Online Neural Gain Tuning for Load Frequency Control of Two-Area Power Systems under Generation Rate Constraints

R. A. Alao, I. A. Aleshinloye*, F. W. Shittu, I. O. Kolawole, A. O. Otuoze
Department of Electrical and Electronics Engineering, University of Ilorin, Ilorin, Nigeria

ABSTRACT: Load frequency control studies often credit performance gains to new controller structures, while the limit on how fast generating units can change their output is assumed rather than tested. This study examines whether online neural gain tuning can improve frequency regulation when that rate limit is active, and how far the limit restricts performance. A two-area, 15 MW system with a photovoltaic source in one area was modelled with an explicit tie-line synchronising integrator, valve position limits and a generation rate constraint of 30% per minute on the turbine mechanical power. A fixed-gain proportional-integral-derivative controller, tuned by particle swarm optimisation over six disturbance scenarios, served as the benchmark. A feed-forward network with two inputs and four hidden neurons adjusted the three gains online from the area control error and its filtered derivative. Its corrections were bounded so that it reduced exactly to the benchmark at zero output weights, and it was tuned with the same optimiser, objective and scenarios, starting from a swarm that contained the benchmark. Both designs were tested on six held-out disturbances and five other rate limits. The neural gain tuner reduced the integral of time-multiplied absolute error by 23.5% and shortened frequency settling from 18% to 22%. It kept an advantage on unseen disturbances within its trained range and fell back to benchmark behaviour beyond it. Removing the rate constraint improved the achievable objective about sixtyfold, showing that the constraint, not the controller, is the dominant limitation. Therefore, online neural gain tuning is a low-risk, bounded enhancement to frequency regulation although the generation rate constraint, not the controller, remains the dominant bottleneck on achievable performance.

KEYWORDS: Area control error; artificial neural network; generation rate constraint; load frequency control; particle swarm optimisation; tie-line power

[Received June. 21, 2026; Revised Sept 15, 2026; Accepted Sept 15, 2026] Print ISSN: 0189-9546 | Online ISSN: 2437-2110

I. INTRODUCTION

In an interconnected power system, frequency is a shared quantity rather than an independent one. When the load changes in one control area, every synchronously connected area is affected, and the tie-line flows between them shift until generation and demand balance again. Load frequency control is the secondary control layer that restores frequency to its nominal value and returns scheduled tie-line flows to their contracted levels after the primary droop response has arrested the initial excursion (Bevrani, 2014; Kundur, 1994). The classical formulation of this problem, together with the area control error signal that drives it, comes from Elgerd & Fosha (1970), and it remains the basis of most published studies on the topic.

Recent literature on this problem, by contrast, concentrates heavily on the controller (Gulzar et al., 2025; Oladipo et al., 2025; Shahzad et al., 2026). Proportional-integral-derivative (PID) structures tuned by metaheuristic search (Can et al., 2026; Khalil et al., 2023; Kumar et al., 2024; Paliwal et al., 2022; Qu et al., 2024; Rahnama, 2020; Younis et al., 2025), fuzzy and type-2 fuzzy supervisors (Razmi et al., 2025; Soliman et al., 2022), sliding mode designs (Biswas et al., 2025; Wang & Liu, 2021), model predictive schemes (Ayman et al., 2026) and reinforcement learning agents (El-Hameed et al., 2025; Li et al., 2025) have all been reported to improve transient frequency behaviour. A common feature of these studies is that the comparison baseline is a conventionally tuned controller rather than one optimised under the same conditions and against the same objective. When the baseline is weak, the reported improvement measures the weakness of the baseline as much as the strength of the proposed method (Halim et al., 2021).

