

Original Article

Artificial Intelligence

Design and Development of an AI-Based Load Frequency Control System for a Power Transmission Grid Integrated with a Solar PV Microgrid

Mohammad Ali* 

School of Science and Engineering, Atlantic International University, Honolulu, HI 96813, USA

Abstract

Displacement of synchronous generation by inverter-interfaced solar photovoltaic (PV) plant reduces inertia and strains fixed-gain secondary load frequency control. This paper develops a supplementary AI-based scheme for a two-area reheat-thermal transmission grid hosting a 40 MW PV plant and a co-located 20 MW / 8.3 MWh battery. It adds to the installed secondary loop feed-forward from the measured PV deviation, local co-dispatch of the battery with a washout and a restoring state-of-charge term, and a bounded fuzzy gain schedule. Governor dead-band and generation rate constraint are enforced with anti-windup, and every controller, proposed and benchmark (PI, PID and a Mamdani fuzzy-PID), is tuned by the same particle swarm procedure on scenarios disjoint from the tests. Under a composite load step and cloud-induced PV loss the scheme reduces the integral of time-weighted absolute error by 15.2 % relative to PI and by 5.5 % relative to the best tuned benchmark. A re-tuned ablation shows that this improvement comes mainly from coordinated battery dispatch and PV feed-forward; the fuzzy gain schedule contributes no measurable performance and is retained only to bound the gain. Noise and jitter are tolerable, but a 1 s communication delay raises the error of ideally tuned designs by up to 110.6 %; re-tuning for the channel with local battery response makes the proposed scheme the most delay-tolerant design. A 30-minute study shows that the battery law holds the state of charge within its window; the other results concern short-term frequency regulation. A residual deep reinforcement learning (TD3) controller chattered without regularisation; a spectral-norm Lipschitz bound removed the chatter consistently, but no learned controller managed the battery's energy over 30 minutes. Results are from simulation.

Keywords: Load frequency control, Area control error, Solar photovoltaic generation, Battery energy storage, Fuzzy gain scheduling, Particle swarm optimisation, Deep reinforcement learning, Communication delay, Low-inertia power systems

INTRODUCTION

Frequency is the most directly observable indicator of the balance between generation and demand in a synchronously interconnected power system. Any mismatch between mechanical power delivered to the turbine shafts and electrical demand appears at once as acceleration or deceleration of the rotating

masses, and hence as a frequency deviation (Kundur, 1994). Control of the balance is hierarchical: primary control acts through governor droop within seconds but leaves an offset; secondary control, Load Frequency Control (LFC) proper, removes that offset and restores scheduled tie-line interchange over tens of seconds to minutes; tertiary control performs economic redispatch (Bevrani,

Correspondence:

Mohammad Ali, School of Science and Engineering, Atlantic International University, Honolulu, HI 96813, USA, ORCID: 0009-0001-2184-7978, Email: malixbd@ieee.org

Received Dates: Oct 04, 2026;

Accepted Date: Oct 09, 2026;

Published Date: Oct 11, 2026:

and multi-area LFC (Yan & Xu, 2019, 2020) and are reviewed by Chen, Qu, Tang, Low, and Li (2022). Learned policies are prone to high-frequency action oscillation; remedies include temporal and spatial smoothness regularisation (Mysore, Mabsout, Mancuso, & Saenko, 2021), Lipschitz bounds by spectral normalisation (Miyato, Kataoka, Koyama, & Yoshida, 2018), and residual formulations that learn a correction to an existing controller (Johannink et al., 2019).

Delay, Storage and the Gap Addressed

Communication delay in the AGC channel reduces the stability margin of the secondary loop, and delay-dependent stability criteria have been derived for LFC (Jiang, Yao, Wu, Wen, & Cheng, 2012; Zhang, Jiang, Wu, He, & Wu, 2013). Battery storage complements variable generation because its sub-second response covers the interval in which thermal plant is slowest (Datta, Kalam, & Shi, 2019), and grid-forming or virtual-inertia converter controls address the primary timescale (Bevrani, Ise, & Miura, 2014; Rocabert, Luna, Blaabjerg, & Rodríguez, 2012). Coordination of storage with secondary control is usually left to separate rule-based logic. The gap addressed here is a bounded supplementary scheme that co-dispatches storage and thermal generation from signals available locally and centrally, evaluated

against equally tuned benchmarks and an equally constrained learned policy, under the non-idealities a field deployment would meet. Table 1 summarises the comparison.

SYSTEM MODELLING

Parameters are listed in Table 2, in per unit on a 2000 MW area base unless stated otherwise.

Generator, Load, Governor and Turbine

The rotational dynamics of an aggregated generating unit follow the swing equation; in per-unit form with inertia constant H (Kundur, 1994),

$$2H/f_0 \cdot d(\Delta f)/dt = \Delta P_m - \Delta P_e - D \cdot \Delta f(1)$$

which gives the power-system transfer function

$$G_{ps}(s) = K_{ps}/(1 + sT_{ps}), \quad K_{ps} = 1/D, \quad T_{ps} = 2H/(f_0 D)(2)$$

As inverter-based generation displaces synchronous plant H falls, so T_{ps} falls and the response to a given imbalance becomes faster and larger. In the penetration study T_{ps} is scaled with the PV penetration ρ as

$$T_{ps}(\rho) = T_{ps0} \cdot (1 - \rho)/(1 - \rho_0)(3)$$

Governor, reheat turbine and droop are represented on the secondary timescale by

Table 1: Comparison of secondary load frequency control strategies.

Control paradigm	Adaptivity	Model dependence	Constraint handling	Actuator activity reported	Principal limitation
Classical integral / PI	None	Low	Post-hoc clipping	Rarely	Fixed gains; decays off design point
LQR / optimal	None	High	Not native	Rarely	Full state feedback, exact model
Sliding mode	Structural	Moderate	Not native	Sometimes	Chattering unsuited to governors
Metaheuristic PID	None (offline)	Moderate	Post-hoc clipping	Rarely	Optimal only at the design point
Model predictive control	Via model update	Very high	Native, explicit	Sometimes	Model identification; solver burden
Fuzzy / neuro-fuzzy	Static or offline	Low	Via rule design	Rarely	Rule growth with input dimension
DDPG / TD3 as reported	Online, continuous	None	Not native	Seldom	Output unbounded; interval often unequal
This work: proposed scheme	Scheduled, bounded	Low	Saturation, rate limit, anti-windup	Yes (Table 6)	Gains tied to storage rating (Sec. VI-I)
This work: residual TD3	Learned, bounded trims	None	Same limiter and anti-windup	Yes (Tables 7, 9)	Chatter unless Lipschitz-bounded; no energy management (Sec. VI-F, VI-H)

Table 2: Plant, PV, storage and controller parameters.

Quantity	Symbol	Value
Rated area capacity / base	P_r	2000 MW
Nominal frequency	f_0	50 Hz
Governor / turbine time constants	T_g / T_t	0.08 s / 0.3 s
Reheater time constant and gain	T_r / K_r	10 s / 0.33
Speed regulation	R	2.4 Hz/p.u.MW
Frequency bias	B	0.425 p.u.MW/Hz
Power-system gain / time constant	K_{ps} / T_{ps0}	120 Hz/p.u.MW / 20 s
Tie-line coefficient	T_{12}	0.0866 p.u.MW/rad
Governor dead-band	DB	0.036 Hz
Generation rate constraint	δ_{GRC}	0.0005 p.u.MW/s (3 %/min)
Governor command limit / integrator clamp	u_{max} / l_{max}	$\pm 0.5 / \pm 1$
Control interval (all controllers)	T_s	0.10 s
Derivative filter time constant	T_f	0.2 s
PV plant rating (nominal penetration)	$P_{pv,sched}$	40 MW (0.02 p.u., 2 %)
PV equivalent lag	T_{pv}	1.8 s
BESS rated power / energy	$P_{BESS,max} / E_{rated}$	20 MW / 8.3 MWh
BESS converter lag	T_b	0.05 s
SOC window / initial and reference	SOC	22–88 % / 60 %
PV-to-storage washout time constant	T_w	30 s
SOC restoring gain	K_{b3}	$2.5e-4$ p.u./%




Figure 2: Linearised block model of one control area. The governor dead-band and the generation rate constraint of Section III-C act on the droop path and on the mechanical power respectively.




Figure 3: (a) Single-diode equivalent circuit. (b) Current–voltage and power–voltage characteristics from (10)–(11) at four irradiance levels.




Figure 4: Structure of the proposed scheme. Governor setpoints are formed centrally and sent over the AGC link; the storage command is formed locally from Δf , ΔP_{pv} and SOC.

$$G_g(s) = 1/(1 + sT_g), \quad G_{rt}(s) = (1 + sK_rT_r)/[(1 + sT_t)(1 + sT_r)] \quad (4)$$

with primary control entering as the static feedback $1/R$ through the governor dead-band of Section III-C.

Tie-Line and Area Control Error

$$\Delta P_{tie}(s) = (2\pi T_{12}/s) \cdot [\Delta f_1(s) - \Delta f_2(s)] \quad (5)$$

$$ACE_1 = B_1 \Delta f_1 + \Delta P_{tie}, \quad ACE_2 = B_2 \Delta f_2 - \Delta P_{tie} \quad (6)$$

with frequency bias $B_i = D_i + 1/R_i$, which makes the areas non-interacting (Cohn, 1957). Figure 2 shows the block model of one area.

