The American Journal of Engineering and Technology
ISSN 2689-0984 Volume 08 - 2026

Physics-Informed Modeling and Stability-Constrained Intelligent Control of Grid-Forming Battery Storage in Low-Inertia Power Systems

Oybek Achilovich Kayumov
Department of Computer Science and Programming, Jizzakh Branch of the National University of Uzbekistan named after Mirzo Ulugbek, Jizzakh, Uzbekistan

Nazokat Rashitovna Kayumova
Department of Computer Science and Programming, Jizzakh Branch of the National University of Uzbekistan named after Mirzo Ulugbek, Jizzakh, Uzbekistan

Received: 28 Feb 2026 | Received Revised Version:14 Mar 2026 | Accepted: 04 Apr 2026 | Published: 30 Apr 2026

Volume 08 Issue 04 2026 | Crossref DOI: 10.37547/tajet/Volume08Issue04-17

Abstract

The declining synchronous inertia of power systems with high shares of inverter-based resources makes the tuning of grid-forming battery energy storage systems (BESS) both more important and more uncertain, because the inertia that the storage must complement is not measured directly and changes with the unit commitment. This paper proposes a framework that couples physics-informed estimation with stability-constrained learning. First, a physics-informed neural network (PINN) embeds the swing and governor equations into the training loss and identifies the synchronous inertia and frequency stiffness from a single disturbance recorded by phasor measurement units, using the measured BESS power as an input. Second, a neural policy maps these estimates to the virtual inertia and damping of a virtual synchronous machine; it is trained with an augmented Lagrangian method subject to closed-form constraints on the rate of change of frequency (RoCoF), frequency nadir, damping ratio and converter power, tightened by the estimation uncertainty. A safety layer certifies each output against the same constraints. In simulations with inverter-based shares of 50–85 %, the PINN estimated inertia within −10.5 % to +7.3 % and frequency stiffness within 0.3 % for measurement noise up to 20 mHz, whereas the RoCoF-based method was biased by +20 % to +49 %. In closed-loop tests on 300 random operating points of a nonlinear model, the proposed controller produced no violation of the 49.5 Hz nadir and 0.5 Hz/s RoCoF limits and reduced the sustained storage reserve by 61 % relative to fixed gains; a policy trained with a fixed penalty violated the limits in 59.7 % of cases.

Keywords: Grid-forming inverter; battery energy storage; physics-informed neural network; inertia estimation; virtual synchronous machine; safe learning; frequency stability; low-inertia power systems.

© **2026**: Nazokat Rashitovna Kayumova, & Oybek Achilovich Kayumov. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). The authors retain copyright and allow others to share, adapt, or redistribute the work with proper attribution.

Cite This Article: Nazokat Rashitovna Kayumova, & Oybek Achilovich Kayumov. (2026). Physics-Informed Modeling and Stability-Constrained Intelligent Control of Grid-Forming Battery Storage in Low-Inertia Power Systems. The American Journal of Engineering and Technology, 8(4), 211–219. https://doi.org/10.37547/tajet/Volume08Issue04-17

1. Introduction

Power systems are replacing synchronous generators with inverter-based resources (IBRs) such as photovoltaic and wind plants. IBRs do not inherently

provide rotational inertia, so the same power imbalance now produces a faster and deeper frequency excursion [1, 2]. Grid-forming converters, and in particular battery energy storage systems (BESS) controlled as virtual synchronous machines (VSMs), can restore part of the lost inertial and damping response [3, 9, 11, 20]. The virtual inertia H_v and damping D_v of such a converter are, however, free design parameters, and selecting them raises two coupled problems.

The first problem is estimation. The amount of virtual inertia that is needed depends on the synchronous inertia that is actually on line, which changes hour by hour with the unit commitment and is difficult to observe directly [15]. The common practice of dividing the power imbalance by the measured RoCoF is sensitive to measurement noise, to the window length and to any fast-acting response, including that of the grid-forming BESS itself. Physics-informed neural networks (PINNs), which add the residual of the governing differential equations to the training loss [5, 6], have been shown to identify power-system dynamic parameters from limited and noisy data [7, 8]. They have rarely been used, though, to supply parameters for the tuning of grid-forming converters.

The second problem is control with guarantees. Adaptive and learning-based tunings of virtual inertia [12, 13] can improve performance, but data-driven controllers that encode stability only through penalty terms provide no assurance that frequency limits will be met in operating conditions that differ from training. In contrast, the safe-learning literature enforces hard constraints through Lagrangian methods [21], constraint-completion layers [17] or safety layers that project actions onto a safe set [16]. For power-system frequency control, the system frequency response (SFR) model [10] offers closed-form, differentiable expressions for RoCoF and nadir that can serve as such constraints.