A second feature of this literature is the way the generation rate constraint (GRC) is treated. Thermal generating units cannot change their mechanical output quickly, so this limit is usually represented as a rate limiter acting on the turbine output (Kalyan et al., 2022; Saini & Ohri, 2023). The constraint is frequently listed among the modelled nonlinearities without evidence that it becomes active during the simulated transient (Ramjug-Ballgobin & Ramlukon, 2021; Saini & Ohri, 2023). A rate limit that never binds contributes nothing to the dynamics, and any explanation of controller behaviour that appeals to it is then unsupported. The third feature concerns neural controllers trained to reproduce the output of an existing controller (Dash et al., 2025; El-Rifaie et al., 2025; Tolba et al., 2026; Yakout et al., 2024). A

*Corresponding author: imranaleshinloye97@gmail.com

doi: http://doi.org/10.63746/njtd.v23i3.4940

network trained on input-output pairs captured from a tuned controller is performing imitation learning (Ross et al., 2011). Such a network is bounded by the performance of the controller it imitates, and small approximation errors compound once the closed loop leaves the state distribution that generated the training data. Reports in which an imitated controller substantially outperforms its own demonstrator therefore require an explanation that is rarely given.

This article addresses these three points together. First, the generation rate constraint is placed on the turbine mechanical power, converted from a stated percentage per minute, and its activation is demonstrated rather than assumed. Second, the benchmark controller is tuned by particle swarm optimisation (PSO) over the same six disturbance scenarios, against the same objective and with the same computational budget as the proposed scheme, and the benchmark is included in the initial swarm of the proposed scheme so that the latter cannot be worse by construction. Third, the artificial neural network (ANN) adjusts the gains of the controller online rather than replacing it, which removes the imitation bound entirely. Performance is reported on the tuning scenarios and separately on six held-out scenarios, including one deliberately outside the trained disturbance range. The remainder of the article is organised as follows. Section II develops the system model, Section III describes the benchmark tuning and the online gain tuner, and Section IV reports the model validation, the activation of the rate constraint, the optimisation statistics, the transient performance, the generalisation and the sensitivity to the rate limit. Section V states the limitations of the work, and Section VI concludes the article.

II. SYSTEM MODEL

The plant is a two-area interconnected power system expressed in per unit on a common 15 MW base at a nominal frequency of 50 Hz. Each area is represented by an equivalent generating unit consisting of a speed governor, a non-reheat turbine and a lumped generator-load block, with primary droop feedback around the governor (Bevrani, 2014; Kundur, 1994; Patra & Patnaik, 2026). A photovoltaic source is connected in Area 1, and the two areas exchange power through a tie-line whose flow is determined by the accumulated frequency difference between them. The structure of the system is shown in Figure 1.