Retained Nonlinearities and Anti-Windup

Two physical nonlinearities are enforced in every simulation reported. The governor dead-band removes the response to small speed deviations, and the generation rate constraint limits the rate of change of delivered mechanical power:

$$\Delta f_{db} = \text{sgn}(\Delta f) \cdot \max(0, |\Delta f| - DB), \quad DB = 0.036 \text{ Hz} \quad (7)$$

$$|d(\Delta P_m)/dt| \leq \delta_{GRC} = 0.0005 \text{ p.u. MW/s (3 \%/min)} \quad (8)$$

With the rate constraint active the plant cannot follow a large command, and unconditional integration of the ACE winds the integrator up; Section VI-E shows that without protection every tuned design is left in a sustained oscillation. Every controller therefore uses clamping anti-windup (Kothare, Campo, Morari, & Nett, 1994): integration is inhibited while the command is saturated, or while the rate constraint is binding, in the direction the error would push further;

$$\int ACE_i \leftarrow \int ACE_i + ACE_i T_s \text{ unless } (s_{u_i} \cdot ACE_i > 0) \text{ or } (-s_{g_i} \cdot ACE_i > 0) \quad (9)$$

where s_{u_i} and s_{g_i} are the signs of active command saturation and active rate limiting (zero when inactive), and the integrator is additionally clamped to $\pm I_{\max}$.

Photovoltaic Plant

The PV cell is represented by the single-diode model of Figure 3(a) (Villalva, Gazoli, & Filho, 2009),

$$I_{pv} = I_{ph} - I_0 \{ \exp[q(V_{pv} + I_{pv} R_s)/(nkT_c N_s)] - 1 \} - (V_{pv} + I_{pv} R_s)/R_{sh} \quad (10)$$

$$I_{ph} = [I_{sc} + K_i(T_c - T_{STC})] \cdot G/G_{STC}, \quad T_c = T_{amb} + G \cdot (NOCT - 20)/800 \quad (11)$$

On the secondary timescale the array, maximum-power-point tracker (Esram & Chapman, 2007) and inverter are aggregated into a first-order lag between available and delivered power,

$$dP_{pv}/dt = (P_{av} - P_{pv})/T_{pv}, \quad \Delta P_{pv} = P_{pv} - P_{pv,sched} \quad (12)$$

so a step loss of irradiance appears at the grid as a ramp. Cloud events are imposed as step reductions of P_{av} .

Battery Energy Storage

The converter follows its command through a first-order lag, and the SOC integrates the delivered power:

$$dP_{BESS}/dt = (u_b - P_{BESS})/T_b \quad (13)$$

$$dSOC/dt = -(100/E_{rated}) \cdot P_{BESS} \quad (14)$$

with $P_{BESS} > 0$ denoting discharge. The power limit $|u_b| \leq P_{BESS,max}$ is enforced by saturation, and the energy window $SOC_{min} \leq SOC \leq SOC_{max}$ by inhibiting further discharge at SOC_{min} and further charge at SOC_{max} .

F. State-Space Form, Controllability and Observability

Assembling the subsystems gives $\dot{x} = Ax + Bu + \Gamma w$ with

$$x = [\Delta f_1, \Delta P_{t1}, \Delta P_{r1}, \Delta X_{g1}, \int ACE_1, \Delta f_2, \Delta P_{t2}, \Delta P_{r2}, \Delta X_{g2}, \int ACE_2, \Delta P_{tie}, \Delta P_{pv}, P_{BESS}]^T \quad (15)$$

$$u = [\Delta u_1, \Delta u_2, u_b]^T, \quad w = [\Delta P_{L1}, \Delta P_{L2}, \Delta P_{av}]^T \quad (16)$$

The controllability and observability of the linearised model (nonlinearities removed) were computed in two independent ways, from the dimension of the controllable (observable) subspace obtained by an orthogonalised Krylov iteration, which remains well conditioned although the eigenvalues span two orders of magnitude, and by the Popov-Belevitch-Hautus test at every eigenvalue; the two agree. The controllability matrix has rank 12 of 13: the PV state is driven only by irradiance and no input can move it. Its mode, at $-1/T_{pv} = -0.556 \text{ s}^{-1}$, is stable, so the pair (A, B) is stabilisable. With the physically measured outputs $y = [\Delta f_1, \Delta f_2, \Delta P_{tie}, \Delta P_{pv}, P_{BESS}]^T$ the observability matrix has rank 11; the two unobservable modes, both at the origin, are the ACE integrators, which are controller-internal and known exactly. Adding them to y gives rank 13. The open-loop spectrum contains a lightly damped inter-area pair at $-0.210 \pm j2.841 \text{ s}^{-1}$ (0.45 Hz, damping ratio 0.074), which is the mode excited by tie-line disturbances in Section VI.

CONTROLLERS

All controllers are digital, sampled every $T_s = 0.1 \text{ s}$ with zero-order hold, and share the anti-windup of (9) and the command limits of Table 2. Derivative terms use the ACE derivative filtered with time constant T_f , so that measurement noise is not amplified without bound.

Proposed Supplementary Scheme

The scheme is supplementary: it retains the integral action that guarantees zero steady-state ACE and adds bounded terms to it. Figure 4 shows the structure. For area i the governor command is

$$u_i = \text{sat}[-(\kappa(ACE_i) \cdot ACE_i + K_I \int ACE_i + K_D \cdot dACE_i/dt) - K_{ff} \Delta P_{pv}] \quad (17)$$

where the feed-forward term is applied only in the area hosting the PV plant. ΔP_{pv} is the measured deviation of delivered PV power from schedule; no forecasting model is used in this work.

Fuzzy Gain Schedule

The proportional gain is scheduled on $|ACE|$ by a two-rule, zero-order Takagi-Sugeno system (Takagi & Sugeno, 1985): IF $|ACE|$ is SMALL THEN $\kappa = \kappa_0$; IF $|ACE|$ is LARGE THEN $\kappa = \kappa_0 + \Delta\kappa$, with complementary memberships $\mu_L = \tanh(\alpha|ACE|)$ and $\mu_S = 1 - \mu_L$. The weighted-average output

$$\kappa = [\mu_S \kappa_0 + \mu_L (\kappa_0 + \Delta\kappa)] / (\mu_S + \mu_L) = \kappa_0 + \Delta\kappa \cdot \tanh(\alpha|ACE|) \quad (18)$$

is exactly the closed form evaluated in simulation. Its three parameters (κ_0 , $\Delta\kappa$, α) are identified by PSO together with the other gains (Section V-B); they are not trained by the hybrid learning rule of an adaptive neuro-fuzzy inference system (ANFIS) (Jang, 1993), and the element is therefore described here as a fuzzy gain schedule rather than as ANFIS. The schedule is bounded by construction, $\kappa \in [\kappa_0, \kappa_0 + \Delta\kappa]$, and Section VI-D shows that it contributes almost nothing to performance.

C. Storage Co-Dispatch from Local Measurements

Because the BESS is co-located with the PV plant, its command can be formed from local measurements without communication delay:

$$u_b = \text{sat}[-K_{b1} \Delta f_1 - K_{b2} \Delta P_{pv}^{hp} + K_{b3} (SOC - SOC_{ref})] \quad (19)$$

$$\Delta P_{pv}^{hp} = \Delta P_{pv} - \Delta P_{pv}^{lp}, \quad T_w \cdot d(\Delta P_{pv}^{lp})/dt = \Delta P_{pv} - \Delta P_{pv}^{lp} \quad (20)$$

The first term is fast frequency support. The second cancels the PV excursion before it propagates into frequency; the washout (20) confines storage to the fast part of a PV change and hands

a sustained deficit to the governors, which receive the full ΔP_{pv} through K_{ff} in (17). The third term restores the SOC. Since the SOC falls when $P_{BESS} > 0$, a restoring term must command charging (negative power) when SOC is below its reference, which requires the positive sign in (19); with the opposite sign the term is positive feedback, and Section VI-H shows the consequence. K_{b3} is fixed by design so that the SOC recovers with a time constant of

$$\tau_{SOC} = E_{rated}/(100 \cdot K_{b3}) \approx 10 \text{ min} \quad (21)$$

long compared with any frequency transient. The washout time constant $T_w = 30 \text{ s}$ is chosen to exceed the duration of cloud transients (seconds) and to be shorter than sustained shading (minutes), and is comparable to the time the rate-limited thermal units need to absorb a deficit of the size considered. Figure 5

Benchmark Controllers

Three benchmarks are used: PI and PID acting on the ACE, and a fuzzy-PID. The fuzzy-PID is a Mamdani fuzzy PD stage (Mamdani & Assilian, 1975) with seven triangular sets on each normalised input (scaled by K_e and K_{de} and saturated at ± 1), a 49-rule MacVicar-Whelan diagonal rule table (MacVicar-Whelan, 1976), min-max inference and centroid defuzzification on an 81-point grid, in parallel with integral action:

$$u_i = -[K_u \cdot FPD(K_e \cdot ACE_i, K_{de} \cdot dACE_i/dt) + K_I \int ACE_i] \quad (22)$$

All three benchmarks have access to the same storage asset under local frequency droop, $u_b = -K_{bd} \Delta f_1$, with K_{bd} tuned along with their other gains. This isolates the value of co-dispatch from the value of owning a battery.

