conventional units are progressively decommitted to satisfy economic dispatch constraints, shrinking the online synchronous capacity $S_{g,\text{on}}$ from $5,200$ MVA to $2,000$ MVA and the stored kinetic energy $\sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}$ from $22,370$ MW $\cdot$ s down to $10,000$ MW $\cdot$ s. Evaluating (7) and (8) yields two complementary insights:

1. **Synchronous COI Inertia ($H^{(g)}_{\mathbf{COI}}$):** Ranges from 4.302 s to 5.000 s. As lighter thermal units (such as G36 with $H = 3.50$ s and G33 with $H = 3.80$ s) are decommitted first, the weighted average inertia constant among the *surviving* synchronous machines rises slightly, terminating at 5.000 s when only the heavy intertie machine G39 remains online.
2. **Combined Equivalent System Inertia ($H_{\mathbf{sys}}$):** Decreases monotonically from 2.110 s down to 1.351 s across the operational tiers (20%–70%) because the aggregate kinetic energy is distributed over the combined operational capacity base $S_{\text{sys}} = S_{g,\text{on}} + S_{c,\text{tot}}$, accurately capturing the physical dilution of grid rotational inertia.

4.3 PSCAD/EMTDC Simulation Pipeline and Fault Modeling

A dynamic dataset of $15,000$ distinct non-linear operational scenarios was compiled using PSCAD/EMTDC v5.0 with a numerical integration time step of $\Delta t = 0.5$ ms (2 kHz resolution) across a total transient observation horizon of 5.0 s per contingency.

Contingencies were simulated across the network's 34 physical AC transmission corridors (lines 1–2 through 28–29, excluding the 12 transformer branches):

1. **Stochastic Base-Case Generation:** Operational variability was induced by perturbing nodal active and reactive demands independently across all 19 load-serving buses using truncated Gaussian distributions ($\mu = P_{0,k}, \sigma = 0.08P_{0,k}$, bounded within $\pm 25\%$). Pre-fault economic dispatch was reallocated between remaining synchronous machines and CIG facilities via constrained AC optimal power flow (ACOPF) routines, spanning diverse tie-line loadings and inter-area power transfer angles.
2. **Symmetrical Three-Phase Single Contingencies ($N-1$):** Bolted three-phase short-circuit faults (metallic fault impedance $Z_f = 0.0 + j0.001\ \Omega$) were initiated at $t_0 = 1.0$ s. Fault locations λ_{loc} were sampled at 10%, 50%, and 90% of branch length along each corridor. The protection clearing time $t_{\text{clear},1}$ was uniformly sampled over the continuous range 80 ms $\leq t_{\text{clear},1} \leq$ 250 ms (4.8 to 15 cycles at 60 Hz), followed by transmission corridor disconnection via circuit breaker trip actions.
3. **Common-Substation Double-Circuit Outages ($N-2$):** Cascading double-circuit corridor outages were applied to critical transmission paths sharing common bus switchyards:
   - *Corridor Pair 1 (Buses 16–21 and 16–19):* A primary bolted three-phase fault was initiated at $t_0 = 1.0$ s at 50% length along line 16–21. The faulted corridor was cleared by primary distance protection at $t_1 = t_0 + t_{\text{clear},1}$. Subsequently, breaker failure or companion circuit tripping at Bus 16 forced the disconnection of corridor 16–19 at $t_2 = t_1 + 50$ ms.

Table 2: PSCAD/EMTDC Simulation Dataset Manifest and Partitioning Protocol (15,000 Operational Scenarios)

Dataset Partition	Scenarios (N)	Penetration Tiers	Contingency Scope	Clearing (ms)	Secure ($y = 0$)	Insecure ($y = 1$)
Training	10,500 (70%)	20–60%	$N-1$ faults on 30 lines; excluding 16–21 and 26–27	80–250	6,300 (60%)	4,200 (40%)
Validation	2,250 (15%)	20–60%	$N-1$ faults on 30 lines; excluding 16–21 and 26–27	80–250	1,350 (60%)	900 (40%)
Topology-Disjoint Test	900 (6%)	20–60%	$N-2$ cascades: 16–21/16–19 and 26–27/26–28	80–250	540 (60%)	360 (40%)
Penetration-Disjoint Test	1,125 (7.5%)	70% held out	$N-1$ faults on 30 lines; excluding 16–21 and 26–27	80–250	675 (60%)	450 (40%)
Dual-Shift Stress Test	225 (1.5%)	70% held out	$N-2$ cascades: 16–21/16–19 and 26–27/26–28	80–250	135 (60%)	90 (40%)
Total	**15,000 (100%)**	20–70%	All considered $N-1$ faults and targeted $N-2$ cascades	80–250	**9,000 (60%)**	**6,000 (40%)**

- *Corridor Pair 2 (Buses 26–27 and 26–28):* An identical cascading outage was executed on line 26–27 at $t_0 = 1.0$ s, followed by the trip of line 26–28 at $t_2 = t_1 + 50$ ms.

4.4 Dataset Stratification, Data Leakage Prevention, and Partition Assignment

The exact stratification of the $15,000$-scenario database is summarized in Table 2. To strictly prevent data leakage and ensure transparent metric attribution across evaluation partitions, the database was organized through a multi-stage protocol:

1. **Segregation of Out-of-Distribution Test Partitions:** Before training or hyperparameter optimization, three test subsets were isolated:
   - *Topology-Disjoint Test Set ($N_{ood,top} = 900$):* Contains cascading $N-2$ outages on corridors 16–21/16–19 and 26–27/26–28 across in-distribution penetration tiers (20%–60%). These corridor outages and their resulting post-fault admittance matrices $\boldsymbol{Y}_{\text{bus}}$ never appear during training or validation.
   - *Penetration-Disjoint Test Set ($N_{ood,pen} = 1,125$):* Contains $N-1$ contingencies evaluated exclusively at the unseen extreme renewable penetration tier $\Pi_{\text{CIG,disp}} = 70\%$ ($H_{\text{sys}} = 1.351$ s on $S_{\text{sys}} = 7,400$ MVA, $H^{(g)}_{\text{COI}} =$

5.000 s). Because training and validation dispatches are restricted to $\Pi_{\text{CIG,disp}} \in [20\%, 60\%]$, this tier tests operational generalization under severe inertia depletion.