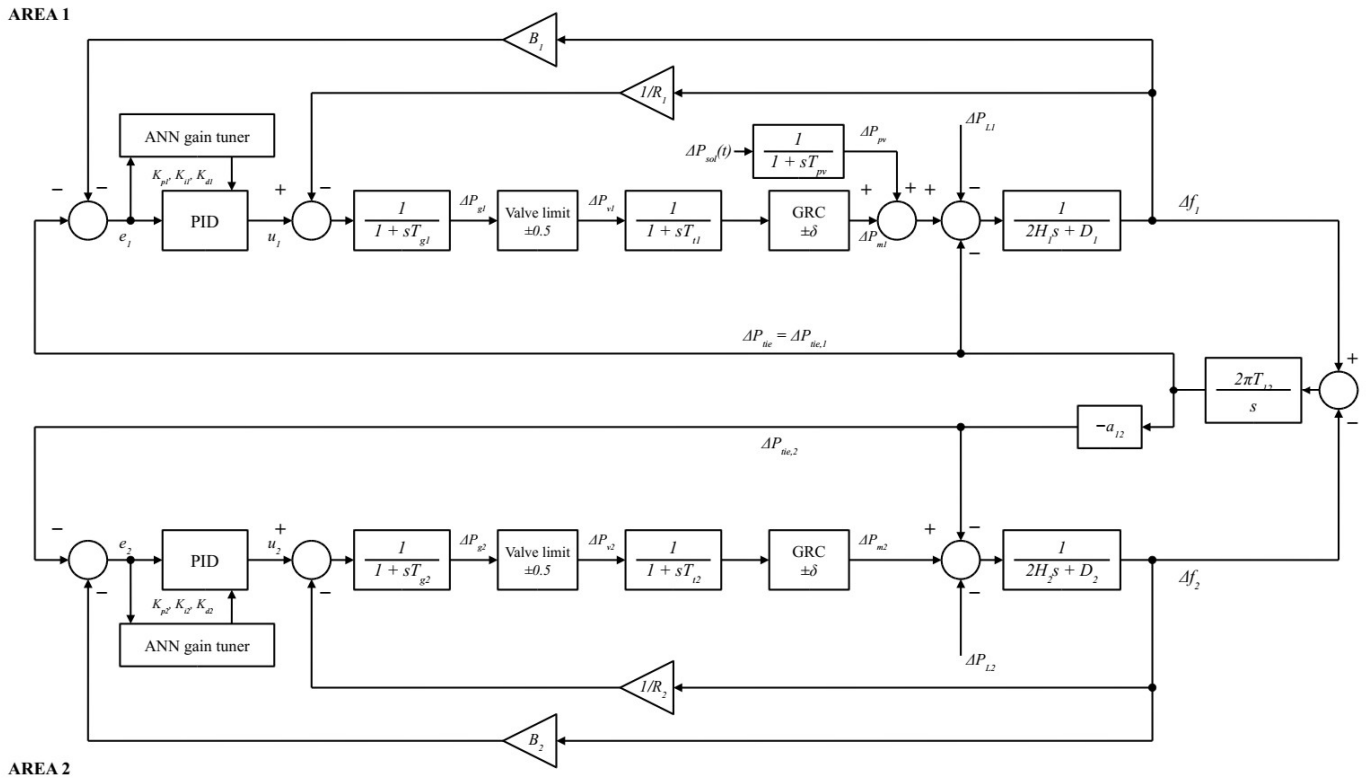


Figure 1 Structure of the two-area load frequency control system

A. Generator-Load and Prime Mover Dynamics

The frequency deviation of an area follows the imbalance between mechanical power, renewable in-feed, tie-line export and load, damped by the frequency sensitivity of the load (Elgerd, 1982; Kundur, 1994). The generator-load relationship for area i is given by Equation (1), in which the tie-line term carries a negative sign for Area 1 and the opposite sign for Area 2 through the area capacity ratio.

$$2H_i \frac{d\Delta f_i}{dt} = \Delta P_{mi} + \Delta P_{\mathrm{pv}} - \Delta P_{\mathrm{tie},i} - \Delta P_{Li} - D_i \Delta f_i \quad (1)$$

where H_i is the equivalent inertia constant of area i, Δf_i is its frequency deviation, ΔP_{mi} is the turbine mechanical power deviation, $\Delta P_{\mathrm{pv}i}$ is the photovoltaic power deviation (equal

to ΔP_{pv} in Area 1 and zero in Area 2), $\Delta P_{\text{tie},i}$ is the tie-line power deviation seen by area i, ΔP_{Li} is the load disturbance, D_i is the load damping coefficient and t is time.

The speed governor is represented as a first-order lag driven by the supplementary control signal and the droop feedback, as expressed in Equation (2). The speed droop R_i sets the steady-state share of an imbalance that the unit picks up without supplementary action (Bevrani, 2014).

$$T_{gi}\frac{d\Delta P_{gi}}{dt} = u_i - \frac{\Delta f_i}{R_i} - \Delta P_{gi} \tag{2}$$

where T_{gi} is the governor time constant, ΔP_{gi} is the governor output and u_i is the supplementary control signal.

The governor output passes through a valve position limit before reaching the turbine (Ramjug-Ballgobin & Ramlukon, 2021). The turbine is a first-order lag whose output is the mechanical power increment, and the generation rate constraint caps the rate at which this increment can change. The combined turbine and rate constraint relationship is stated in Equation (3) and discussed further in Section II.C.

$$\frac{d\Delta P_{mi}}{dt} = \text{sat}\left[\frac{\Delta P_{vi} - \Delta P_{mi}}{T_{ti}}, \pm\delta\right] \tag{3}$$

where ΔP_{vi} is the valve position after the limit of ± 0.5 pu MW, T_{ti} is the turbine time constant, δ is the generation rate limit, taken as the same in both areas, and $\text{sat}[\cdot, \pm\delta]$ clips its first argument to the interval from $-\delta$ to δ.