E. Residual Deep Reinforcement Learning Variant

A TD3 agent (Fujimoto et al., 2018) supplies bounded trims on top of the PSO-tuned PI loop and its storage droop, a residual formulation (Johannink et al., 2019):

$$u_i = u_i^{PI} + 0.03 \cdot a_i, \quad u_b = -K_{bd} \Delta f_1 + P_{BESS, max} \cdot a_3, \quad a \in [-1, 1]^3 \quad (23)$$

so that a zero action reproduces the tuned PI baseline exactly. The trim authority of $\pm 0.03 \text{ p.u.}$ is three times the test load step; a much larger authority ($\pm 0.30 \text{ p.u.}$) was found in preliminary training to make most of the action space destructive. The observation contains only measurable quantities,

$$s_t = [\Delta f_1, \Delta f_2, \Delta P_{tie}, ACE_1, ACE_2, \int ACE_1, \int ACE_2, dACE_1/dt, dACE_2/dt, \Delta P_{pv}, SOC - SOC_{ref}, d(SOC - SOC_{ref})/dt] \quad (24)$$

$$dt, \Delta P_{pv}, SOC - SOC_{ref}, d(SOC - SOC_{ref})/dt] \quad (24)$$

where the PV ramp rate is the filtered derivative of the measured PV deviation (time constant 0.5 s). No element contains information about the future, so the agent sees exactly what a deterministic controller could measure. The reward is

$$r_t = -[w_1(\Delta f_1^2 + \Delta f_2^2) + w_2 \Delta P_{tie}^2 + w_3(ACE_1^2 + ACE_2^2) + w_4 \|a_t\|^2 + w_5 \|a_t - a_{t-1}\|^2 + w_6 (SOC - SOC_{ref})^2] \quad (25)$$

with $w = (120, 70, 45, 0.012, 0.06, 2 \times 10^{-5})$. The actor's output layer is initialised near zero so that training starts at the baseline, and the initial policy is retained as a candidate during selection.

F. Remedies for Actuator Chatter

Because the penalty on action increments in (25) shapes the return but does not constrain the mapping from state to action, four further remedies are evaluated. Two act on the trained policy at run time: a first-order low-pass filter and a per-step rate limiter on each action channel,

$$\tilde{a}_t = \tilde{a}_{t-1} + [T_s/(\tau_a + T_s)](a_t - \tilde{a}_{t-1}), \quad \tilde{a}_t = \tilde{a}_{t-1} + \text{sat}_{\{\pm \Delta a\}}(a_t - \tilde{a}_{t-1}) \quad (26)$$

Two act during training. Conditioning for action policy smoothness (CAPS) (Mysore et al., 2021) adds temporal and spatial penalties to the actor objective,

$$L_\mu = -Q(s, \mu(s)) + \lambda_T \|\mu(s_t) - \mu(s_{t+1})\|^2 + \lambda_S \|\mu(s) - \mu(s + \epsilon)\|^2, \quad \epsilon \sim N(0, 0.05^2 I) \quad (27)$$

where the spatial term is a local Lipschitz penalty. A hard global bound is obtained by spectral normalisation (Miyato et al., 2018), rescaling every actor weight matrix after each update so that its largest singular value does not exceed c:

$$W_i \leftarrow W_i \cdot \min(1, c/\sigma_{\max}(W_i)) \quad \square \quad \|\mu(s) - \mu(s')\| \leq c^3 \|s - s'\| \quad (28)$$

since ReLU and tanh are 1-Lipschitz. Here $\lambda_T = \lambda_S = 1$ and $c = 1$ on the normalised state.

SIMULATION STUDY

Implementation

The plant of Section III and all controllers are implemented once, in a single module, and integrated by explicit Euler at a 0.01 s step. Each controller sample is processed by one routine that draws measurement noise, forms the commands, queues the governor setpoints on the delayed AGC link, applies the local storage command and advances the plant to the next sampling instant. The reinforcement learning environment calls the same routine,




Figure 5: Scheduled proportional gain (18) with the PSO-identified parameters, and the constant gain identified when scheduling is removed and the scheme re-tuned (Section VI-D).

so with a zero action it reproduces the tuned PI benchmark exactly; this was checked, bit for bit, with and without noise, jitter and communication delay before any agent was trained. The agent, including back-propagation and the Adam optimiser (Kingma & Ba, 2015), is written in NumPy without a deep learning framework.

Tuning Protocol

Every controller is tuned by the same PSO procedure (Kennedy & Eberhart, 1995) with a swarm of 30 particles and 40 iterations (1,230 evaluations), an inertia weight decreasing linearly from 0.9 to 0.4, acceleration coefficients of 1.5, a velocity clamp of 20 % of each range, and the same random seed. The objective is the mean ITAE over two tuning scenarios that are disjoint from the test cases: a 0.8 % load step in area 1 with a later 0.5 % reduction in area 2 and a 60 % PV loss lasting 5 s, and a 1.0 % step in area 2 with a 50 % PV loss lasting 4 s. Any excursion beyond 0.5 Hz is penalised. Table 3 lists the search bounds and the gains found; parameters that ended on a bound are marked.

Test Cases, Pre-Roll and Metrics

Two disturbance cases are reported. TC-1 is a clean 1 % load step in area 1 at $t = 2$ s. TC-2 is a composite case: the same step, a 0.75 % step in area 2 at 30 s, a 70 % loss of available PV for 4 s from 10 s and a 40 % loss for 3 s from 45 s. Separating the two matters, because settling time measured on TC-2 would measure recovery from a cloud event rather than from the load step.

Every test run is preceded by 120 s of undisturbed operation, and only the following window is scored. The deterministic controllers sit exactly at equilibrium during this pre-roll, so their results are

unchanged by it. A learned policy, however, may produce a non-zero action at equilibrium; when a run starts from rest, that action creates a transient that the baseline integrators take tens of seconds to cancel, and if a disturbance arrives in that interval the transient can be scored as regulation performance. Section VI-F shows that this effect is large enough to reverse the ranking of controllers, which is why the pre-roll is applied to every test.

Regulation is assessed by peak deviation, settling time to a ± 5 mHz band and the integral indices of the error e formed from Δf_1 , Δf_2 and ΔP_{tie} ,

$$ISE = \int e^2 dt, IAE = \int |e| dt, ITAE = \int t|e| dt, ITSE = \int te^2 dt \quad (29)$$

over the 60 s window. ITAE is the headline index because its time weighting penalises persistent error. Actuator activity is assessed by the total variation of the area-1 governor command at the control instants and by its root-mean-square increment,

$$TV(u) = \sum_k |u(k) - u(k-1)|, RMS\Delta u = [\text{mean}(u(k) - u(k-1))^2]^{\frac{1}{2}} \quad (30)$$

because a controller that regulates well by moving the valve continuously is not deployable, and regulation indices alone do not reveal it. Table 4

RESULTS AND DISCUSSION

Clean Load Step (TC-1)

Table 5 reports the step response. The storage converter, rated at 20 MW, can by itself cover the 20 MW step, so first-swing deviations are small for every design: 9.2 mHz for PI, 7.4 mHz for PID, 7.4 mHz for fuzzy-PID, 10.2 mHz for proposed scheme. On ITAE the fuzzy-PID is best (0.866), and the proposed scheme (3.816) is not: with no PV event in TC-1 its feed-forward paths are idle, and its gains, tuned on scenarios that contain cloud events,

Table 3: PSO Search bounds and identified gains (* = On A Search Bound).

Design	Free parameters and bounds	Identified gains	Tuning ITAE
PI	$K_p, K_i \in [0.01, 3]; K_{bd} \in [0, 3]$	$K_p=1.281, K_i=1.179, k_{bd}=1.204$	7.659
PID	as PI; $K_D \in [0, 1.5]$	$K_p=0.263, K_i=3.000^*, K_d=0.687, k_{bd}=2.817$	5.644
Fuzzy-PID	$K_e \in [0.2, 10]; K_{de} \in [0.05, 5]; K_u, K_i \in [0.01, 3]; K_{bd} \in [0, 3]$	$K_e=0.774, K_{de}=3.592, K_u=0.133, K_i=2.968, k_{bd}=3.000^*$	5.593
Proposed	$\kappa_0 \in [0.1, 3]; \Delta\kappa \in [0, 2]; \alpha \in [0.5, 20]; K_i \in [0.01, 3]; K_D \in [0, 1.5]; K_{ff} \in [0, 2]; K_{b1} \in [0, 3]; K_{b2} \in [0, 1.5]$	$\kappa_0=1.964, d\kappa_0=0.928, \alpha=4.109, K_i=2.716, K_d=1.310, K_{ff}=1.248, k_{b1}=1.005, k_{b2}=0.944$	4.903

Table 4: TD3 agent and training configuration.