This paper combines the two ideas. Its contributions are: **(i)** a PINN that identifies synchronous inertia and frequency stiffness from a single measured event while treating the grid-forming BESS power as a known input; **(ii)** a neural tuning policy for H_v and D_v trained with an augmented Lagrangian method under SFR-based constraints on RoCoF, nadir, damping ratio and converter power, tightened by the PINN uncertainty; **(iii)** a safety layer that certifies or corrects each policy output and reports infeasibility; and **(iv)** a closed-loop evaluation on a nonlinear model over 300 random operating points, compared with fixed gains, a penalty-trained policy and the same policy fed by RoCoF-based estimates.

2. Methods

The overall framework is shown in Fig. 1. A physics-informed estimator processes PMU records of a disturbance, a constrained neural policy converts the estimates into VSM gains, and a safety layer certifies the gains before they are sent to the grid-forming BESS.

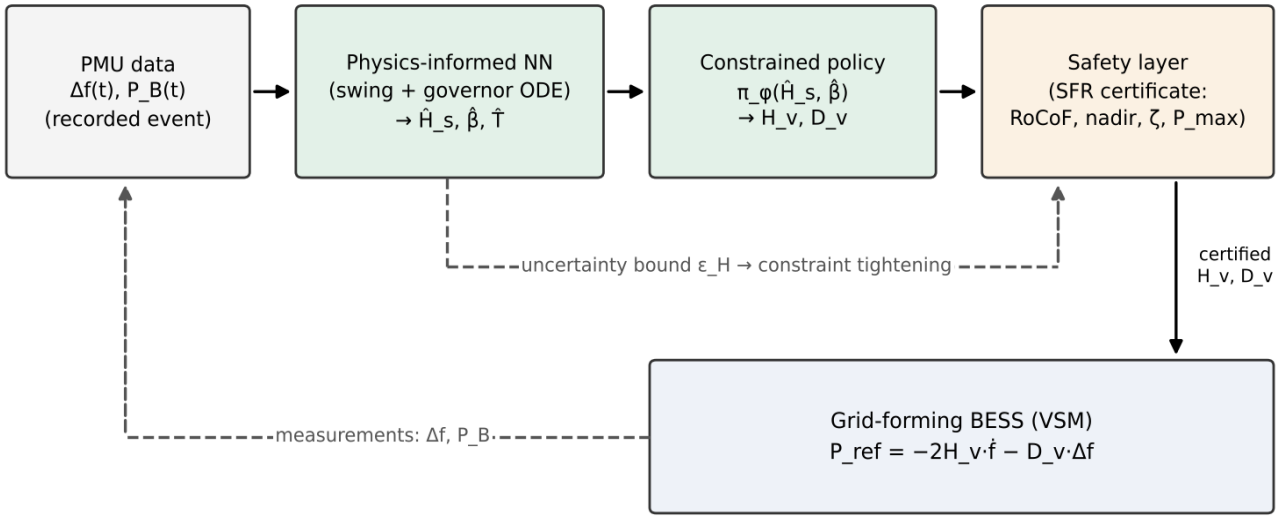


Fig. 1. Proposed framework: physics-informed estimation, stability-constrained policy and safety layer.

2.1. System model

The power system is represented by an aggregated single-area model [10, 18] in per-unit (p.u.) quantities on the system base, with $f_0 = 50$ Hz. For an IBR share η, the synchronous inertia is $H_s = H_0(1 - \eta)$ and the primary-control gain is $K = (1 - \eta)/R$. The nonlinear plant used for data generation and validation is

$$2H_s \cdot d\Delta f/dt = \Delta P_m + \Delta P_B - \Delta P_L - D \cdot \Delta f \quad (1)$$

$$T_g \cdot dx_g/dt = -K \cdot \Delta f - x_g, \quad T_t \cdot d\Delta P_m/dt = x_g - \Delta P_m \quad (2)$$

$$T_c \cdot d\Delta P_B/dt = -(2H_v \cdot d\Delta f/dt + D_v \cdot \Delta f) - \Delta P_B, \quad |\Delta P_B| \le P_{max} \quad (3)$$

where ΔP_L is a step loss of generation, D is the load damping, T_g and T_t are the governor and turbine time constants, and (3) is a VSM-controlled grid-forming BESS with inner-loop time constant T_c and power limit P_{max}. Nominal values are $T_g = 0.2$ s, $T_t = 0.3$ s, $T_c = 20$ ms and $P_{max} = 0.12$ p.u.