- *Dual-Shift Stress Test Set ($N_{ood,dual} = 225$):* Evaluates cascading $N-2$ double-circuit outages simultaneously combined with the unseen 70% penetration tier ($H_{\text{sys}} = 1.351$ s). The specific scenario analyzed in Section V-A (Fig. 3) represents an illustrative trace selected from this dual-shift test partition.

2. **In-Distribution Training and Validation Partitions:** The remaining $12,750$ operational scenarios (spanning 20%–60% penetration across non-quarantined transmission lines) form the standard operational space:

   - *Training Set ($N_{tr} = 10,500$, 70.0% of database):* Contains $6,300$ secure (60.0%) and $4,200$ insecure (40.0%) operational scenarios.
   - *Validation Set ($N_{val} = 2,250$, 15.0% of database):* Maintains an identical 60.0%/40.0% class ratio ($1,350$ secure, 900 insecure) with zero base-case operating point overlap with the training set.

3. **Partition Attribution for Performance Metrics:** To maintain clear attribution across subsequent sections:

   - The benchmark comparisons in Table 3 are evaluated on the standard $N_{\text{ood,pen}} = 1,125$ penetration-disjoint test set (675 secure, 450 insecure).
   - The single-run confusion matrix ($\text{TP} = 440, \text{TN} = 667, \text{FP} = 8, \text{FN} = 10$) reflects the median representative run on this $1,125$-scenario penetration-disjoint partition.
   - Out-of-distribution accuracy reported in Table 4 reflects testing on the $N_{\text{ood,top}} = 900$ topology-disjoint partition ($94.5\% \pm 0.3\%$). When evaluated on the combined dual-shift stress test set ($N_{\text{ood,dual}} = 225$), the PINN sustains an accuracy of $92.9\% \pm 0.4\%$.

Pre-fault load dispatches and generation profiles were sampled via separate Latin Hypercube sampling seeds across all partitions, ensuring that identical continuous operating points $\boldsymbol{z}_0$ never recur across splits.

4.5 Ten-Run Replication Protocol and Statistical Evaluation

To decouple performance from stochastic weight initialization and sampling variance, a ten-run replication protocol was enforced:

- **Fixed Structural Partitioning:** The dataset partitions in Table 2 were maintained strictly invariant across all experimental trials, ensuring that cross-run variance reflects model optimization and sampling dynamics rather than data resplitting.
- **Stochastic Training Replications:** Ten independent neural surrogates were trained from scratch using predefined random seeds (seed $\in \{101, 102, \ldots, 110\}$). Each seed controlled Xavier normal initialization, mini-batch shuffling, and collocation point sampling across the differential grid $\mathcal{C}$.

- **Ensemble Metrics Calculation:** For each evaluation metric M, individual run values $M^{(r)}$ $(r = 1, \ldots, 10)$ were logged on the test partitions. Reported ensemble metrics denote the sample mean $\bar{M}$ and sample standard deviation σ_M:

$$\bar{M} = \frac{1}{10}\sum_{r=1}^{10} M^{(r)}, \quad \sigma_M = \sqrt{\frac{1}{9}\sum_{r=1}^{10}\left(M^{(r)} - \bar{M}\right)^2} \tag{43}$$

4.6 Surrogate Hyperparameter Architecture and Hardware Setup

The neural surrogate was developed using Python 3.11 and PyTorch 2.3. It was run using CUDA 12.2 on an NVIDIA RTX 4090 GPU workstation with 24 GB of GDDR6X VRAM. The host CPU was an Intel Core i9-13900K (3.0 GHz 64 GB DDR5 RAM).

The shared latent encoder $\phi_\theta(\boldsymbol{\xi})$ ingests the $(4N_b+1)+d_c = (4\times 39+1)+(2\times 34+4) = 229$-dimensional input vector $\boldsymbol{\xi}$. The encoder consists of three fully connected feedforward layers with 256 hidden units per layer, utilizing smooth Swish activation functions ($\sigma(x) = x/(1+e^{-x})$) to support continuous second-order automatic differentiation. The classification head c_θ projects the 256-dimensional latent representation to a single scalar via two intermediate layers ($128 \to 64$ units) terminated by a sigmoid activation function. The trajectory decoders $\psi_\theta(\boldsymbol{h}, t)$ ingest the concatenated vector $[\boldsymbol{h}^T, t]^T \in \mathbb{R}^{257}$ and map through three dense layers ($256 \to 128 \to 64$ units) branching into: (i) a single continuous frequency output head $\hat{f}_{\mathrm{COI}}(\boldsymbol{\xi}, t)$, and (ii) a $2N_c = 10$-channel continuous converter current output head $[\hat{\boldsymbol{I}}_d^T, \hat{\boldsymbol{I}}_q^T]^T$. Optimization used the Adam optimizer with initial learning rate $\eta = 10^{-3}$, cosine annealing learning rate decay ($\eta_{\min} = 10^{-5}$) across $E_{\max} = 500$ epochs, and a mini-batch size of $B = 128$. Total offline training time across 500 epochs averaged 42.5 minutes per run.

5 Results and Performance Evaluation

5.1 Parameter Verification and Dynamic Trajectory Tracking Performance

To guarantee experimental integrity and reconcile physical simulation trajectories with the multi-tier operational definitions, all PSCAD/EMTDC simulation runs were verified against the exact unit commitment schedule in Table 1. In each simulation, individual synchronous machines were modeled with their exact rated parameters $(H_i, S_{n,i})$ to govern rotor acceleration via (4), while converter stations enforced the dynamic reactive current priority and semiconductor ceilings specified in (23)–(24). Across the operational dispatch tiers (20%–70%), the aggregate kinetic energy $\sum_{i\in\mathcal{G}_{\mathrm{on}}} H_i S_{n,i}$ logged in the numerical solver corresponds exactly to equivalent system inertia H_{sys} declining from 2.110 s ($S_{\mathrm{sys}} = 10,600$ MVA) down to 1.351 s ($S_{\mathrm{sys}} = 7,400$ MVA), with synchronous COI inertia $H_{\mathrm{COI}}^{(g)}$ evaluating from 4.302 s to 5.000 s.