B. Tie-Line Power and Area Control Error

The tie-line power deviation is not an algebraic function of the area frequencies. It is the integral of their difference scaled by the synchronising coefficient of the interconnection, as shown in Equation (4) (Elgerd & Fosha, 1970; Kundur, 1994). This distinction matters because an algebraic or controller-generated substitute for this relationship changes the plant rather than the controller, and a comparison built on such a substitute is then not a comparison of controllers.

$$\frac{d\Delta P_{\text{tie}}}{dt} = 2\pi T_{12}(\Delta f_1 - \Delta f_2) \tag{4}$$

where ΔP_{tie} is the tie-line power deviation from Area 1 to Area 2 and T_{12} is the synchronising coefficient. The deviation seen by each area is $\Delta P_{\text{tie},1} = \Delta P_{\text{tie}}$ and $\Delta P_{\text{tie},2} = -a_{12}\Delta P_{\text{tie}}$, where a_{12} is the area capacity ratio, equal to unity when both areas are expressed on a common base.

Each area forms its control signal from the area control error, which combines the local tie-line deviation with the local frequency deviation weighted by the frequency bias factor. The bias factor is set to the area frequency response characteristic, as given in Equation (5) (Bevrani, 2014; Elgerd & Fosha, 1970).

$$B_i = D_i + \frac{1}{R_i} \tag{5}$$

where B_i is the frequency bias factor of area i. With this bias, the area control error is formed as in Equation (6).

$$\text{ACE}_i = \Delta P_{\text{tie},i} + B_i\Delta f_i \tag{6}$$

where ACE_i is the area control error of area i.

The supplementary controller acts on the negative of the area control error, so that a frequency deficit produces a positive command to the governor. This sign convention is stated explicitly in Equation (7), because writing the control law without it describes a positive feedback loop.

$$u_i = C_i(s) \cdot (-\text{ACE}_i) \tag{7}$$

where $C_i(s)$ is the transfer function of the supplementary controller in area i and s is the Laplace variable.

C. Generation Rate Constraint

The generation rate constraint expresses the physical inability of a generating unit to change its mechanical output faster than a given rate (Kalyan et al., 2022). It is usually stated as a percentage of rated capacity per minute and must be converted to per unit per second before use, a conversion that determines whether the constraint binds at all (Saini & Ohri, 2023). A value of 30% per minute is adopted here, which gives the limit in Equation (8).

$$\delta = \frac{0.30}{60} = 5.000 \times 10^{-3}\ \text{pu/s} \tag{8}$$

This value lies within the range of fast-ramping plant. Published figures place a heavy-duty F-class gas turbine near 10% per minute of rated capacity (Gonzalez-Salazar et al., 2018), fast-ramp control packages raise gas turbine capability to at least 16.6% per minute (Park et al., 2021), and reciprocating gas engines reach an effective 50% per minute (Wärtsilä, 2024). A 15 MW system of the kind modelled here would be served by plant in this class. The constraint is applied to the turbine mechanical output rather than to the valve position, so that it limits the quantity it is meant to limit (Kalyan et al., 2022). Its activity is examined in Section IV.B rather than assumed.

D. Photovoltaic Source

The photovoltaic plant in Area 1 is represented as a first-order lag between available solar power and injected electrical power, as given in Equation (9) (Bakeer et al., 2022). The input is an irradiance profile rather than a constant, so the source can be exposed to a cloud transient in which output falls to 40% of nominal over 1.5 s, holds for 6 s and recovers over a further 1.5 s (Agüera-Pérez et al., 2024).

$$T_{\text{pv}}\frac{d\Delta P_{\text{pv}}}{dt} = \Delta P_{\text{sol}}(t) - \Delta P_{\text{pv}} \tag{9}$$

where T_{pv} is the photovoltaic time constant, ΔP_{pv} is the injected photovoltaic power deviation and $\Delta P_{\mathrm{sol}}(t)$ is the available solar power set by the irradiance profile.