2.2. Physics-informed inertia estimation

A fully connected network with one hidden layer of 64 tanh units maps time to the scaled frequency deviation and mechanical power, $[\Delta\hat{f}(t), \Delta\hat{P}_m(t)] = N_\theta(\varphi(t))$. The input transformation $\varphi(\tau) = \ln(1 + \tau/\tau_0)/\ln(1 + 1/\tau_0)$, with $\tau = t/T_w$ and $\tau_0 = 0.03$, stretches the first instants after the event, where the inertial information is concentrated. The physics model is the reduced SFR form of (1)–(2) with a single lag T:

$$r_1 = 2H \cdot d\Delta\hat{f}/dt - \Delta\hat{P}_m + D \cdot \Delta\hat{f} + \Delta P_L - \Delta P_{B,meas}, \quad r_2 = T \cdot d\Delta\hat{P}_m/dt + K \cdot \Delta\hat{f} + \Delta\hat{P}_m \quad (4)$$

The unknown physical parameters (H, D, K, T) are learned jointly with θ in log form. The loss is

$$L = (1/N_d) \cdot \Sigma(\Delta\hat{f} - \Delta f_{meas})^2 + w_r \cdot (1/N_c) \cdot \Sigma(r_1^2 + r_2^2) + w_{ic} \cdot \|[\Delta\hat{f}(0), \Delta\hat{P}_m(0)]\|^2 \quad (5)$$

with $N_d = 501$ PMU samples (50 frames/s over 10 s), $N_c = 600$ collocation points, $w_{ic} = 10$ and w_r ramped from 0 to 1 over the first 30 % of the iterations. Because the measured BESS power enters r_1 as a known input, H estimates the synchronous inertia alone. All gradients, including the time derivatives of the network outputs, were derived analytically and verified against finite differences (maximum relative error $5 \cdot 10^{-8}$). Training used Adam [19] for $4 \cdot 10^4$ iterations. As a baseline, the RoCoF method estimates $H = (\Delta P_L - \bar{P}_B)/(2|\dot{f}|)$, where $\dot{f}$ is the least-squares slope over a 200 ms or 500 ms window. Because D and K are only weakly separable from one event, the identifiable stiffness $\beta = D + K$ is reported.

2.3. Stability certificate

With $T_c \to 0$ the grid-forming BESS adds H_v and D_v to the system, and the frequency deviation obeys the second-order transfer function $-(1 + sT)/(as^2 + bs + c)$, with $a = 2(H_s + H_v)T$, $b = 2(H_s + H_v) + (D + D_v)T$ and $c = \beta + D_v$. From it, the initial RoCoF, the damping ratio and the nadir follow in closed form:

$$RoCoF_0 = f_0 \Delta P_d / (2(H_s + H_v)), \quad \zeta = b / (2\sqrt{(ac)}) \quad (6)$$

$$t_n = \text{atan2}(T\omega_d, \sigma T - 1)/\omega_d, \quad f_{min} = f_0[1 - \Delta P_d \cdot y(t_n)] \quad (7)$$

where $\sigma = b/(2a)$, $\omega_d = \sqrt{(c/a - \sigma^2)}$ and $y(t) = 1/c + e^{(-\sigma t)}(C_1 \cos \omega_d t + C_2 \sin \omega_d t)$, $C_1 = -1/c$, $C_2 = (T/a - \sigma/c)/\omega_d$. $\Delta P_d = 0.1$ p.u. is the design disturbance. The certificate $g(H_s, \beta, H_v, D_v) \le 0$ collects four normalized constraints: $RoCoF_0 \le 0.5$ Hz/s, $f_{min} \ge 49.5$ Hz, $\zeta \ge 0.4$ and $1.12 \cdot \max\{H_v/(H_s + H_v), D_v/c\} \cdot \Delta P_d \le 0.9 P_{max}$. The factor 1.12 accounts for the converter-delay overshoot observed in the nonlinear model. Robustness to estimation error is obtained by evaluating RoCoF, nadir and power at $H_s(1 - \varepsilon_H)$ and $\beta(1 - \varepsilon_\beta)$, and damping at $H_s(1 + \varepsilon_H)$, with $\varepsilon_H = 0.12$ and $\varepsilon_\beta = 0.03$. These bounds cover the largest PINN errors found in Section 3.1.