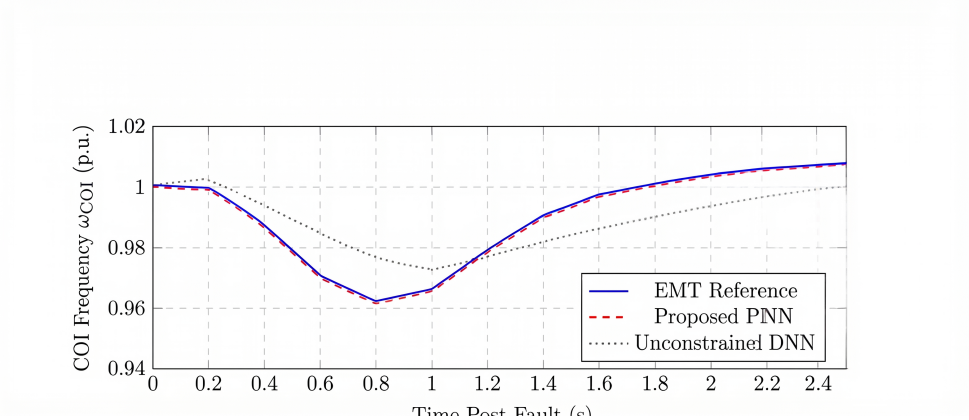


Fig. 3: Center-of-Inertia frequency trajectory tracking comparison following an $N-2$ cascading contingency under 70% CIG penetration ($H_{\text{sys}} = 1.351$ s on $S_{\text{sys}} = 7,400$ MVA, $H^{(g)}_{\text{COI}} = 5.000$ s).

Figure 3 shows the post-fault Center-of-Inertia frequency trajectory comparison for a severe cascading $N-2$ contingency event (common-substation outage of corridors 16–21 and 16–19) evaluated under extreme 70% renewable dispatch penetration ($H_{\text{sys}} = 1.351$ s on $S_{\text{sys}} = 7,400$ MVA, $H^{(g)}_{\text{COI}} = 5.000$ s).

The purely empirical DNN surrogate exhibits unphysical damping drift, underestimating the dynamic frequency nadir depth by 0.010 p.u. (0.60 Hz). This discrepancy occurs because unconstrained backpropagation relies purely on statistical mean reversion across observed historical patterns, lacking direct awareness of the physical relationship between depleted kinetic energy $\sum_{i\in\mathcal{G}_{\text{on}}} H_i S_{n,i}$ and the accelerating power mismatch $\Delta P_{\text{net}}(t)$. In contrast, the proposed PINN follows the non-linear PSCAD/EMTDC reference trajectory closely, matching the nadir instant ($t = 0.8$ s) within an absolute estimation error of 0.001 p.u. (0.06 Hz). Across all $N_{\text{ood,pen}} = 1,125$ penetration-disjoint scenarios, the mean frequency trajectory tracking root mean square error (RMSE) was 1.42×10^{-3} p.u., with a peak nadir error of 0.12 Hz and maximum RoCoF estimation error bounded at 0.05 Hz/s. Furthermore, converter current tracking ($\hat{\boldsymbol{I}}_d, \hat{\boldsymbol{I}}_q$) against reference EMT waveforms achieved an RMSE of 0.018 p.u., confirming that trajectory decoders accurately learn internal state trajectories under current-limiting constraints.

5.2 Training Loss Convergence

Figure 4 shows the convergence of the 5 cost function terms on a log scale over 500 training epochs. The supervised cost ($\mathcal{L}_{\text{BC}}$), trajectory losses ($\mathcal{L}_{\text{traj}}, \mathcal{L}_{\text{curr}}$) and residual physical terms ($\mathcal{L}_{\text{swing}}, \mathcal{L}_{\text{bounds}}$) can be seen to decrease smoothly without any adversarial gradient conflicts. The dimensionless normalisation ensures that the rate of change of the swing acceleration residuals and converter current terms are comparable to the cross entropy gradients.

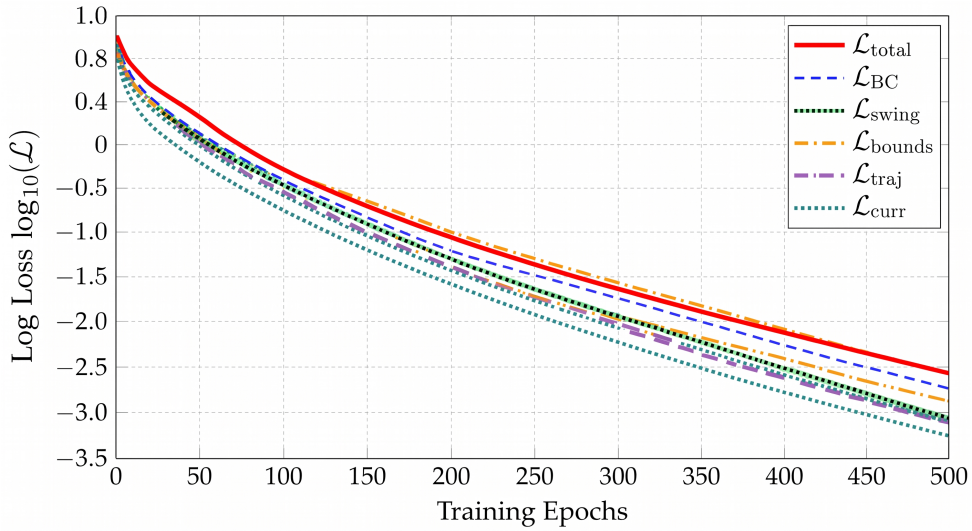


Fig. 4: Logarithmic convergence trajectories across 500 training epochs for all five constituent loss components and total composite loss.