Item	Value
Actor / critics	12-64-64-3 (ReLU, tanh) / two 15-64-64-1 (ReLU, linear)
Learning rates; optimiser	1×10^{-4} actor, 1×10^{-3} critics; Adam, gradient-norm clip 1.0
Discount; soft update; replay; batch	0.995; 0.005; 1×10^5 ; 128
Exploration; target smoothing; delay	OU σ 0.35→0.05; σ 0.20 clipped at 0.5; actor every 2 critic steps
Action and trim authority	$a \in [-1, 1]^3$; governors ± 0.03 p.u.; storage $\pm P_{BESS, \max}$
Initialisation	output layer $\times 0.01$, so the initial policy $\approx$ PI baseline
Episodes / length / interval	450 / 60 s / 0.1 s, starting from equilibrium
Randomised per episode	load steps ($\pm 1.5\%$, $\pm 1.2\%$) and times in both areas, up to two cloud events, PV penetration 2–40 %, $T_{ps} \pm 25\%$
Checkpoint selection	every 10 episodes, min mean ITAE over 3 validation scenarios disjoint from TC-1 and TC-2
Variants and seeds	unregularised; CAPS ($\lambda_T = \lambda_S = 1$); spectral norm ($c = 1$); 3 seeds each
Training time per run	7.1–15.5 min (one CPU core, NumPy)

Table 5: Transient performance: TC-1 clean step and TC-2 composite disturbance

Controller	TC-1 peak (mHz)	TC-1 overshoot (mHz)	TC-1 settling ± 5 mHz (s)	TC-1 ITAE	TC-2 peak Δf_1 (mHz)	TC-2 peak Δf_2 (mHz)	TC-2 peak ΔP_{tie} (MW)
PI (PSO)	9.2	1.4	6.53	5.301	40.6	44.4	10.3
PID (PSO)	7.4	0.2	0.53	0.905	39.3	43.0	11.6
Fuzzy-PID (PSO)	7.4	0.2	0.53	0.866	39.6	43.4	12.3
Proposed (PSO)	10.2	1.0	6.74	3.816	35.1	39.1	9.6
TD3 (residual)	12.1	0.5	8.42	6.906	37.7	43.4	10.2
TD3 + CAPS	9.3	0.7	6.32	2.757	12.8	24.4	8.4
TD3 + spectral norm	9.1	1.1	6.42	3.414	30.5	32.9	10.9

are not optimal for a pure load step. Section VI-D returns to this. Every deterministic design returns to within 1 mHz of nominal by the end of the window.

Composite Disturbance (TC-2)

Figure 6 shows the responses to the composite case. The deepest swing follows the 70 % cloud event at 10 s, which removes 28 MW of PV output while the storage can supply 20 MW and the thermal units ramp at only 1 MW/s; the deterministic designs reach first-swing minima between 35.1 and 40.6 mHz. The proposed scheme has peaks of 35.1 mHz in area 1, 39.1 mHz in area 2 and 9.6 MW on the tie-line, against 40.6 mHz, 44.4 mHz and 10.3 MW for PI and 39.3 mHz, 43.0 mHz and 11.6 MW for the PID, the best tuned benchmark. The learned agents are discussed in Section VI-F.

Integral Indices, Actuator Activity and Storage Use

Table 6 and Figure 7 give the integral indices. On ITAE the proposed scheme improves on the PI baseline by 15.2 % and on the best tuned benchmark, the PID, by 5.5 %. These margins are modest. When the benchmarks are tuned by the same procedure and given the same storage asset, a well-tuned droop already captures most of what the battery can do; the remaining gain comes from using the battery and the governors in a coordinated way during PV events.

The proposed scheme moves the governor command 2.6 times as much as the PI baseline (total variation 0.315 against 0.120), a modest increase for the regulation gained.

Figure 8 shows the storage. Under co-dispatch the converter reaches 100.0 % of rating, as it does under the benchmarks' tuned droop, and the SOC moves only between 58.68 % and 60.00 % in the 60 s window. A transient case of this length says nothing about energy management, which is why Section VI-H adds a 30-minute study.

Where the Improvement Comes From

Removing a component and re-simulating with the remaining gains unchanged confounds the value of the component with the cost of a mistuned remainder. Each reduced design was therefore re-tuned by PSO with the same budget and scenarios (Table 7, Figure 9). On TC-2, removing the PV feed-forward costs +18.8 % in ITAE and replacing storage co-dispatch by a tuned frequency droop costs +13.9 %, while removing the fuzzy gain schedule changes ITAE by +1.6 %. On the tuning scenarios the re-tuned design without scheduling is slightly better than the full one (4.806 against 4.903); since the full parameter space contains the unscheduled design as the special case $\Delta\kappa = 0$, the difference reflects the harder search in a larger space, not a disadvantage

of scheduling as such. Removing scheduling without re-tuning changes ITAE by -0.3 %. Figure 5 shows why: over TC-2 the magnitude of the ACE never exceeds 0.020 p.u., across which the identified schedule moves the gain only from 1.96 to 2.04 (3.8 %).

The improvement on PV-driven disturbances therefore comes from storage co-dispatch and PV feed-forward. The fuzzy gain schedule contributes nothing measurable to performance; its only role is to bound the proportional gain, which a fixed gain also does. Without storage at all every design fails: the re-tuned scheme reaches 227.4 mHz and an ITAE of 455.5, because the rate-limited thermal units cannot follow a cloud event of this size.

The ablation also qualifies the main result. On the clean step TC-1, every reduced design that retains storage does better than the full scheme (Table 7). The full design's gains are specialised for PV-driven disturbances, and the benefit reported here is specific to them.

Measurement Noise, Communication Delay, Sampling Jitter and Nonlinearities

Table 8 and Figure 10 report every design under non-ideal conditions on TC-2, with ITAE taken over the same 60 s window as before. Frequency measurements carry 3 mHz rms noise and tie-line measurements 1 MW rms, of the order of the frequency-error limits for synchrophasor measurement (Institute of Electrical and Electronics Engineers [IEEE], 2011); jitter perturbs each sampling instant by a Gaussian amount with a standard deviation of 20 % or 50 % of the 0.1 s interval; and the AGC link delays governor setpoints by up to 2 s. Storage commands formed from local measurements are not delayed (Section IV-C). The learned agents are central controllers, so their governor and storage trims are both delayed. A combined case applies the noise, a jitter of 30 % and a 0.5 s delay together. Stochastic cases are averaged over five realisations. Each case is also run for 300 s without measurement noise, and a sustained oscillation is recorded (marked † in Table 8) when the rms area-1 deviation over the last 75 s exceeds 2 mHz and has not fallen below 90 % of its value over the preceding 75 s.

Noise raises the ITAE of the deterministic designs by 17.8–29.7 %, a moderate penalty; the derivative terms act on a filtered ACE derivative, which keeps them from amplifying it. Jitter of 50 % changes ITAE by at most 12.0 %, the PID being the most sensitive. Communication delay is the dominant practical risk. At 1.0 s the ITAE of the designs tuned on an ideal channel rises by factors of 1.48–2.11, and the degradation is not monotonic in the delay: a pure delay adds phase lag in proportion to frequency, so its effect depends on which modes the loop excites, and the lightly damped inter-area mode of Section III-F is sensitive to particular delays. The ideal-channel proposed design develops a sustained




Figure 6: Responses to the composite disturbance TC-2: 1 % load step in area 1 at 2 s, 0.75 % in area 2 at 30 s, and PV losses of 70 % (10–14 s) and 40 % (45–48 s), shaded. All controllers sampled at 0.1 s with dead-band and GRC enforced, after a 120 s undisturbed pre-roll.




Figure 7: Integral performance indices under TC-2; ITAE printed above its bar.




Figure 8: Storage power and state of charge under TC-2 for the proposed scheme and for the PI benchmark's droop.




Figure 9: Re-tuned ablation under TC-2 (the no-storage design is omitted for scale; see Table 7).

Table 6: Integral indices and actuator activity under TC-2.

Controller	ISE ($\times 10^{-3}$)	IAE	ITAE	ITSE ($\times 10^{-3}$)	I T A E reduction vs PI (%)	Governor total variation	RMS command increment	Peak BESS power (%)
PI (PSO)	11.63	0.760	17.340	205.92	baseline	0.120	0.00039	100.0
PID (PSO)	10.72	0.675	15.566	190.73	+10.2	0.183	0.00057	100.0
Fuzzy-PID (PSO)	11.09	0.688	15.998	199.99	+7.7	0.187	0.00058	100.0
Proposed (PSO)	9.10	0.684	14.710	163.40	+15.2	0.315	0.00120	100.0
TD3 (residual)	11.47	0.807	19.356	215.90	-11.6	0.513	0.00238	100.0
TD3 + CAPS	3.23	0.509	12.836	74.44	+26.0	0.124	0.00037	100.0
TD3 + spectral norm	6.83	0.629	14.823	130.25	+14.5	0.122	0.00040	100.0

Table 7: Re-tuned ablation of the proposed scheme.

Configuration (each re-tuned by PSO)	Tuning ITAE	TC-2 ITAE	Change vs full (%)	TC-1 ITAE	Change vs full (%)	TC-2 peak (mHz)
Full proposed scheme	4.903	14.710	reference	3.816	reference	35.1
Without PV feed-forward	5.684	17.480	+18.8	0.898	-76.5	39.4
Without fuzzy gain scheduling	4.806	14.939	+1.6	1.990	-47.9	40.1
Storage on frequency droop (no co-dispatch)	5.471	16.753	+13.9	0.863	-77.4	39.9
Without feed-forward and scheduling	5.860	17.619	+19.8	0.954	-75.0	40.1
Without storage	321.114	455.519	+2996.7	234.364	+6041.1	227.4

Table 8: Itae over the 60 s window of tc-2 under non-ideal conditions; robust designs re-tuned with a 1.0 s AGC delay; Combined = Noise, 30 % JITTER and 0.5 S delay; † Sustained Oscillation.