2.4. Stability-constrained policy learning

The policy π_φ is a multilayer perceptron (2–24–24–2, tanh) that maps normalized $(\hat{H}_s, \hat{\beta})$ to $(H_v, D_v) \in [0, 8\text{ s}] \times [0, 40\text{ p.u.}]$ through sigmoid outputs. Its objective is the storage burden $J = H_v/(H_s + H_v) + D_v/(\beta + D_v)$. The first term is the share of the imbalance carried by the inertial (short) response, and the second is the share carried in steady state until secondary control acts, which dominates the energy drawn from the battery. Training solves

$$\min_\varphi \; E[J] + \Sigma_k \lambda_k \cdot E[\max(0, g_k + \delta)] + (\rho/2) \cdot \Sigma_k E[\max(0, g_k + \delta)^2] \quad (8)$$

over 4000 operating points sampled from $H_0 \in [3, 7]$ s, $\eta \in [0.3, 0.9]$, $R \in [0.04, 0.06]$ and $D \in [0.5, 1.5]$, with mini-batches of 512, $\rho = 5$, a small margin $\delta = 0.03$ and multiplier updates $\lambda_k \leftarrow \lambda_k + \rho \cdot E[\max(0, g_k)]$ every 200 iterations (augmented Lagrangian method [21]). Gradients of (8) with respect to the policy outputs were computed by central differences of the closed-form expressions and back-propagated through the network. A comparison policy was trained identically but with a fixed penalty weight of 0.5 and no multiplier updates, which mimics the reward shaping commonly used in learning-based control.

2.5. Safety layer

At run time each policy output is checked against $g \leq 0$. If the RoCoF constraint is violated, H_v is raised to its closed-form bound $H_v \geq f_0 \Delta P_d / (2 \cdot \text{RoCoF}_{lim}) - H_s(1 - \varepsilon_H)$. If the nadir or damping constraint is violated, D_v is raised by bisection. If no admissible pair exists within the power limit, the layer reports the operating point as infeasible, which means that more storage power is required [16].

2.6. Evaluation

Estimation was tested for $\eta \in \{50, 70, 85\}$ %, i.e., $H_s \in \{2.5, 1.5, 0.75\}$ s with $H_0 = 5$ s and $R = 0.05$, Gaussian PMU noise $\sigma \in \{0, 5, 10, 20\}$ mHz and two noise realizations per case (30 for the RoCoF method). During these events the BESS operated with $H_v = 2$ s and $D_v = 8$ p.u. Control was tested on 300 operating points drawn from the training ranges with a different random seed. The full plant (1)–(3) was simulated for $\Delta P_L = 0.1$ p.u. with SciPy's RK45 solver (maximum step 2 ms). Five schemes were compared: (A) fixed gains $H_v = 3$ s, $D_v = 10$ p.u.; (B) the penalty-trained policy; (C) the constrained policy without the safety layer; (D) the proposed scheme (C + safety layer); and (E) the proposed scheme fed by RoCoF-based (500 ms) estimates instead of PINN estimates. Estimation errors were applied to each operating point according to the statistics in Section 3.1. The reported metrics are the nadir, RoCoF over 500 ms (RoCoF_{500}) and over 100 ms (RoCoF_{100}), saturation of the converter, the peak BESS power and the sustained reserve $D_v/(\beta + D_v) \cdot \Delta P_L$.

3. Results

3.1. Inertia estimation

Table 1 and Fig. 2 summarize the estimation results. The PINN reproduced the true frequency trajectory with an RMSE of 1.3–3.7 mHz, below the measurement noise. The inertia error was −8.6 % to −10.5 % at $\eta = 50$ %, −2.4 % to −3.6 % at $\eta = 70$ % and +6.5 % to +7.3 % at $\eta = 85$ %. These errors changed by less than 2 percentage points as the noise increased from 0 to 20 mHz. The stiffness β was identified within 0.3 % in all cases. The 500 ms RoCoF method had a low variance but a large positive bias (+20.5 %, +28.7 % and +48.5 %), because the governor and the BESS already counteract the imbalance within the window. The 200 ms window halved the bias, but its scatter grew to ±41 % at 20 mHz.

Table 1. Inertia and stiffness estimation (mean of two noise realizations)

η, %	True H_s, s	PINN $\hat{H}_s$, s (σ = 0 / 10 / 20 mHz)	PINN error, %	$\hat{\beta}$ / β, p.u.	RoCoF 200 ms, % (σ = 10)	RoCoF 500 ms, % (σ = 10)
50	2.50	2.284 / 2.263 / 2.238	−8.6 … −10.5	11.01 / 11.0	+15.9 ± 14.8	+20.5 ± 3.4
70	1.50	1.465 / 1.456 / 1.445	−2.4 … −3.6	7.01 / 7.0	+21.2 ± 12.1	+28.6 ± 3.1
85	0.75	0.805 / 0.801 / 0.799	+6.5 … +7.3	4.01 / 4.0	+38.0 ± 11.0	+48.5 ± 3.1