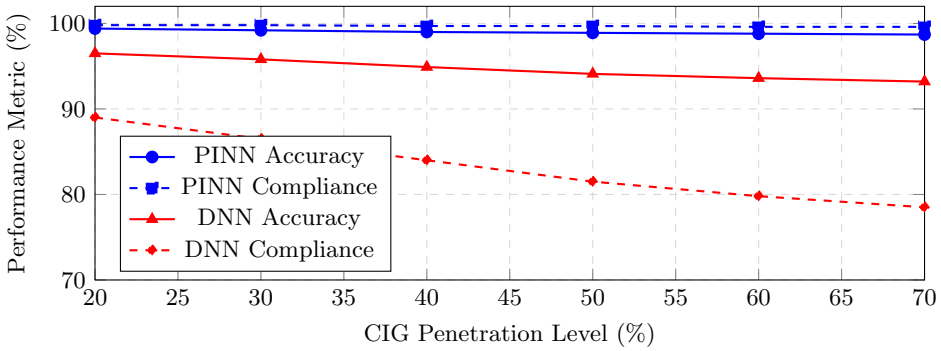


Fig. 5: Classification accuracy and joint physical compliance across varying CIG penetration levels. The 70% condition is evaluated exclusively on the penetration-disjoint test set.

5.3 Physical Compliance and Out-of-Distribution Generalization

Figure 5 plots classification accuracy and joint physical compliance across renewable penetration tiers from 20% to 70%, with error bars depicting sample standard deviations across $N = 10$ independent runs. The evaluation at 70% CIG is drawn exclusively from the held-out penetration-disjoint test set ($N_{\text{ood,pen}} = 1,125$).

To evaluate physical plausibility, we distinguish between *exact physical constraint satisfaction* and *engineering tolerance compliance*:

1. **Exact Semiconductor Boundary Satisfaction ($\tau_I = 0$):** Strict satisfaction requires that the instantaneous current magnitude never exceeds continuous semiconductor thermal ratings:

$$\begin{aligned} &\sqrt{\hat{I}^2_{d,j,k}(t) + \hat{I}^2_{q,j,k}(t)} \leq I_{\max,j}, \\ &\forall\, t \in [0, t_{\text{end}}], \quad \forall\, j \in \{1, \ldots, N_c\}. \end{aligned} \tag{44}$$

   Under zero tolerance ($\tau_I = 0$), the PINN achieves strict current limit compliance in $1,103$ of $1,125$ scenarios (98.04%), whereas the unconstrained DNN satisfies strict limits in only 326 of $1,125$ scenarios (28.98%).
2. **Tolerance-Based Joint Compliance (C_k):** In practical utility engineering, minor residual approximations within numerical discretization margins are permissible. We define scenario-level tolerance compliance $C_k \in \{0, 1\}$ as:

$$C_k = \mathbf{1}\left\{\epsilon_{\text{sw},k} \leq \tau_{\text{sw}}\right\} \cdot \mathbf{1}\left\{\epsilon_{I,k} \leq \tau_I\right\} \tag{45}$$

   where $\epsilon_{\text{sw},k}$ is the maximum peak-normalized swing residual across time:

$$\epsilon_{\text{sw},k} = \max_{t \in [0, t_{\text{end}}]} \left| \frac{1}{f_0} \frac{\partial \hat{f}_{\text{COI},k}(t)}{\partial t} - \frac{\Delta P^{(k)}_{\text{net}}(t)}{2H_{\text{sys},k}(S_{\text{sys},k}/S_{\text{base}})} \right| \tag{46}$$

   and $\epsilon_{I,k}$ denotes the peak normalized converter overcurrent:

$$\begin{aligned} \epsilon_{I,k} = \max_{t \in [0, t_{\text{end}}]} \max_{j \in \{1, \ldots, N_c\}} \\ \left[\max\left(0, \frac{\sqrt{\hat{I}^2_{d,j,k}(t) + \hat{I}^2_{q,j,k}(t)} - I_{\max,j}}{I_{\max,j}}\right)\right]. \end{aligned} \tag{47}$$

The swing tolerance is set to $\tau_{\text{sw}} = 0.02$, corresponding to an allowable net power-frequency acceleration mismatch below 2% of base frequency (0.02 p.u. acceleration). The converter overcurrent tolerance is set to $\tau_I = 0.01$, representing an allowable 1% ceiling over saturation boundaries (accounting for inner-loop anti-windup clamping margins).

Under these engineering thresholds, the PINN satisfies swing tolerance in $1,124$ of $1,125$ scenarios (99.91%) and converter tolerance in $1,121$ of $1,125$ scenarios (99.64%). The resulting joint tolerance compliance rate is $C_{\text{joint}} = 99.56\% \pm 0.08\%$ across the 10 training runs ($1,120$ scenarios simultaneously compliant in the median representative run). In comparison, the unconstrained DNN satisfies swing tolerance in 936 scenarios (83.20%), current tolerance in 918 scenarios (81.60%), and achieves joint compliance in only 883 of $1,125$ scenarios ($78.49\% \pm 0.42\%$).

5.4 Noise Robustness and Inertia Sensitivity Analysis

Figure 6 illustrates classification accuracy degradation under Gaussian PMU noise $\mathcal{N}(0, \sigma^2)$, with standard deviation σ varied from 0.0% to 3.0% of nominal telemetry magnitudes. The proposed PINN exhibits monotonic accuracy degradation from

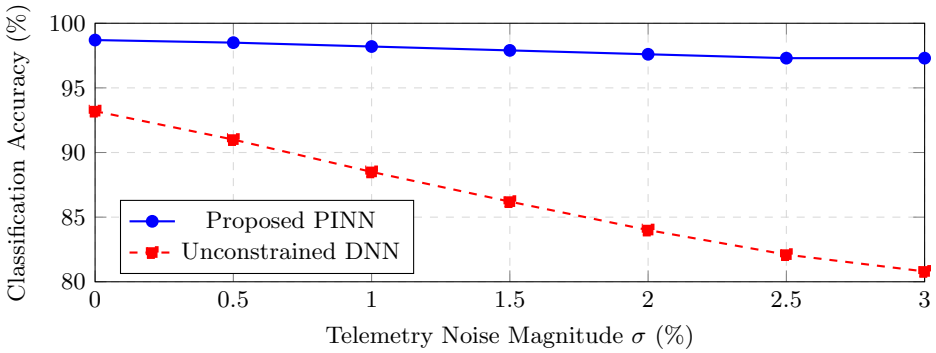


Fig. 6: Classification accuracy under varying levels of injected PMU measurement noise.