This is a linear representation of the plant dynamics with a time-varying input. It is not a detailed model of the array, the converter or the maximum power point tracker, and no claim of electrical realism at that level is made. Its purpose is to supply a disturbance whose shape differs from a step, so that controller behaviour can be assessed against something other than an idealised load change. The complete parameter set is listed in Table 1.

Table 1 System parameters

Parameter	Symbol	Area 1	Area 2	Unit
Equivalent inertia	H_i	0.083	0.101	pu MWs
Damping coefficient	D_i	0.015	0.016	pu MW/Hz
Governor time constant	T_{gi}	0.08	0.06	s
Turbine time constant	T_{ti}	0.40	0.44	s
Speed droop	R_i	3.00	2.73	Hz/pu MW
Frequency bias factor	B_i	0.34833	0.38233	pu MW/Hz
Valve position limits	ΔP_{vi}	±0.5	±0.5	pu MW
Generation rate constraint	δ	5.000×10^{-3}	5.000×10^{-3}	pu/s
Synchronising coefficient	T_{12}	0.08	–	pu MW/Hz
Area capacity ratio	a_{12}	1.0	–	–
PV time constant	T_{pv}	1.8	–	s
Base power / frequency	S, f_0	15 / 50	15 / 50	MW / Hz

III. CONTROLLER DESIGN

In this section, two controllers are described and compared. The first is a fixed-gain proportional-integral-derivative controller in each area, with the six gains chosen by particle swarm optimisation. The second keeps that controller structure but recomputes its three gains at every sample from measured signals, using a small feed-forward network whose weights are also chosen by particle swarm optimisation. Both controllers are tuned against the same objective, over the same six scenarios, with the same optimiser and the same number of independent runs. This construction removes the possibility that a measured difference reflects unequal tuning effort rather than a difference between the control laws.

A. Fixed-Gain Benchmark and Its Tuning

The benchmark controller in each area implements the parallel form with a filtered derivative term, as expressed in Equation (10) (Åström & Hägglund, 2006). Computing the gains outside the integrator, rather than inside it, allows the same expression to serve the adaptive controller of Section III.B with the constant gains replaced by time-varying ones.

$$u_i = K_{pi}e_i + K_{ii}q_i + K_{di}D_{f,i} \tag{10}$$

where $e_i = -\mathrm{ACE}_i$ is the control error, K_{pi}, K_{ii} and K_{di} are the proportional, integral and derivative gains of area i, q_i is the integral of the error and $D_{f,i}$ is its filtered derivative. The integral state evolves as in Equation (11).

$$\frac{dq_i}{dt} = e_i \tag{11}$$

The filtered derivative is produced by the first-order filter in Equation (12), with the filter coefficient N set to 10.

$$\frac{dD_{f,i}}{dt} = N\left(\frac{de_i}{dt} - D_{f,i}\right) \tag{12}$$

Particle swarm optimisation searches the six-dimensional gain vector (Kennedy & Eberhart, 1995). Each particle updates its velocity from its own best position and the swarm best, with a linearly decreasing inertia weight, as given in Equation (13) (Shi & Eberhart, 1998).

$$v_k(n+1) = wv_k(n) + c_1r_1\big(\mathrm{pbest}_k - x_k(n)\big) + c_2r_2\big(\mathrm{gbest} - x_k(n)\big) \tag{13}$$

where v_k and x_k are the velocity and position of particle k, n is the iteration counter, w is the inertia weight, c_1 and c_2 are the cognitive and social coefficients, r_1 and r_2 are random numbers drawn uniformly from [0, 1], pbest_k is the best position found by particle k and gbest is the best position found by the swarm. The position of each particle then follows from Equation (14).

$$x_k(n+1) = x_k(n) + v_k(n+1) \tag{14}$$

Positions are clipped to the search bounds and velocities to one fifth of the bound range (Clerc & Kennedy, 2002).