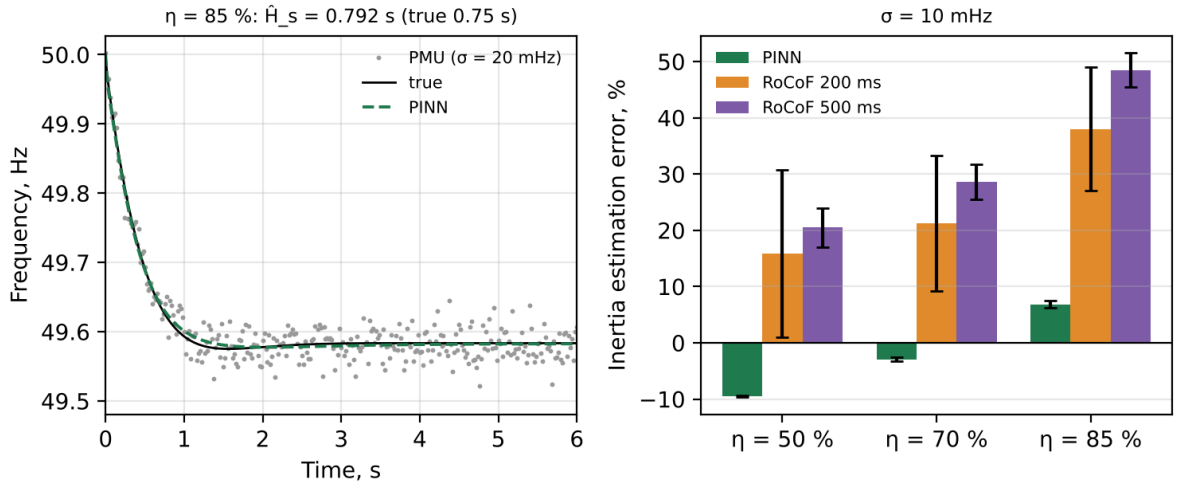


Fig. 2. Left: PINN reconstruction of a noisy event at η = 85 %. Right: inertia estimation error of the PINN and the RoCoF method (σ = 10 mHz; bars show ±1 standard deviation).

3.2. Policy training

Fig. 3 compares the training of the two policies. With the augmented Lagrangian method, the share of training points that violate the tightened certificate fell from 96 % to 0 % within 2000 iterations and remained at zero. The multipliers settled on the RoCoF and nadir constraints ($\lambda \approx 0.75$ and 1.1), whereas the damping constraint never became active and the power constraint was active only near the boundary. The fixed-penalty policy violated the certificate at 83–100 % of the points throughout training. On 10 000 fresh in-distribution points, the constrained policy satisfied the certificate everywhere, so the safety layer made no corrections.

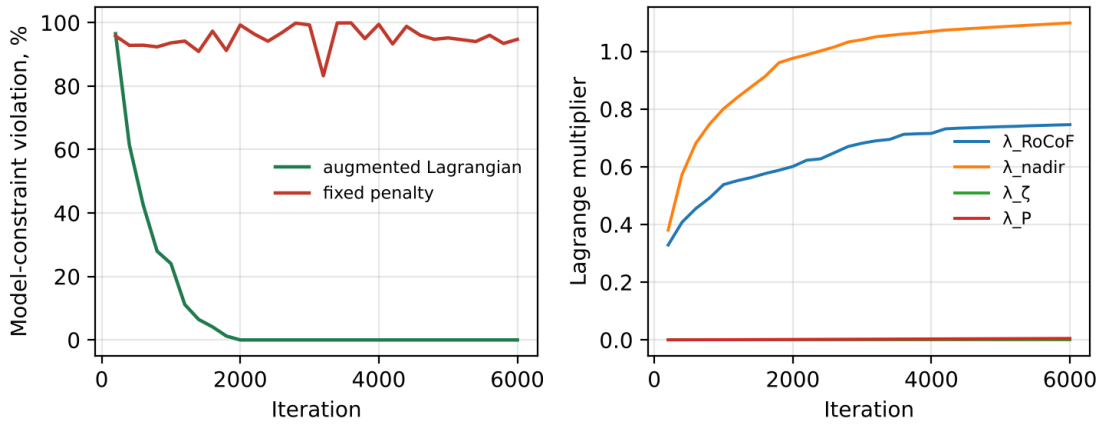


Fig. 3. Training history: share of points violating the model certificate (left) and Lagrange multipliers of the constrained policy (right).

The learned policy is physically interpretable (Fig. 4). H_v decreases almost linearly with $\hat{H}$_s, following the RoCoF bound in (6) with the uncertainty margin. D_v is zero for $\beta \gtrsim 10$ p.u. and rises as β falls. This is consistent with the necessary condition $c = \beta + D_v \geq f_0 \Delta P_d / \Delta f_{lim} = 10$ p.u., which the steady-state frequency must satisfy for the nadir to stay above 49.5 Hz.