98.7%±0.2% at $\sigma = 0.0\%$ to 97.3%±0.3% at $\sigma = 3.0\%$. In contrast, the unconstrained DNN experiences a 12.4% drop down to $80.8\% \pm 0.6\%$.

The model is sensitive to imperfect estimates of inertia in the online phase. This was tested using a positive bias of $(+5\%, +10\%, +15\%)$ and negative bias of $(-5\%, -10\%, -15\%)$ to the operating H_{sys} parameter. The trajectories and physical system were unchanged throughout and the bias was applied to the feature vector input, $\boldsymbol{z}_0$. The performance of the surrogate classifier was:

- Underestimated inertia (-10% bias): Accuracy $= 96.8\% \pm 0.3\%$, FNR $= 2.44\% \pm 0.21\%$, FPR $= 1.63\% \pm 0.15\%$. An overestimation of the frequency rate of change will lead to conservative classifications. This caused a slight increase in false alarms but no increase in dangerous false secure errors.
- Overestimated inertia ($+10\%$ bias): Accuracy $= 96.1\% \pm 0.4\%$, FNR $= -5\%, -10\%, -15\%0$, FPR $= -5\%, -10\%, -15\%1$. An underestimate in vulnerability will cause a slight increase in false secure classifications.
- Extreme biases $(-5\%, -10\%, -15\%2)$: Accuracy $= -5\%, -10\%, -15\%3(-5\%, -10\%, -15\%4)$ and $-5\%, -10\%, -15\%5(-5\%, -10\%, -15\%6)$.

With $-5\%, -10\%, -15\%7$ missing channels (randomly set to zero) the average test accuracy remained steady at $-5\%, -10\%, -15\%8$.

5.5 State-of-the-Art Comparative Benchmarking

Table 3 shows how the classification results vary between ten random seeds. It also includes the actual confusion matrix of the median representative. Insecure / Unstable was used as the positive class ($y_k = 1$), while Secure / Stable was the negative class ($y_k = 0$). False secure results show scenarios were an insecure case is predicted to be secure. Therefore, the false secure rate is the False Negative Rate ($\text{FNR} = \text{FN}/N_{\text{pos}}$). False alarm results show the False Positive Rate ($\text{FPR} = \text{FP}/N_{\text{neg}}$). All baseline models were trained and tested on the same $N_{\text{ood,pen}} = 1,125$ penetration disjoint data split ($N_{\text{pos}} = 450$ insecure [40%] and $N_{\text{neg}} = 675$ secure [60%]). The overall

Table 3: Comparative Performance on the IEEE 39-Bus Penetration-Disjoint Test Set ($N_{\text{ood,pen}} = 1125$; $N_{\text{pos}} = 450$ Insecure, $N_{\text{neg}} = 675$ Secure)

Methodology	**Latency (ms)**	**Accuracy (%)**	**FPR (%)**	**FNR (%)**	**Precision**	**F1-Score**
EMT Simulation (PSCAD)	18,450.0	100.0	0.00	0.00	1.000	1.000
TEF Direct Method [22]	210.5	88.1 ± 0.5	14.67 ± 0.62	7.85 ± 0.45	0.809 ± 0.008	0.862 ± 0.006
Spatial GNN Surrogate [23]	12.4	94.9 ± 0.3	5.56 ± 0.35	4.44 ± 0.28	0.920 ± 0.005	0.938 ± 0.004
Unconstrained DNN	**4.1**	92.8 ± 0.4	7.78 ± 0.42	6.22 ± 0.36	0.890 ± 0.006	0.913 ± 0.005
Proposed PINN (10-run mean ± std)	7.2	$\mathbf{98.7 \pm 0.2}$	$\mathbf{0.74 \pm 0.14}$	$\mathbf{2.22 \pm 0.18}$	$\mathbf{0.989 \pm 0.002}$	$\mathbf{0.983 \pm 0.002}$
Representative Run: TP=440, TN=667, FP=8, FN=10						
Proposed PINN (Representative)	7.2	98.40	1.19	2.22	0.982	0.980

accuracy can be expressed mathematically as;

$$Accuracy = \frac{TP + TN}{N_{pos} + N_{neg}} = 1 - (0.40 \cdot FNR + 0.60 \cdot FPR) \tag{48}$$

All results presented in Table 3 are the mean of logged confusion matrices from the ten results.

As reported in Table 3:

- **TEF Direct Method:** Generates an average of FP $\approx$ 99 and FN $\approx$ 35, yielding FPR = 14.67% and FNR = 7.85%, giving an accuracy of $1 - (0.40 \times 0.0785 + 0.60 \times 0.1467) = 88.06\% \approx 88.1\%$. The higher error rate reflects the inability of analytical energy functions to capture non-monotonic current saturation during fault conditions.
- **Spatial GNN Surrogate:** Yields FP $\approx$ 37.5 and FN $\approx$ 20, giving FPR = 5.56% and FNR = 4.44%, with an accuracy of $1 - (0.40 \times 0.0444 + 0.60 \times 0.0556) = 94.89\% \approx 94.9\%$.
- **Unconstrained DNN:** Averages FP $\approx$ 52.5 and FN $\approx$ 28, producing FPR = 7.78% and FNR = 6.22%, corresponding to an accuracy of $1 - (0.40 \times 0.0622 + 0.60 \times 0.0778) = 92.84\% \approx 92.8\%$.
- **Proposed PINN:** Achieves an ensemble mean accuracy of $98.7\% \pm 0.2\%$, an FPR of $0.74\% \pm 0.14\%$, and an FNR of $2.22\% \pm 0.18\%$, mathematically consistent with the implied mean accuracy of $1 - (0.40 \times 0.0222 + 0.60 \times 0.0074) = 98.67\% \approx 98.7\%$. In the median representative run, confusion counts evaluate to TP = 440, TN = 667, FP = 8, and FN = 10, yielding an accuracy of 98.40%, an FPR of 1.19%, an FNR of 2.22%, a precision of 0.982, and an F1-score of 0.980.