The optimiser settings are listed in Table 2. Bounds of zero to twenty are used for the benchmark gains, which is wide enough that no gain reached the upper limit in any of the ten runs. Ten independent runs are performed with different random seeds so that the best value, mean and standard deviation of the objective can be reported, which a single run cannot provide (Halim et al., 2021).

Table 2 Optimiser settings

Setting	Fixed-gain benchmark	Neural gain tuner
Decision variables	6 gains	54 network weights
Swarm size	40	45
Iterations	60	70
Independent runs	10	10
Search bounds	[0, 20]	[−3, 3]
Cognitive / social coefficients	2.0 / 2.0	2.0 / 2.0
Inertia weight	0.9 down to 0.4	0.9 down to 0.4
Initial swarm	Random	Benchmark seeded as particle 1

B. Online Neural Gain Tuner

The proposed controller does not replace the proportional-integral-derivative law. It recomputes the three gains at each sample from two locally measurable signals, the area control error and its filtered derivative. A network with two inputs, four hyperbolic tangent neurons and three linear outputs produces a bounded correction to each gain, as expressed in Equation (15) (Hornik et al., 1989).

$$K_j(t) = K_{0,j} + \mathrm{span}_j \cdot \tanh\left(z_j(t)\right) \qquad (15)$$

where $j = 1,2,3$ indexes the proportional, integral and derivative gains, $K_j(t)$ is the gain applied at time t, $K_{0,j}$ is the corresponding benchmark gain, $z_j(t)$ is the jth network output and span_j is the larger of $K_{0,j}$ and 0.5. The result is clipped below at zero so that no gain becomes negative.

Two properties follow from this construction. First, the gains cannot leave a bounded interval around the benchmark values whatever the network outputs, which limits the damage an out-of-distribution input can do. Second, setting the output layer weights and biases to zero makes every correction vanish, so the network reduces exactly to the fixed-gain benchmark. That configuration is placed in the initial swarm, which means the optimiser starts from a solution that already matches the benchmark and can only improve on it.

The network is trained by particle swarm optimisation directly against the closed-loop objective, not by supervised regression against the output of an existing controller (Safari et al., 2021; Tolba et al., 2026). This distinction is the reason the imitation bound described in Section I does not apply (Ross et al., 2011). The network is not attempting to reproduce any controller, so its performance is not capped by one. Its weights are searched in the interval from −3 to 3, and the input signals are scaled by the reciprocal of their standard deviations measured on the standard response.

C. Objective Function and Scenario Design

The objective is the integral of time-multiplied absolute error (ITAE) accumulated over both area frequencies and the tie-line power and summed across the scenario set, as given in Equation (16) (Sahib & Ahmed, 2016). Weighting by time penalises errors that persist and tolerates the initial deviation, which is set by inertia and the rate limit rather than by the controller.

$$J = \sum_{s \in S} \int_0^T (|\Delta f_1| + |\Delta f_2| + |\Delta P_{\mathrm{tie}}|) t \, dt \qquad (16)$$

where J is the objective, S is the set of tuning scenarios, s is a scenario in that set and $T = 60$ s is the simulation horizon. A terminal penalty proportional to the residual deviations at the end of the horizon is added so that a controller which fails to settle within 60 s is not rewarded.

Six scenarios are used for tuning and six further scenarios are reserved for testing. The tuning set spans load steps from 0.003 to 0.025 pu in either area, a simultaneous disturbance in both areas and the photovoltaic cloud transient. Tuning across this range rather than on a single disturbance is necessary. When the objective contains only a disturbance in Area 1, the optimiser drives the integral gain of Area 2 to zero, because an undisturbed area needs no integral action, and the resulting controller then fails on any disturbance in that area. Both scenario sets are listed in Table 3. The held-out set includes one scenario, H6, a load step of 0.035 pu that is deliberately larger than anything seen during tuning.