Condition	PI	PID	Fuzzy- PID	Proposed	TD3	TD3 + CAPS	TD3 + SN	PI robust	PID robust	Fuzzy robust	Proposed robust
Ideal	17.34	15.57	16.00	14.71	19.36	12.84	14.82	19.89	20.74	21.94	15.64
Noise	20.82	20.19	20.37	17.33	20.11	17.09	19.06	21.61	24.40	21.13	17.79
Jitter 20%	17.21	16.97	16.66	14.85	19.78	12.78	14.78	19.78	20.70	22.21	15.58
Jitter 50%	17.35	17.43	17.18	14.93	20.58	12.88	14.77	19.65	20.74	22.63	15.61
Delay 0.2 s	17.61	22.56	22.33	16.14	20.62	15.35	17.12	21.82	22.66	20.73	17.49
Delay 0.5 s	21.43	27.65	26.16	20.43	20.20	16.33	17.36	28.60	23.94	24.31	16.87
Delay 1.0 s	31.12	32.79	23.65	27.28	20.17	15.67	29.12	29.93	23.84	31.03	17.10
Delay 2.0 s	25.67	18.67	19.77	23.44†	43.94	40.93	21.35	20.41	31.99	21.76	17.93
Combined	23.57	22.60	22.75	19.58	18.98	17.95	20.29	25.88	25.06	27.01	19.30

Table 9: Governor activity of the learned controller and its mitigation under tc-2 (representative runs, then all runs).

Configuration	Peak Δf_1 (mHz)	ITAE	Validation ITAE	Governor total variation	RMS command increment	Change in total variation (%)
TD3, raw policy output	37.7	19.36	13.19	0.513	0.00238	reference
TD3 + low-pass filter, 0.5 s	29.6	20.89	13.35	0.476	0.00149	-7.3
TD3 + low-pass filter, 1.0 s	27.8	19.50	14.77	0.352	0.00101	-31.4
TD3 + low-pass filter, 2.0 s	22.8	17.56	14.09	0.213	0.00064	-58.5
TD3 + rate limiter, 0.05 per step	30.3	20.84	12.84	0.357	0.00095	-30.4
TD3 + rate limiter, 0.02 per step	28.8	18.35	11.03	0.252	0.00062	-51.0
TD3 trained with CAPS smoothness	12.8	12.84	9.40	0.124	0.00037	-75.9
TD3 with spectral-norm Lipschitz bound	30.5	14.82	6.67	0.122	0.00040	-76.2
PI baseline (zero trim)	40.6	17.34	15.18	0.120	0.00039	-76.7
Proposed scheme	35.1	14.71	13.87	0.315	0.00120	-38.7
TD3 unregularised, all 3 runs (mean $\pm$ s.d.)	35.0 $\pm$ 3.2	18.18 $\pm$ 1.06	15.87 $\pm$ 4.39	0.317 $\pm$ 0.194	-	-
TD3 + CAPS, all 3 runs (mean $\pm$ s.d.)	39.0 $\pm$ 23.1	22.99 $\pm$ 9.75	14.38 $\pm$ 5.28	0.737 $\pm$ 0.722	-	-
TD3 + spectral norm, all 3 runs (mean $\pm$ s.d.)	25.2 $\pm$ 4.7	14.69 $\pm$ 0.21	6.91 $\pm$ 0.35	0.117 $\pm$ 0.025	-	-

oscillation from a delay of 1.25 s, persisting at about 6.5 mHz rms after 300 s. PSO with an ideal channel drives the integral gains high (Table 8), which leaves little phase margin once delay is inserted (Jiang et al., 2012; Zhang et al., 2013).

Re-tuning every design with a 1.0 s delay in the loop, by the same PSO procedure over a 90 s horizon, reduces this sensitivity, and the delay-robust proposed scheme is the most tolerant by a clear margin: its ITAE stays between 15.64 and 18.07 over 0–2 s with no sustained oscillation at any delay, against worst cases of 32.51 (PI), 31.99 (PID), 36.75 (Fuzzy) for the other delay-robust designs, and it costs 6.3 % in ideal-channel ITAE. Its fast actions do not travel over the delayed link. Where the fast storage response is implemented matters as much as how the controller is tuned: when the PI design was tuned with its storage droop also routed over a 1 s delayed link, the best design PSO found had a tuning ITAE of 702.5, against 17.08 with local storage, with K_P on its lower bound and the droop gain reduced to 0.013; secondary control and storage support were effectively switched off. The practical rules are to tune the secondary loop for the latency of the channel it will use, and to implement fast storage response in the plant controller. Central AGC commonly issues setpoints every few seconds; the 0.1 s interval used here is representative of a plant-level controller, and the delay study approximates the effect of the slower central channel.

Under the combined case the proposed scheme gives an ITAE of 19.58, against 22.60–23.57 for the benchmarks and 19.30 for its delay-robust version. The learned agents give 18.98 (unregularised TD3), 17.95 (TD3 + CAPS), 20.29 (TD3 + spectral norm) in that case, but they were trained on an ideal channel and degrade at longer delays: at 2.0 s their ITAE is 43.94 (unregularised TD3), 40.93 (TD3 + CAPS), 21.35 (TD3 + spectral norm), against 17.93 for the delay-robust proposed scheme.

The generation rate constraint and the dead-band are present in every result in this paper. Under the rate constraint the thermal units need at least 20 s to absorb a 1 % load step whatever the controller does, so the storage, not the control law, sets the first swing. The anti-windup of (9) is essential: removed from the tuned designs, it leaves every one of them in a sustained oscillation over 300 s, with rms area-1 deviations over the last 75 s of 355.2 mHz (PI), 392.9 mHz (PID), 396.6 mHz (fuzzy-PID), 7.9 mHz (proposed), against zero with it.

The Learned Controller: Evaluation, Actuator Chatter and Its Mitigation

Three TD3 variants were trained, each with 3 seeds: unregularised, with CAPS, and with a spectral-norm bound. For each run the trained checkpoint with the lowest validation ITAE is used. That checkpoint beat the PI baseline from which training starts (validation ITAE 15.18) in 2 of 3 unregularised runs, 2 of 3 CAPS runs and 3 of 3 spectral-norm runs; a deployment rule that keeps the initial policy as a candidate would have retained the baseline in the others. In each variant the run with the lowest validation ITAE is used as its representative (Figure 11, Table 9), and all runs are reported. Validation ITAE fluctuates widely between checkpoints, and training return is a poor guide to engineering performance.

Evaluating these policies exposed a pitfall that affects any comparison of learned and conventional LFC. A policy that simply holds a constant trim of 0.65 on the area-2 governor, with no feedback at all, scores a TC-2 ITAE of 13.76 and a peak of 18.7 mHz when the run starts from equilibrium, better than every controller in this study; with the 120 s pre-roll of Section V-C the

same policy scores 17.34 and 40.6 mHz, exactly the PI baseline. The trim pre-loads area 2 while the integrators are still cancelling it, and in a test that begins with a deficit this transient is scored as regulation. The trained agents learned such biases: all 3 spectral-norm policies hold near-constant positive governor trims (mean area-2 trim 0.40–0.64, standard deviation at most 0.10), and the representative one scores a TC-2 ITAE of 13.18 without the pre-roll and 14.82 with it. Because the training episodes and the validation scenarios all start from equilibrium, the agents were rewarded for this; all results below use the pre-roll.

Table 9 reports regulation and governor activity under TC-2. The representative unregularised policy has a governor total variation of 0.513, 4.3 times that of PI and 1.6 times that of the proposed scheme, and an ITAE of 19.36, worse than PI; across its 3 runs the ITAE is 18.18 ± 1.06 and the total variation 0.317 ± 0.194 . Its reward penalises action increments, and the policy chatters regardless. Run-time filtering reduces activity in proportion to the bandwidth removed: low-pass filters of 0.5, 1 and 2 s cut total variation by 7.3, 31.4 and 58.5 %, and rate limits of 0.05 and 0.02 per step by 30.4 and 51.0 %. The 2 s filter also improves this policy's ITAE, to 17.56, because the chatter was itself degrading regulation, but every filter adds a lag the policy was not trained with.