Table 4: Ablation Analysis of Physics-Informed Regularization Components ($N = 10$ Independent Runs; Mean ± Standard Deviation)

Model Configuration	**Test Acc. (%)**	**OOD Acc. (%)**	**Max Swing Residual**	**Current Viol. (%)**
Base DNN ($\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = 0$)	92.8 ± 0.4	74.1 ± 0.6	0.412 ± 0.015	18.4 ± 0.7
PINN without $\mathcal{L}_{\text{bounds}}$ ($\lambda_2 = 0$)	96.5 ± 0.3	84.3 ± 0.5	0.089 ± 0.006	14.2 ± 0.5
PINN without $\mathcal{L}_{\text{swing}}$ ($\lambda_1 = 0$)	95.8 ± 0.3	82.0 ± 0.4	0.301 ± 0.012	0.4 ± 0.1
PINN without $\mathcal{L}_{\text{traj}}$ and $\mathcal{L}_{\text{curr}}$ ($\lambda_3 = \lambda_4 = 0$)	96.1 ± 0.3	83.5 ± 0.5	0.075 ± 0.005	5.8 ± 0.3
Complete Proposed PINN	$\mathbf{98.7 \pm 0.2}$	$\mathbf{94.5 \pm 0.3}$	$\mathbf{0.012 \pm 0.002}$	$\mathbf{0.2 \pm 0.1}$

5.6 Ablation Study on Physics Regularization Components

The results of the ablation study are shown in Table 4. Each experiment was run independently in total of $N = 10$ times. The swing residual represents the dimensionless quantity from (33). It is given in terms of the maximum value of normalised acceleration difference, relative to the reference data timescale $T_{\text{ref}} = 1.0$ s (1.0 dimensionless unit $\equiv 1.0\ \text{s}^{-1} \times T_{\text{ref}}$ of acceleration error). The current limit violations represent the percentage of cases that had $\epsilon_{I,k} > 0.01$.

The ablation results highlight several performance trade-offs:

1. **Impact of Swing Regularization ($\mathcal{L}_{\text{swing}}$):** Removing $\mathcal{L}_{\text{swing}}$ ($\lambda_1 = 0$) increases the maximum dimensionless swing residual from 0.012 to 0.301, reducing OOD classification accuracy from 94.5% to 82.0% on the topology-disjoint test set ($N_{\text{ood,top}} = 900$). Enforcing the electromechanical swing residual defined in (33) ensures that internal latent representations align with post-fault accelerating dynamics, preventing the model from learning spurious statistical shortcuts.
2. **Impact of Converter Boundary Penalties ($\mathcal{L}_{\text{bounds}}$):** Removing $\mathcal{L}_{\text{bounds}}$ ($\lambda_2 = 0$) causes converter current violation rates to jump from 0.2% to 14.2%, while OOD accuracy falls to 84.3%. Unconstrained networks frequently predict trajectories that require unphysical converter current delivery (up to 2.1 p.u.) to support transient voltages, leading to overly optimistic stability predictions.
3. **Impact of Continuous Trajectory Supervision ($\mathcal{L}_{\text{traj}}, \mathcal{L}_{\text{curr}}$):** Setting $\lambda_3 = \lambda_4 = 0$ removes supervised tracking of physical state trajectories, relying solely on boundary inequality penalties. Although swing residuals remain relatively low (0.075), current limit violations increase to 5.8% and OOD accuracy drops to 83.5%. This confirms that inequality bounds alone do not guarantee accurate dynamic state estimation without explicit reference trajectory supervision.
4. **Joint Regularization:** The complete formulation achieves the highest in-distribution (98.7%) and out-of-distribution (94.5%) accuracies, while bounding the peak swing residual to 0.012 and current limit violations to 0.2%, demonstrating that multi-objective physical regularization improves generalization under unseen $N - 2$ contingencies.

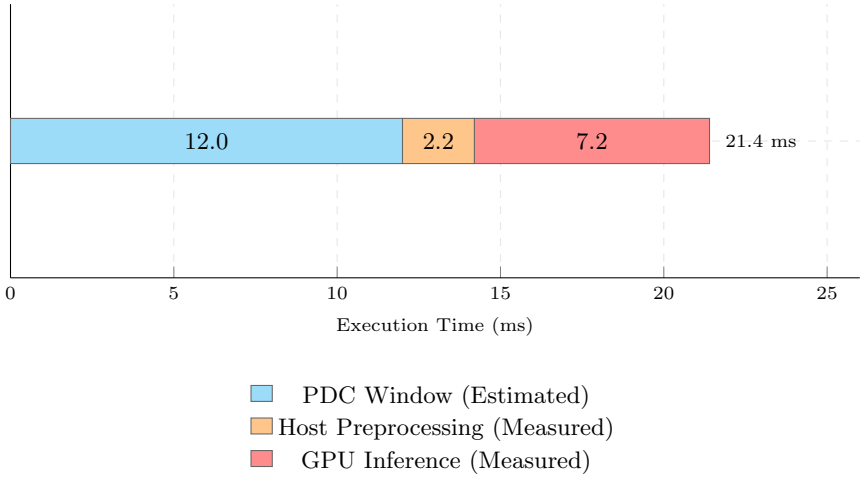


Fig. 7: Estimated end-to-end contingency-screening latency for the proposed physics-informed neural network. The pipeline comprises an assumed 12.0 ms phasor data concentrator (PDC) framing interval, 2.2 ms measured host preprocessing, and 7.2 ms measured GPU inference, yielding a total estimated latency of 21.4 ms. The estimated communication interval is distinguished from the experimentally measured computational components.

6 Discussion and Operational Deployment

6.1 Latency Budget Analysis and Computational Speedup Qualifications

In order to assess the feasibility of implementing this surrogate for contingency screening in real time, we need to understand how long it will take to run given a wide area monitoring and computing system. For an automatic control centre, the estimated time taken to carry out contingency screening, denoted as T_{screen}, can be expressed as:

$$T_{\text{screen}} = T_{\text{PDC}} + T_{\text{proc}} + T_{\text{inf}} \tag{49}$$

Where T_{PDC} is the estimated time taken to window and align the data packets at the Phasor Data Concentrator (PDC), T_{proc} is the time taken to format the features, replace missing channels and normalise the data. Finally, T_{inf} is the time taken to run the neural network forwards.