Table 3 Tuning and held-out scenarios

Tuning set	ΔP_{L1} (pu)	ΔP_{L2} (pu)	Held-out set	ΔP_{L1} (pu)	ΔP_{L2} (pu)
T1: step, Area 1	0.003	0	H1: step, Area 1	0.020	0
T2: step, Area 1	0.010	0	H2: step, Area 1	0.005	0
T3: step, Area 1	0.025	0	H3: step, Area 2	0	0.015
T4: step, Area 2	0	0.012	H4: both areas	0.005	0.010

T5: both areas	0.008	0.015	H5: PV cloud plus load	0.010	0
T6: PV cloud only	0	0	H6: step, Area 1 (extrapolation)	0.035	0

IV. RESULTS AND DISCUSSION

This section reports the validation of the model against analytic results, evidence that the rate constraint is active, the outcome of both optimisations, the transient performance on the nominal disturbance, the behaviour on held-out scenarios and the sensitivity to the value of the rate limit. All simulations use a fixed-step fourth-order Runge-Kutta integrator at 2 ms for the reported results and 20 ms within the optimiser inner loop, the latter verified to agree with the former to within 0.1% on the objective.

A. Model Validation

With supplementary control disabled, the steady-state response to a load step is set entirely by the combined area frequency response characteristics and can be computed in closed form (Elgerd, 1982; Kundur, 1994). For a 0.01 pu step in Area 1, the analytic frequency deviation is the negative of the step divided by the sum of the two bias factors whilst the analytic tie-line deviation is that quantity weighted by the bias factor of Area 2. The simulation returns −0.013687 Hz and −0.005232 pu against analytic values of −0.013687 Hz and −0.005232 pu respectively. The agreement to six decimal places confirms that the tie-line integrator, the bias factors and the droop paths are correctly connected.

B. Activation of the Generation Rate Constraint

Figure 2 compares the rate the controller demands of the turbine with the rate the constraint permits. The demanded rate exceeds the limit by more than two orders of magnitude in the first seconds after the disturbance, and the constraint is active for 10.4% of the 60 s horizon. The lower panel shows that the mechanical power rises linearly during that interval, which is the signature of a binding rate limiter rather than a first-order lag (Kalyan et al., 2022).

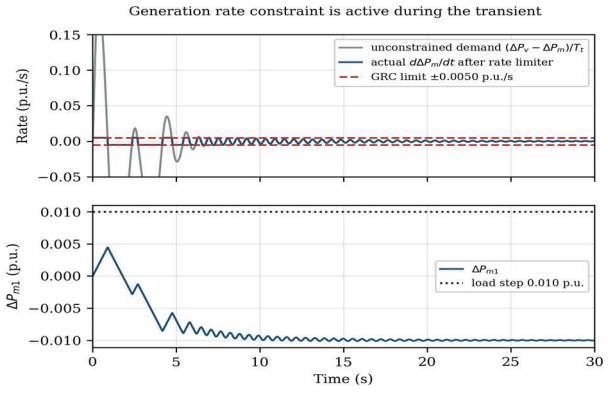


Figure 2 Demanded and permitted turbine ramp rate, and the resulting mechanical power

Figure 3 isolates the effect of the constraint by simulating the same controller with the rate limiter removed. The peak frequency deviation changes only slightly, since the first excursion is governed by inertia before the turbine can respond, but the recovery differs markedly once the constraint is in place. This establishes that the constraint is a material feature of the plant in this study and not a nominal inclusion.