The two training-time regularisers behave very differently. The spectral-norm bound, a hard limit on the policy's Lipschitz constant, gave consistent policies: across its 3 runs the TC-2 ITAE is 14.69 ± 0.21 against 14.71 for the proposed scheme, the total variation is 0.117 ± 0.025 , about that of PI and 2.7 times smaller than the proposed scheme's, and the peak area-1 deviation is 21.7–30.5 mHz against 35.1 mHz. CAPS, a soft penalty, was unreliable. Its representative run gives the lowest TC-2 ITAE of any controller tested (12.84, peak 12.8 mHz, total variation 0.124), but its other runs give ITAE of 23.87 and 32.27, and one of them has the highest governor activity in the study (1.533); across runs the ITAE is 22.99 ± 9.75 . The reviewer's point is confirmed: reward shaping alone did not yield smooth actions. A hard Lipschitz bound did, reliably and with regulation equal to the deterministic scheme on this test; a soft smoothness penalty did so only on some seeds; and output filtering reduced activity at a cost in lag. These are short-horizon results, and Section VI-H shows that no learned agent managed the battery's energy over 30 minutes. Figure 1

Solar Penetration, and Why the Disturbance Is Held Fixed

PV penetration is raised from 2 % to 40 % with T_{ps} scaled by (3), so the inertia time constant falls from 20.0 s to 12.2 s. Two disturbance models are compared (Table 10, Figure 13). In the first, the cloud event removes a fixed 28 MW at every penetration, isolating the effect of falling inertia; in the second it removes 70 % of the whole fleet, so the loss grows to 560.0 MW at 40 %.

With the disturbance fixed, peak deviation rises with penetration for every controller: by 35.7 % for the proposed scheme and by 50.7, 82.8 and 91.7 % for PI, PID and fuzzy-PID, and by 48.5 % for the representative spectral-norm agent. The proposed scheme degrades least among the deterministic designs, because its feed-forward paths act on the disturbance before inertia shapes the response. The representative CAPS agent holds its peak between 10.5 and 19.4 mHz across the range, although Section VI-F showed that the other CAPS runs regulate TC-2 poorly. With the disturbance scaled, every controller exceeds the 0.8 Hz first-stage under-frequency load-shedding margin from 20.0 % penetration, and at 40 % the computed deviation reaches 3.65 Hz

Table 10: Peak area-1 frequency deviation (mhz) against pv penetration: cloud loss fixed at 28 mw, and scaled with the fleet.

PV pen. (%)	T_ps (s)	PI fixed	PID fixed	Fuzzy-PID fixed	Proposed fixed	TD3 + SN fixed	Scaled loss (MW)	PI scaled	PID scaled	Fuzzy-PID scaled	Proposed scaled	TD3 + SN scaled
2.0	20.00	40.6	39.3	39.6	35.1	30.5	28.0	40.6	39.3	39.6	35.1	30.5
10.0	18.37	41.7	42.8	43.2	35.3	30.3	140.0	593.9	602.6	603.0	582.9	578.5
20.0	16.33	44.2	47.2	49.2	36.7	31.2	280.0	> 800	> 800	> 800	> 800	> 800
30.0	14.29	53.9	57.5	58.2	43.2	39.8	420.0	> 800	> 800	> 800	> 800	> 800
40.0	12.24	61.2	71.9	76.0	47.7	45.3	560.0	> 800	> 800	> 800	> 800	> 800




Figure 10: : ITAE against AGC communication delay for the deterministic designs tuned on an ideal channel (dotted) and re-tuned with a 1.0 s delay (solid). Storage response is local and undelayed in all cases.




Figure 11: Validation ITAE during training for the three TD3 variants (representative run bold, other runs faint), with the PI baseline from which each starts and the proposed scheme.

Table 11: documented large simultaneous pv output reductions in north America.

Event	Date	Region	PV reduction	Mechanism reported	Source
Blue Cut Fire	16 Aug. 2016	Southern California	≈1,200 MW	Inverter trips on mis-measured low frequency (≈700 MW) and momentary cessation (≈450 MW) after line faults	NERC (2017)
Canyon 2 Fire	9 Oct. 2017	Southern California	≈900 MW	Momentary cessation; instantaneous trips on sub-cycle transient over-voltage	NERC (2018)
San Fernando	7 Jul. 2020	Southern California	≈1,000 MW	Momentary cessation with slow recovery; ac over-current and low-voltage trips	NERC (2020)
Odessa	9 May 2021	Texas	1,112 MW	PLL loss of synchronism, ac over-voltage trips, momentary cessation	NERC (2021)
Odessa	4 Jun. 2022	Texas	1,711 MW (+844 MW synchronous)	Ac over-current, phase-jump anti-islanding, ac over-voltage trips	NERC (2022)
Southwest Utah	10 Apr. 2023	Western Interconnection	921 MW (9 plants)	Failure to ride through a normally cleared 345 kV fault	NERC (2023)




Figure 12: Area-1 governor command under TC-2: (a) representative unregularised TD3 policy, raw and low-pass filtered; (b) representative policies trained with CAPS and with a spectral-norm bound, and the proposed scheme.




Figure 13: : Peak area-1 deviation against PV penetration: (a) cloud loss held at 28 MW, isolating falling inertia; (b) loss scaled to 70 % of the fleet, with the 0.8 Hz first-stage load-shedding margin.




Figure 14: : Thirty-minute stochastic run: (a) available PV and area-1 load deviation; (b) storage SOC for the storage law (19)–(21), with each of its elements reversed or removed, and for the representative learned agents.

Table 12: Thirty-minute stochastic run: storage energy management.

Case	SOC min (%)	SOC max (%)	SOC end (%)	rms Δf_1 (mHz)	max $ \Delta f_1 $ (mHz)	IAE	Time at SOC limit (%)	Storage throughput (MWh)
PI, storage on droop	58.36	61.16	58.56	5.53	58.4	15.85	0.0	3.05
Proposed storage law (19)–(21)	57.42	60.87	59.93	3.18	25.2	12.79	0.0	2.95
Proposed, without washout	26.66	60.53	32.06	4.22	51.9	13.90	0.0	4.10
Proposed, without SOC term	56.13	60.93	58.26	3.11	24.1	12.70	0.0	2.95
Proposed, SOC term reversed	51.52	61.26	51.52	3.07	24.5	12.63	0.0	3.05
No washout and SOC term reversed	22.00	60.57	22.20	80.11	570.0	102.21	10.6	4.94
As above, 5× SOC gain	22.00	60.66	22.00	71.95	369.1	130.14	15.9	5.12
TD3, unregularised	22.88	70.41	22.88	9.21	71.0	22.24	0.0	5.96
TD3 + CAPS	60.00	71.19	70.14	5.23	67.4	14.16	0.0	3.30
TD3 + spectral norm	29.44	62.90	29.44	4.25	42.2	14.56	0.0	4.30

for the proposed scheme and 3.66 Hz for PI, where the linearised model no longer applies and load shedding would already have acted. The controllers are indistinguishable there because a 20 MW battery and a 1 MW/s ramp cannot replace hundreds of megawatts within seconds. That is a reserve and storage-sizing problem, not a secondary-control problem.

The fixed disturbance is the main case for two reasons. First, cloud-driven variability does not scale with fleet size: irradiance fluctuations at separated sites are weakly correlated on timescales of seconds to minutes, so the relative variability of aggregate output falls as a fleet disperses (Lave, Kleissl, & Arias-Castro, 2012). A 70 % simultaneous loss is realistic for one plant under a cloud edge, not for a regional fleet. Second, the large simultaneous PV losses actually recorded had a different cause. Table 11 lists six documented North American events, from investigations by the North American Electric Reliability Corporation, in which between about 900 MW and 1,711 MW of PV output was lost within seconds following transmission faults, because inverters ceased current injection or tripped on protection settings (North American Electric Reliability Corporation [NERC], 2017, 2018, 2020, 2021, 2022, 2023). These are ride-through failures, addressed by inverter performance requirements (IEEE, 2018) and by contingency reserves sized to the credible loss; they unfold faster than secondary control can act, and the scaled case shows the outcome if they are left to it.

State-of-Charge Management over 30 Minutes

To assess energy management, every design was run for 30 minutes under a stochastic profile (Figure 14(a)): random-walk load variation with occasional steps in both areas, a 25 % decline in available PV, twelve cloud episodes of 5–40 s and 20–70 % depth, and a sustained 35 % shading between minutes 15 and 25. Table 12 summarises the outcome.

With the storage law (19)–(21) the SOC stays between 57.42 % and 60.87 % and ends at 59.93 %, and the rms area-1 deviation is 3.18 mHz, against 5.53 mHz for the PI benchmark with storage on droop. Two elements of the law are essential. Without the washout, the battery is commanded to replace the sustained PV deficit as though it were transient: the SOC falls to 26.66 % during the shading period and ends the run at 32.06 %, and the rms deviation rises to 4.22 mHz. With the restoring term reversed in sign it reinforces drift instead of opposing it, and the SOC leaves its reference with the ten-minute time constant of (21), ending at 51.52 % and still falling. Removing the SOC term altogether leaves a milder drift, to 58.26 %, over this run.

With both faults together the battery reaches its 22.00 % floor and is held there for 10.6 % of the run; the scheme's gains were identified with the battery available, and without it the system is poorly damped, as the no-storage design of Section VI-D already showed: the rms deviation rises to 80.11 mHz with excursions to 570.0 mHz. A five times larger reversed SOC gain reaches the floor sooner, is held there for 15.9 % of the run, and gives an rms deviation of 71.95 mHz.

No learned agent manages the battery. The unregularised agent drives the SOC down to 22.88 % by the end of the run (rms deviation 9.21 mHz); the CAPS agent drives the SOC to 71.19 % and ends at 70.14 % (rms deviation 5.23 mHz); the spectral-norm agent drives the SOC down to 29.44 % by the end of the run (rms deviation 4.25 mHz). Trained on 60 s episodes with a small SOC penalty, the agents never experienced the consequence of energy drift. This is the clearest argument in the paper for an explicit, designed state-of-charge law, at least with short-episode training;

a learned controller would need episodes long enough for energy to matter, or an explicit SOC constraint.