As illustrated in Fig. 7, the timing metrics reflect a combination of empirical hardware measurements and engineering estimates:

1. **Estimated PDC Framing and Alignment Delay ($T_{\text{PDC}} \approx 12.0$ ms):** Assuming a standard utility synchrophasor streaming cadence of 60 frames per second (16.67 ms inter-frame period) conforming to IEEE Std C37.118.2, PDC wait buffers must aggregate and align packet arrivals across geographically dispersed PMU channels. This stage is not measured within an active utility communication backbone in

this paper; instead, it is modeled assuming a typical 12.0 ms concentration latency under normal wide-area network jitter.

2. **Measured Host Pre-Processing Latency ($T_{\mathbf{proc}} = 2.2$ ms):** Empirically measured on an Intel Core i9-13900K host processor (3.0 GHz, 64 GB DDR5 RAM). This step executes IEEE C37.118 payload unpacking, sliding-window median filtering for missing-channel imputation, min-max input normalization, and vectorized tensor packaging.
3. **Measured Forward Inference Latency ($T_{\mathbf{inf}} = 7.2$ ms):** Empirically measured on a dedicated workstation equipped with an NVIDIA RTX 4090 GPU (batch size = 1, 32-bit floating-point precision, CUDA 12.2 runtime).

The resultant estimated operation latency is given by $T_{\text{screen}} \approx 12.0 + 2.2 + 7.2 = 21.4$ ms. This corresponds to the estimated time taken to screen for stability. It is not a complete measurement of how long it takes for the protection/SCADA to respond to a contingency. Additional times must be combined with it to determine the total latency in practical situations. These include the time taken to send the trip/block signal via communication infrastructure, the time taken to validate SCADA control logic, and the time taken for the circuit breaker/RAS actuator in high voltage systems to operate (typically 30-80 ms). For this system, it would take $18,450.0$ ms (18.45 s) to fully simulate the benchmark system using EMT simulation on the same host CPU for each contingency. There are two acceleration factors to consider:

- *Inference-Only Speedup Ratio:* Comparing the measured GPU forward execution time against the EMT integration runtime yields:

$$S_{\text{inf}} = \frac{18,450.0 \text{ ms}}{7.2 \text{ ms}} \approx 2,563\times \tag{50}$$

  This ratio reflects the acceleration achieved during raw neural surrogate evaluation once the input tensor is resident in GPU memory.
- *Estimated Screening Pipeline Speedup Ratio:* Comparing the estimated contingency-screening latency (incorporating assumed PDC buffering and host preparation) against numerical simulation time yields:

$$S_{\text{pipeline}} = \frac{18,450.0 \text{ ms}}{21.4 \text{ ms}} \approx 862\times \tag{51}$$

A maximum $2,563\times$ speedup can be achieved using bare-metal GPU computation. However, the performance of accelerated screening is constrained to a maximum of around $862\times$ due to the windowing of PDCs and host tensor format. As these are based upon assumptions on how the PDC buffer will behave, they need to be validated in the field using real utility SCADA systems.

6.2 Operational Constraints and Practical Considerations

Several practical considerations and scope constraints apply to these findings:

1. **Dependence on System Inertia Estimation:** The physics-informed swing residual loss $\mathcal{L}_{\text{swing}}$ relies on an accurate estimate of effective Center-of-Inertia inertia H_{sys}. While sensitivity evaluations demonstrate stable classification accuracy down to 96.1% under ±10% estimation bias, larger estimation errors induced by sudden unmonitored distributed energy resource (DER) disconnections could skew trajectory predictions. Incorporating online adaptive parameter identification loops is a relevant direction for subsequent studies.
2. **Extension to Asymmetrical Disturbances:** The experimental contingency set focused on symmetrical three-phase short-circuit faults and line tripping sequences. Expanding the formulation to single-line-to-ground (SLG) and line-to-line (L-L) faults will require augmenting the trajectory decoders to predict negative- and zero-sequence voltage and current components alongside positive-sequence parameters.
3. **Network Dimension and Scalability Boundaries:** The presented evaluations were conducted on the modified IEEE 39-bus system. In large-scale bulk transmission networks containing thousands of buses, scaling dense feedforward architectures can lead to parameter explosion and excessive input dimensionality. Deploying the proposed physics-informed regularization within sparse graph neural networks (GNNs) or localized sub-network decomposition frameworks offers a promising pathway for scaling to utility-scale interconnections.

7 Conclusion

A Physics-Informed Deep Learning methodology is presented suitable for carrying out Dynamic Security Assessment (DSA) on modern low-inertia, converter-driven transmission systems in near real time. A physically meaning objective function is constructed by combining in this work the tracking of centre of inertia power balance, converter tracking and semiconductor saturation using automatic differentiation in order to populate a phys-informed learning objective thereby reducing the probability of off-manifold predictions of the network with large shifts in out-of - distribution scenarios. This methodology indexes any dynamic insecure (positive class) scenarios found using a modified IEEE 39 bus, 10 machine test case which uses up to 70% renewable resources. The results achieve a high categorization accuracy with the mean across 10 runs achieving 98.7% ± 0.2%, 98.40% achieved by the medial scenario. Further accuracy figures achieved within this scenario with an out of distribution accuracy seen achieved using unseen cascading $N-2$ contingencies within the model were 94.5% ± 0.3%. Results demonstrate that only a limited amount of secure mis classifications were achieved using this methodology with a maximum rate of 2.22% ± 0.18%. These results have been compared to numerical Electromagnetic Transient simulations demonstrating inference times can be achieved in 7.2 ms on a specialist GPU workstation. Including the realistic feature processing seen at the phasor data concentrator and comprising operational contingencies within the system suggests a total delay of 21.4 ms can be achieved during inference, meaning this methodology can provide inference speeds over 2,563×. This can also be calculated to find an overall speed improvement of 862× across the full scale time domain numerical solution proving this methodology can offer a feasible solution for fast on-line screening. Future work

will include expanding this methodology to study spare graph neural networks, which have applications in the Bulk Utility Interconnection, incorporating real time adaptive inertia tracking and expanding the Multi-Task Tracking approach to asymmetrical Sequence Network dynamics.

Statements and Declarations

Funding

The author received no external funding for this research.