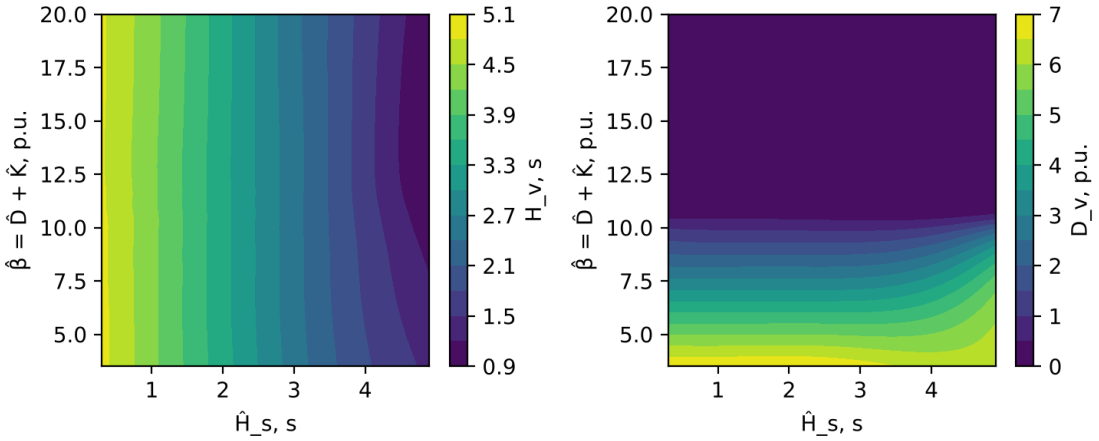


Fig. 4. Virtual inertia (left) and damping (right) chosen by the constrained policy as functions of the estimated inertia and stiffness.

3.3. Closed-loop validation on the nonlinear model

Table 2 and Fig. 5 report the results for the 300 test operating points. The proposed scheme (D) met the nadir and $RoCoF_{500}$ limits at every point, with a worst-case nadir of 49.52 Hz and a maximum $RoCoF_{500}$ of 0.44 Hz/s. Compared with the fixed gains (A), it reduced the mean sustained reserve from 5.34 % to 2.09 % (−61 %) and the mean peak BESS power from 6.82 % to 6.34 %. The penalty-trained policy (B) used the least reserve but violated the limits at 59.7 % of the points, with nadirs as low as 49.13 Hz. The fixed gains passed the nadir limit everywhere because D_v = 10 p.u. is generous, but they exceeded the $RoCoF_{500}$ limit at 1.3 % of the points.

Table 2. Closed-loop results on 300 random operating points (ΔP_L = 0.1 p.u.)

Scheme	Violations, %	Min nadir, Hz	Max $RoCoF_{500}$, Hz/s	$RoCoF_{100}$ > 0.5, %	Mean P_pk, %	Sustained reserve, %	Certificate violated (true params), %
A. Fixed gains	1.3	49.60	0.52	61.3	6.82	5.34	52.7
B. Penalty policy	59.7	49.13	0.59	91.7	5.86	0.97	62.3
C. Constrained policy	0	49.52	0.44	58.0	6.34	2.09	0
D. Proposed (C + SL)	0	49.52	0.44	58.0	6.34	2.09	0
E. Proposed, RoCoF est.	0	49.51	0.49	98.7	5.87	2.09	64.3

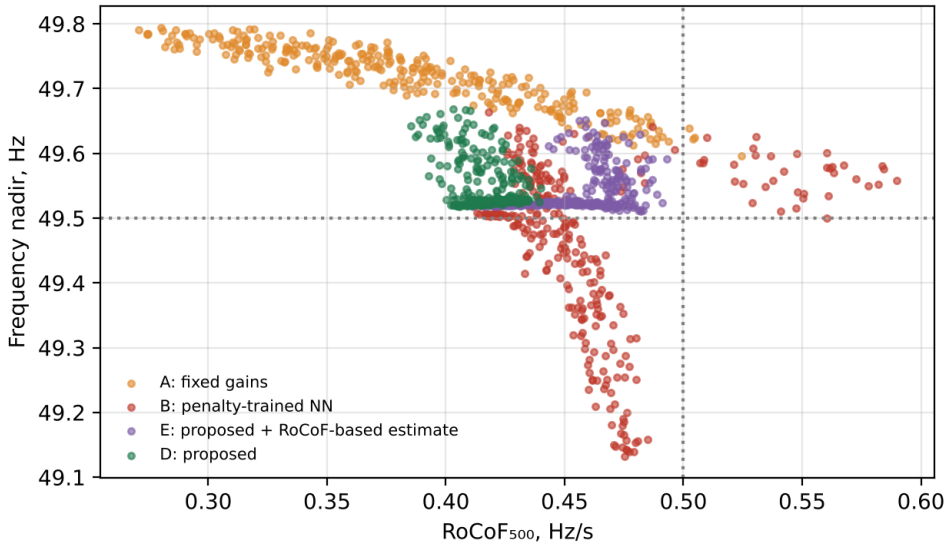


Fig. 5. Frequency nadir versus $RoCoF_{500}$ on the nonlinear model for 300 operating points; dotted lines mark the limits.