The 60 s results of Sections VI-A to VI-G therefore concern short-term frequency regulation only. The 30-minute study shows that the storage law (19)–(21) keeps the battery well within its energy window under persistent disturbance of the kind modelled; it does not establish behaviour over hours, where state-of-charge management interacts with energy scheduling and tertiary control, which are outside the scope of this paper.

Storage Rating

Because storage sets the first swing, its rating matters more than any control choice. With gains held at their 20 MW values, a 10 MW converter gives TC-2 peaks of 81.7 mHz for the proposed scheme and 81.0, 83.3 and 84.3 mHz for PI, PID and fuzzy-PID; a 40 MW converter gives 10.8, 12.4, 6.4 and 6.3 mHz. At 10 MW the proposed scheme is no longer best on ITAE (50.63 against 39.85 for fuzzy-PID), because its gains were identified for the larger converter. A scheme of this kind should be tuned for the storage it will command.

Limitations

This is a simulation study of a linearised two-area equivalent with the nonlinearities of Section III-C; no hardware-in-the-loop or field validation was performed. A two-area system does not exhibit the multiple inter-area modes of a large interconnection. Governor, turbine and reheat parameters are standard values rather than those of a particular utility, and PV and storage are aggregated first-order models; converter control, fault ride-through and voltage are outside scope. Every deterministic controller was tuned with one PSO run and one seed; a different optimiser or budget could change individual gains. Each learned variant was trained with 3 seeds, which shows the run-to-run variance but does not characterise it fully, and the agents were trained on episodes that start from equilibrium, which rewarded the start-up bias described in Section VI-F; training with a pre-roll, or with disturbances of both signs, would remove that incentive. The delay-robust designs were tuned for a single 1.0 s delay, whereas a real AGC channel presents a distribution of delays and a dispatch interval of seconds.

CONCLUSION

A supplementary load frequency control scheme was developed for a two-area reheat-thermal grid hosting a 40 MW PV plant and a co-located 20 MW / 8.3 MWh battery. It adds to the installed secondary loop feed-forward from the measured PV deviation, local co-dispatch of the battery with a washout and a restoring SOC term, and a bounded fuzzy gain schedule. With governor dead-band, generation rate constraint and anti-windup enforced, and every controller tuned by one PSO procedure on scenarios disjoint from the tests, the scheme reduced ITAE under a composite load-and-cloud disturbance by 15.2 % relative to PI and by 5.5 % relative to the best tuned benchmark, the PID. On a clean load step it was not the best design. These margins are modest.

The improvement comes mainly from coordinated battery dispatch and PV feed-forward, not from the fuzzy gain schedule. In a re-tuned ablation, replacing co-dispatch by a tuned droop raised ITAE by 13.9 % and removing PV feed-forward by 18.8 %, while removing the fuzzy schedule changed it by +1.6 %; the schedule is retained only because it bounds the proportional gain. Over 30 minutes the battery law kept the SOC within 57.42–60.87 % with an rms frequency deviation of 3.18 mHz, lower than the PI benchmark and every learned agent, and both its washout and the sign of its restoring term proved essential; the other results

concern short-term regulation only.

For deployment, measurement noise raised ITAE moderately and sampling jitter little, but communication delay was the dominant risk: at 1 s the ITAE of designs tuned for an ideal channel rose by up to 110.6 %, and the ideal-channel proposed design oscillated persistently from 1.25 s. Re-tuned for the channel, with fast battery response kept local, the proposed scheme was the most delay-tolerant design. Battery rating, more than the control law, sets the first swing, and the large simultaneous PV losses recorded in practice are inverter ride-through events that call for reserves and inverter performance requirements rather than better secondary control.

A residual TD3 agent sharing the same plant code, interface and interval chattered without regularisation despite a reward penalty on action changes. A hard spectral-norm Lipschitz bound gave consistent policies across seeds, with regulation equal to the deterministic scheme on the short test and governor activity near that of PI; a soft CAPS penalty produced the best single controller but was unreliable across seeds; output filtering reduced activity at a cost in lag. Testing learned policies from rest overstated their performance, and a pre-roll was needed for a fair comparison. Over 30 minutes no learned agent managed the battery's energy. Future work should validate the scheme in hardware-in-the-loop, extend it to interconnections with several inter-area modes, design for distributions of communication delay and realistic AGC intervals, and train learned controllers with pre-rolled episodes long enough, or constraints explicit enough, for energy management to be learned.

ACKNOWLEDGEMENT

The author thanks the Editor and the anonymous reviewer for careful and constructive comments, which substantially improved the rigour and clarity of the paper.

In accordance with the IEEE policy on artificial intelligence (AI)-generated content, the author discloses that an AI system, Claude (Anthropic), was used in the preparation of this article, as follows:

Code (Sections III–VI). The simulation, controller-tuning, reinforcement-learning and analysis code that produced every numerical result was written with Claude.

REFERENCES

- Bevrani, H. (2014). *Robust power system frequency control* (2nd ed.). Cham, Switzerland: Springer. <https://doi.org/10.1007/978-3-319-07278-4>
- Bevrani, H., Ise, T., & Miura, Y. (2014). Virtual synchronous generators: A survey and new perspectives. *International Journal of Electrical Power & Energy Systems*, 54, 244–254. <https://doi.org/10.1016/j.ijepes.2013.07.009>
- Chen, X., Qu, G., Tang, Y., Low, S., & Li, N. (2022). Reinforcement learning for selective key applications in power systems: Recent advances and future challenges. *IEEE Transactions on Smart Grid*, 13(4), 2935–2958. <https://doi.org/10.1109/TSG.2022.3154718>
- Cohn, N. (1957). Some aspects of tie-line bias control on interconnected power systems. *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems*, 75(3), 1415–1436.
- Datta, U., Kalam, A., & Shi, J. (2019). The relevance of large-scale battery energy storage (BES) application in providing primary frequency control with increased wind energy penetration. *Journal of Energy Storage*, 23, 9–18. <https://doi.org/10.1016/j.est.2019.02.013>
- Elgerd, O. I., & Fosha, C. E. (1970). Optimum megawatt-frequency control of multiarea electric energy systems. *IEEE Transactions on Power Apparatus and Systems*, PAS-89(4), 556–563. <https://doi.org/10.1109/TPAS.1970.292602>
- Ersdal, A. M., Imsland, L., & Uhlen, K. (2016). Model predictive load-frequency control. *IEEE Transactions on Power Systems*, 31(1), 777–785. <https://doi.org/10.1109/TPWRS.2015.2412614>
- Esram, T., & Chapman, P. L. (2007). Comparison of photovoltaic array maximum power point tracking techniques. *IEEE Transactions on Energy Conversion*, 22(2), 439–449. <https://doi.org/10.1109/TEC.2006.874230>
- Fosha, C. E., & Elgerd, O. I. (1970). The megawatt-frequency control problem: A new approach via optimal control theory. *IEEE Transactions on Power Apparatus and Systems*, PAS-89(4), 563–577. <https://doi.org/10.1109/TPAS.1970.292603>

Figures and tables (Figures 1–14, Tables 1–12). These were prepared with Claude; the numerical figures and tables were generated directly from the simulation results.

Text (all sections). Claude drafted and edited the text, including the response to the reviewer, and assisted in formatting and checking the references.

The author defined the topic and scope of the study, directed the revision in response to the review, and reviewed, verified and approved all content. The author takes full responsibility for the article.

FUNDING: This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflicts of interest : The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

DATA AVAILABILITY : This is a simulation study and no experimental data were generated. All information needed to reproduce the results is reported in the manuscript: the plant, photovoltaic and battery parameters (Table 2), the particle swarm search bounds and identified gains (Table 3), the reinforcement-learning agent and training configuration (Table 4), and the disturbance and test scenarios (Section V). The documented photovoltaic output reductions in Table 11 are taken from publicly available reports of the North American Electric Reliability Corporation, cited in the reference list.

The simulation code, the tuned controller gains, the trained agent weights and the numerical results underlying every figure and table are available from the corresponding author on reasonable request.

AUTHORS' CONTRIBUTIONS: Mohammad Ali is the sole author. Following the CRediT taxonomy, his contributions are: conceptualization, methodology, software, validation, formal analysis, investigation, data curation, visualization, writing – original draft, writing – review and editing, and project administration. The use of an AI system in software, analysis and drafting is disclosed in the Acknowledgement; the AI system is not an author. The author has read and approved the final manuscript.