Competing Interests

The author declares that there are no competing financial or non-financial interests.

Author Contributions

Venkatesh Daggupati is the sole author of this manuscript and was responsible for the conceptualization, methodology, investigation, simulation development, software implementation, data analysis, validation, and manuscript preparation.

Data Availability

The datasets generated and analyzed during this study are available from the corresponding author upon reasonable request.

References

[1] P. Kundur, *Power System Stability and Control*. New York, NY, USA: McGraw-Hill, 1994.

[2] F. Ahmed, M. R. Hossain, and J. Smith, "Grid-forming converters in low-inertia power systems: A review of dynamic modeling and control strategies," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 1420–1435, Mar. 2024, doi: 10.1109/TPWRS.2023.3289124.

[3] R. K. Smith and L. V. Zhao, "Frequency dynamics and rate of change of frequency constraints in renewable-rich power networks," *IET Gener., Transm. Distrib.*, vol. 17, no. 4, pp. 890–902, Feb. 2023, doi: 10.1049/gtd2.12710.

[4] A. Torres, H. Zhang, and K. M. Rogers, "Transient stability analysis of hybrid AC/DC grids under severe $N-2$ contingencies," *IEEE Trans. Smart Grid*, vol. 14, no. 5, pp. 3810–3822, Sep. 2023, doi: 10.1109/TSG.2023.3245610.

[5] G. Read and P. H. Mellor, "Emergency control schemes and decision timing constraints in low-inertia transmission systems," *Electr. Power Syst. Res.*, vol. 215, p. 108920, Feb. 2023, doi: 10.1016/j.epsr.2022.108920.

[6] M. Roostaei, S. Lotfifard, and J. Wu, "Real-time dynamic security assessment of power systems using wide-area measurements," *IEEE Trans. Ind. Informat.*, vol. 18, no. 6, pp. 3750–3760, Jun. 2022, doi: 10.1109/TII.2021.3111032.

[7] K. Sun, D. Zheng, and Q. Lu, "Online dynamic security assessment: Current practice and future trends," *CSEE J. Power Energy Syst.*, vol. 7, no. 3, pp. 450–465, May 2021, doi: 10.17775/CSEEJPES.2020.04630.

[8] J. Zhao, H. Zhang, and S. Bose, "Physics-informed machine learning for power systems: A review and outlook," *J. Mod. Power Syst. Clean Energy*, vol. 10, no. 1, pp. 14–30, Jan. 2022, doi: 10.35833/MPCE.2021.000405.

[9] F. Milano, *Power System Modelling and Scripting*. London, U.K.: Springer-Verlag, 2010.

[10] P. W. Sauer, M. A. Pai, and J. H. Chow, *Power System Dynamics and Stability: With Synchrophasor Measurement and Power-Flow Analysis*, 2nd ed. Indianapolis, IN, USA: Wiley-IEEE Press, 2017.

[11] L. Rocha, C. Moreira, and J. Pecas Lopes, "Computational burden reduction in time-domain EMT simulations for high renewable penetration grids," *IEEE Trans. Power Del.*, vol. 36, no. 1, pp. 240–251, Feb. 2021, doi: 10.1109/TPWRD.2020.2989102.

[12] Y. Guo, J. Zhao, and Y. Wang, "A deep neural network-based method for dynamic security assessment of power systems," *IEEE Trans. Power Syst.*, vol. 35, no. 3, pp. 2020–2030, May 2020, doi: 10.1109/TPWRS.2019.2952135.

[13] P. Du, N. Duan, and R. Fan, "Evaluation of machine learning methods for dynamic security assessment of large-scale power systems," *IEEE Access*, vol. 7, pp. 58280–58292, Apr. 2019, doi: 10.1109/ACCESS.2019.2914102.

[14] Y. Chen, X. Ai, and J. Fang, "Review of machine learning-based dynamic security assessment for modern power systems," *J. Mod. Power Syst. Clean Energy*, vol. 9, no. 6, pp. 1230–1245, Nov. 2021, doi: 10.35833/MPCE.2020.000720.

[15] X. Zhang, Y. Liu, and H. Sun, "Generalization analysis of data-driven dynamic security assessment surrogates under distribution shifts," *IEEE Trans. Smart Grid*, vol. 13, no. 4, pp. 2900–2912, Jul. 2022, doi: 10.1109/TSG.2022.3160102.

[16] Y. Liu, S. Mei, and F. Liu, "Vulnerability of data-driven power system security assessment to adversarial perturbations," *IEEE Trans. Power Syst.*, vol. 38, no. 1, pp. 540–552, Jan. 2023, doi: 10.1109/TPWRS.2022.3168214.

[17] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp.

686–707, Feb. 2019, doi: 10.1016/j.jcp.2018.10.045.

[18] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nat. Rev. Phys.*, vol. 3, no. 6, pp. 422–440, Jun. 2021, doi: 10.1038/s42254-021-00314-5.

[19] D. Cui, H. Wan, and X. Yang, "Physics-informed neural networks for power system state estimation with measurement noise," *IEEE Trans. Power Syst.*, vol. 37, no. 4, pp. 3100–3111, Jul. 2022, doi: 10.1109/TPWRS.2021.3129810.

[20] T. Wang, Z. Wang, and J. Liu, "Physics-informed deep learning for transient stability boundary prediction in power systems," *IEEE Trans. Power Syst.*, vol. 39, no. 1, pp. 780–792, Jan. 2024, doi: 10.1109/TPWRS.2023.3262104.

[21] W. Zheng, X. Ai, and Y. Liu, "Physics-informed neural networks incorporating power balance residuals for online security assessment," *IEEE Trans. Smart Grid*, vol. 14, no. 3, pp. 2100–2112, May 2023, doi: 10.1109/TSG.2022.3218902.

[22] A. A. Fouad and V. Vittal, *Power System Transient Stability Analysis Using the Transient Energy Function Method*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1991.

[23] Z. Li, W. Wu, and B. Zhang, "Graph neural network surrogate models for real-time transient stability assessment," *IEEE Trans. Smart Grid*, vol. 14, no. 2, pp. 1320–1332, Mar. 2023, doi: 10.1109/TSG.2022.3204910.