Scheme E shows the role of the estimator. Fed by the biased RoCoF-based estimates, the same policy and safety layer chose on average 16 % less virtual inertia (2.92 s versus 3.47 s). It still passed the $RoCoF_{500}$ limit, but with almost no margin (maximum 0.49 Hz/s). Moreover, its outputs violated the certificate evaluated with the true parameters at 64.3 % of the points, which means that the formal guarantee was lost. The 100 ms RoCoF, which is sensitive to the converter delay that the certificate neglects, exceeded 0.5 Hz/s at 58 % of the points for scheme D, but by at most 0.048 Hz/s (mean 0.503 Hz/s). For schemes A, B and E the maxima were 0.81, 0.66 and 0.58 Hz/s.

3.4. Out-of-distribution operating points

To test the safety layer outside the training range, 500 points with $H_0 \in [2, 3]$ s and $\eta \in [0.90, 0.95]$ were evaluated ($H_s = 0.10$–0.29 s). The safety layer resolved every RoCoF and nadir conflict, but at 67.6 % of these points no admissible gains existed within $P_{max} = 0.12$ p.u., because the inertial response alone would require more than 90 % of the design imbalance. The layer reported these points as infeasible instead of silently returning unsafe gains. Such a report tells the operator that additional storage power or a reduced design contingency is required.

4. DISCUSSION

The results support the central premise of this work: when a learning-based controller tunes grid-forming storage, the quality of its physical inputs matters as much as the controller itself. The PINN reduced the inertia-estimation error from 20–49 % (RoCoF method) to below 11 %, consistent with earlier PINN identification studies [7, 8]. This made it possible to use a modest, explicit uncertainty margin. With the RoCoF-based inputs, the same certified policy remained compliant only because the 500 ms metric is less strict than the instantaneous RoCoF in the certificate, and the guarantee did not hold. The PINN errors were systematic rather than random: negative at high inertia and positive at low inertia. The likely cause is the mismatch between the single-lag physics model and the two-stage governor of the plant. A richer physics model, or the reporting of PINN uncertainty through ensembles, would allow ε_H to be reduced.

The comparison between schemes B and D illustrates why penalty-based learning is insufficient for security-critical tuning. Both policies share the architecture, data and objective. Only the treatment of constraints differs, yet the penalty policy traded frequency security for a smaller reserve at most operating points. The augmented Lagrangian method found multipliers that make the certificate binding, and the safety layer turns the remaining model guarantee into a run-time check that also detects infeasibility. Because the certificate uses closed-form SFR expressions, both training and checking are computationally light and the policy remains interpretable (Fig. 4).

The study has limitations. (i) The certificate assumes an ideal converter ($T_c \to 0$), so the 100 ms RoCoF can slightly exceed the limit. Including T_c in the certificate or adding a delay-dependent margin would address this. (ii) The aggregated single-area model omits network

topology, inter-area oscillations, voltage dynamics and the location of storage [14, 22]. (iii) D was assumed nominal in the damping term because only β is identifiable from one event. (iv) State of charge and degradation were not modelled. (v) All data were synthetic, and no hardware validation has yet been performed. Future work will extend the PINN to multi-machine models and ensemble uncertainty, include the converter delay in the certificate, and test the framework on hardware-in-the-loop platforms.

5. CONCLUSIONS

A physics-informed and stability-constrained framework for tuning grid-forming battery storage in low-inertia power systems was developed and evaluated. The main findings are as follows:

1. The PINN, which embeds the swing and governor equations and treats the measured BESS power as an input, estimated synchronous inertia within −10.5 % to +7.3 % and frequency stiffness within 0.3 % from one event, and was almost insensitive to PMU noise up to 20 mHz. The conventional RoCoF method was biased by +20 % to +49 %.
2. A neural policy trained with an augmented Lagrangian method under closed-form SFR constraints satisfied the uncertainty-tightened certificate at all 10 000 in-distribution test points and learned a physically interpretable mapping.
3. On a nonlinear model, the proposed scheme met the 49.5 Hz nadir and 0.5 Hz/s RoCoF limits at all 300 test points and reduced the sustained storage reserve by 61 % relative to fixed gains. A penalty-trained policy violated the limits at 59.7 % of the points.
4. The safety layer detected operating points at which no admissible gains exist within the converter rating, which turns the controller into a tool for assessing whether the installed storage is sufficient.