- Fujimoto, S., van Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 1587–1596). Stockholm, Sweden: PMLR. Retrieved from <https://arxiv.org/abs/1802.09477>
- Ibraheem, Kumar, P., & Kothari, D. P. (2005). Recent philosophies of automatic generation control strategies in power systems. *IEEE Transactions on Power Systems*, 20(1), 346–357. <https://doi.org/10.1109/TPWRS.2004.840438>
- Institute of Electrical and Electronics Engineers. (2011). *IEEE standard for synchrophasor measurements for power systems* (IEEE Std C37.118.1-2011). New York, NY: Author. <https://doi.org/10.1109/IEEESTD.2011.6111219>
- Institute of Electrical and Electronics Engineers. (2018). *IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces* (IEEE Std 1547-2018). New York, NY: Author. <https://doi.org/10.1109/IEEESTD.2018.8332112>
- Jang, J.-S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(3), 665–685. <https://doi.org/10.1109/21.256541>
- Jiang, L., Yao, W., Wu, Q. H., Wen, J. Y., & Cheng, S. J. (2012). Delay-dependent stability for load frequency control with constant and time-varying delays. *IEEE Transactions on Power Systems*, 27(2), 932–941. <https://doi.org/10.1109/TPWRS.2011.2172821>
- Johannink, T., Bahl, S., Nair, A., Luo, J., Kumar, A., Loskyll, M., & Levine, S. (2019). Residual reinforcement learning for robot control. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 6023–6029). Montreal, Canada: IEEE. <https://doi.org/10.1109/ICRA.2019.8794127>
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of ICNN'95 – International Conference on Neural Networks* (Vol. 4, pp. 1942–1948). Perth, Australia: IEEE. <https://doi.org/10.1109/ICNN.1995.488968>
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. Paper presented at the 3rd International Conference on Learning Representations, San Diego, CA. Retrieved from <https://arxiv.org/abs/1412.6980>
- Kothare, M. V., Campo, P. J., Morari, M., & Nett, C. N. (1994). A unified framework for the study of anti-windup designs. *Automatica*, 30(12), 1869–1883. [https://doi.org/10.1016/0005-1098\(94\)90048-5](https://doi.org/10.1016/0005-1098(94)90048-5)
- Kundur, P. (1994). *Power system stability and control*. New York, NY: McGraw-Hill.
- Lave, M., Kleissl, J., & Arias-Castro, E. (2012). High-frequency irradiance fluctuations and geographic smoothing. *Solar Energy*, 86(8), 2190–2199. <https://doi.org/10.1016/j.solener.2011.06.031>
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. Paper presented at the 4th International Conference on Learning Representations, San Juan, Puerto Rico. Retrieved from <https://arxiv.org/abs/1509.02971>
- MacVicar-Whelan, P. J. (1976). Fuzzy sets for man-machine interaction. *International Journal of Man-Machine Studies*, 8(6), 687–697. [https://doi.org/10.1016/S0020-7373\(76\)80030-2](https://doi.org/10.1016/S0020-7373(76)80030-2)
- Mamdani, E. H., & Assilian, S. (1975). An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7(1), 1–13. [https://doi.org/10.1016/S0020-7373\(75\)80002-2](https://doi.org/10.1016/S0020-7373(75)80002-2)
- Mi, Y., Fu, Y., Wang, C., & Wang, P. (2013). Decentralized sliding mode load frequency control for multi-area power systems. *IEEE Transactions on Power Systems*, 28(4), 4301–4309. <https://doi.org/10.1109/TPWRS.2013.2277131>
- Milano, F., Dörfler, F., Hug, G., Hill, D. J., & Verbič, G. (2018). Foundations and challenges of low-inertia systems. In *2018 Power Systems Computation Conference (PSCC)* (pp. 1–25). Dublin, Ireland: IEEE. <https://doi.org/10.23919/PSCC.2018.8450880>
- Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. Paper presented at the 6th International Conference on Learning Representations, Vancouver, Canada. Retrieved from <https://arxiv.org/abs/1802.05957>
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- Mysore, S., Mabsout, B., Mancuso, R., & Saenko, K. (2021). Regularizing action policies for smooth control with reinforcement learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1810–1816). Xi'an, China: IEEE. <https://doi.org/10.1109/ICRA48506.2021.9561138>
- North American Electric Reliability Corporation. (2017). *1,200 MW fault induced solar photovoltaic resource interruption disturbance report*. Retrieved from https://www.nerc.com/pa/rrm/ea/Documents/1200_MW_Fault_Induced_Solar_Photosvtaic_Resource_Interruption_Final.pdf
- North American Electric Reliability Corporation. (2018). *900 MW fault induced solar photovoltaic resource interruption disturbance report*. Retrieved from <https://www.nerc.com/pa/rrm/ea/Documents/900%20MW%20Solar%20Photovoltaic%20Resource%20Interruption%20Disturbance%20Report.pdf>
- North American Electric Reliability Corporation. (2020). *San Fernando disturbance: Southern California event: July 7, 2020* (Joint NERC and WECC staff report). Retrieved from https://www.nerc.com/pa/rrm/ea/Pages/July_2020_San_Fernando_Disturbance_Report.aspx
- North American Electric Reliability Corporation. (2021). *Odessa disturbance: Texas events: May 9, 2021 and June 26, 2021* (Joint NERC and Texas RE staff report). Retrieved from https://www.nerc.com/pa/rrm/ea/Documents/Odessa_Disturbance_Report.pdf
- North American Electric Reliability Corporation. (2022). *2022 Odessa disturbance: Texas event: June 4, 2022* (Joint NERC and Texas RE staff report). Retrieved from https://www.nerc.com/globalassets/our-work/reports/white-papers/nerc_2022_odessa_disturbance_report-1.pdf

- North American Electric Reliability Corporation. (2023). *2023 Southwest Utah solar loss disturbance* (Joint NERC and WECC staff report). Retrieved from https://www.nerc.com/pa/rrm/ea/Documents/NERC_2023_Southwest_UT_Disturbance_Report.pdf
- Pandey, S. K., Mohanty, S. R., & Kishor, N. (2013). A literature survey on load-frequency control for conventional and distribution generation power systems. *Renewable and Sustainable Energy Reviews*, 25, 318–334. <https://doi.org/10.1016/j.rser.2013.04.029>
- Rocabert, J., Luna, A., Blaabjerg, F., & Rodríguez, P. (2012). Control of power converters in AC microgrids. *IEEE Transactions on Power Electronics*, 27(11), 4734–4749. <https://doi.org/10.1109/TPEL.2012.2199334>
- Sahu, R. K., Panda, S., & Chandra Sekhar, G. T. (2015). A novel hybrid PSO-PS optimized fuzzy PI controller for AGC in multi area interconnected power systems. *International Journal of Electrical Power & Energy Systems*, 64, 880–893. <https://doi.org/10.1016/j.ijepes.2014.08.021>
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., & Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning* (pp. 387–395). Beijing, China: PMLR.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). Cambridge, MA: MIT Press.
- Takagi, T., & Sugeno, M. (1985). Fuzzy identification of systems and its applications to modeling and control. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-15(1), 116–132. <https://doi.org/10.1109/TSMC.1985.6313399>
- Tan, W. (2010). Unified tuning of PID load frequency controller for power systems via IMC. *IEEE Transactions on Power Systems*, 25(1), 341–350. <https://doi.org/10.1109/TPWRS.2009.2036463>
- Tielens, P., & Van Hertem, D. (2016). The relevance of inertia in power systems. *Renewable and Sustainable Energy Reviews*, 55, 999–1009. <https://doi.org/10.1016/j.rser.2015.11.016>
- Ulbig, A., Borsche, T. S., & Andersson, G. (2014). Impact of low rotational inertia on power system stability and operation. *IFAC Proceedings Volumes*, 47(3), 7290–7297. <https://doi.org/10.3182/20140824-6-ZA-1003.02615>
- Villalva, M. G., Gazoli, J. R., & Filho, E. R. (2009). Comprehensive approach to modeling and simulation of photovoltaic arrays. *IEEE Transactions on Power Electronics*, 24(5), 1198–1208. <https://doi.org/10.1109/TPEL.2009.2013862>
- Wood, A. J., Wollenberg, B. F., & Sheblé, G. B. (2013). *Power generation, operation, and control* (3rd ed.). Hoboken, NJ: Wiley.
- Yan, Z., & Xu, Y. (2019). Data-driven load frequency control for stochastic power systems: A deep reinforcement learning method with continuous action search. *IEEE Transactions on Power Systems*, 34(2), 1653–1656. <https://doi.org/10.1109/TPWRS.2018.2881359>
- Yan, Z., & Xu, Y. (2020). A multi-agent deep reinforcement learning method for cooperative load frequency control of a multi-area power system. *IEEE Transactions on Power Systems*, 35(6), 4599–4608. <https://doi.org/10.1109/TPWRS.2020.2999890>
- Yousef, H. A., Al-Kharusi, K., Albadi, M. H., & Hosseinzadeh, N. (2014). Load frequency control of a multi-area power system: An adaptive fuzzy logic approach. *IEEE Transactions on Power Systems*, 29(4), 1822–1830. <https://doi.org/10.1109/TPWRS.2013.2297432>
- Zhang, C.-K., Jiang, L., Wu, Q. H., He, Y., & Wu, M. (2013). Delay-dependent robust load frequency control for time delay power systems. *IEEE Transactions on Power Systems*, 28(3), 2192–2201. <https://doi.org/10.1109/TPWRS.2012.2228281>

©2026 Ali M, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License 4.0 International License.

Cite this article as: Ali M, Design and Development of an AI-Based Load Frequency Control System for a Power Transmission Grid Integrated with a Solar PV Microgrid. *Glob Open Access J Sci*. 2026;2(3):71–87.