Declarations

Funding: this research received no external funding. Conflict of interest: the authors declare no conflict of interest. Data and code availability: the simulation code and data used in this study are available from the corresponding author on reasonable request. Author contributions: O. A. Kayumov — conceptualization, methodology, software, formal analysis, writing – original draft; N. R. Kayumova — validation, visualization, data curation, writing – review and editing.

References

1. F. Milano, F. Dörfler, G. Hug, D. J. Hill, G. Verbič, "Foundations and challenges of low-inertia systems," in Proc. Power Systems Computation Conference (PSCC), Dublin, 2018, pp. 1–25.
2. F. Dörfler, D. Groß, "Control of low-inertia power systems," Annual Review of Control, Robotics, and Autonomous Systems, vol. 6, pp. 415–445, 2023.
3. Y. Lin, J. H. Eto, B. B. Johnson et al., "Research Roadmap on Grid-Forming Inverters," NREL/TP-5D00-73476, National Renewable Energy Laboratory, Golden, CO, 2020.
4. N. Hatziargyriou et al., "Definition and classification of power system stability – revisited & extended," IEEE Trans. Power Syst., vol. 36, no. 4, pp. 3271–3281, 2021.
5. M. Raissi, P. Perdikaris, 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, 2019.
6. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, "Physics-informed machine learning," Nature Reviews Physics, vol. 3, pp. 422–440, 2021.
7. G. S. Misyris, A. Venzke, S. Chatzivasileiadis, "Physics-informed neural networks for power systems," in Proc. IEEE PES General Meeting, 2020.
8. J. Stiasny, G. S. Misyris, S. Chatzivasileiadis, "Physics-informed neural networks for non-linear system identification for power system dynamics," in Proc. IEEE Madrid PowerTech, 2021.
9. U. Tamrakar, D. Shrestha, M. Maharjan, B. P. Bhattarai, T. M. Hansen, R. Tonkoski, "Virtual inertia: Current trends and future directions," Applied Sciences, vol. 7, no. 7, art. 654, 2017.
10.] P. M. Anderson, M. Mirheydar, "A low-order system frequency response model," IEEE Trans. Power Syst., vol. 5, no. 3, pp. 720–729, 1990.
11. Q.-C. Zhong, G. Weiss, "Synchronverters: Inverters that mimic synchronous generators," IEEE Trans. Ind. Electron., vol. 58, no. 4, pp. 1259–1267, 2011.
12. J. Alipoor, Y. Miura, T. Ise, "Power system stabilization using virtual synchronous generator with alternating moment of inertia," IEEE J. Emerg. Sel. Topics Power Electron., vol. 3, no. 2, pp. 451–458, 2015.

13. B. K. Poolla, D. Groß, F. Dörfler, "Placement and implementation of grid-forming and grid-following virtual inertia and fast frequency response," IEEE Trans. Power Syst., vol. 34, no. 4, pp. 3035–3046, 2019.
14. U. Markovic, O. Stanojev, P. Aristidou, E. Vrettos, D. Callaway, G. Hug, "Understanding small-signal stability of low-inertia systems," IEEE Trans. Power Syst., vol. 36, no. 5, pp. 3997–4017, 2021.
15. P. Tielens, D. Van Hertem, "The relevance of inertia in power systems," Renewable and Sustainable Energy Reviews, vol. 55, pp. 999–1009, 2016.
16. G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, Y. Tassa, "Safe exploration in continuous action spaces," arXiv:1801.08757, 2018.
17. P. L. Donti, D. Rolnick, J. Z. Kolter, "DC3: A learning method for optimization with hard constraints," in Proc. Int. Conf. Learning Representations (ICLR), 2021.
18. P. Kundur, Power System Stability and Control. New York: McGraw-Hill, 1994.
19. D. P. Kingma, J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.
20. Q. Salem, B. Bany Fawaz, R. Aljarrah, "Grid forming converters for low inertia systems – capabilities and limitations: A critical review," IEEE Open J. Ind. Electron. Soc., 2025.
21. D. P. Bertsekas, Constrained Optimization and Lagrange Multiplier Methods. New York: Academic Press, 1982.
22. M. Paolone et al., "Fundamentals of power systems modelling in the presence of converter-interfaced generation," Electric Power Systems Research, vol. 189, art. 106811, 2020.