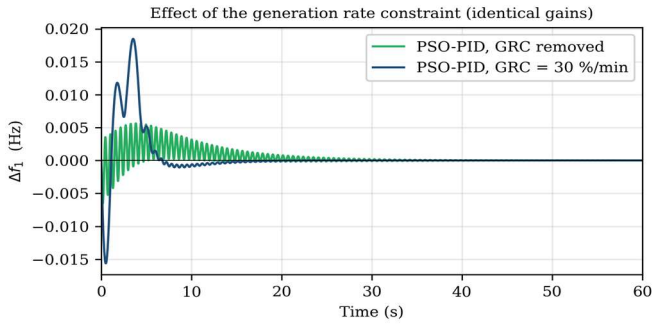


Figure 3 Area 1 frequency deviation with and without the generation rate constraint, using identical gains

C. Optimisation Results

Convergence of both optimisations is shown in Fig. 4. The benchmark search is highly reproducible, with a standard deviation of 0.0597 across ten runs and no gain reaching the upper search bound in any run. The proportional gain of Area 2 converged to zero in all ten runs, giving an integral-derivative controller in that area. This is a consistent structural outcome rather than a search artefact and is reported as found.

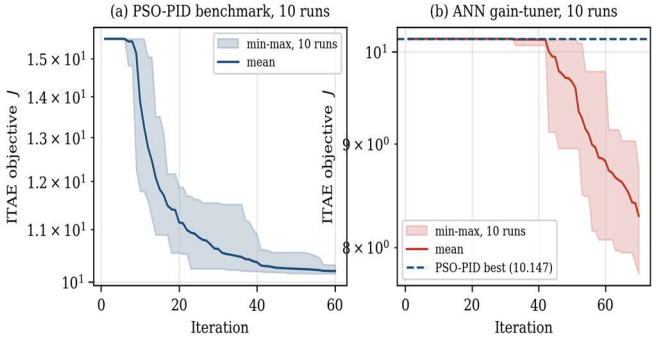


Figure 4 Convergence of the objective J over ten independent runs: (a) fixed-gain benchmark tuning and (b) neural gain tuner

The best benchmark gains are listed in Table 4, and the objective statistics of both optimisations over the ten runs are compared in Table 5.

Table 4 Tuned benchmark gains, best of ten runs

Area	K_{pi}	K_{ii}	K_{di}
Area 1	2.7073	0.5151	3.8187
Area 2	0.0000	0.1950	0.5799

Table 5 Objective statistics over ten independent runs

Controller	Best J	Mean J	Std	Worst J
PSO-tuned fixed-gain PID	10.1469	10.1961	0.0597	10.3125
Neural gain tuner	7.7631	8.2921	0.2933	8.7579
Improvement	23.5%	18.7%	–	–

The worst of the ten neural runs, at 8.7579, is still better than the best of the ten benchmark runs, at 10.1469. Since the benchmark configuration is a member of the initial swarm in every neural run, the improvement cannot be attributed to a weak baseline. Table 6 gives the breakdown by tuning scenario and shows that the advantage is not uniform. The neural controller loses 9.0% on the simultaneous disturbance and gains between 10 and 50% elsewhere, with the largest absolute reductions in the objective on the two hardest cases, the 0.025 pu step and the cloud transient.

Table 6 Objective by tuning scenario

Scenario	Fixed-gain PID	Neural gain tuner	Change
T1: 0.003 pu, Area 1	0.4905	0.2452	−50.0%
T2: 0.010 pu, Area 1	0.4926	0.3714	−24.6%
T3: 0.025 pu, Area 1	4.3001	3.3559	−22.0%
T4: 0.012 pu, Area 2	0.5754	0.5147	−10.5%
T5: 0.008 / 0.015 both areas	0.7145	0.7785	+9.0%
T6: PV cloud transient	3.5687	2.5029	−29.9%
Total	10.1417	7.7687	−23.4%

D. Transient Performance

Figure 5 shows the responses of both controllers to the nominal 0.010 pu step in Area 1, with identical axes for the two traces so that scaling does not distort the comparison. Table 7 lists the corresponding metrics, with settling time measured as the last instant at which the signal leaves a band of 2% of its own peak. Defining the band relative to the peak rather than to the final value is necessary because the final value of every signal here is zero.

Table 7 Transient metrics for a 0.010 pu step in Area 1

Metric	Fixed-gain PID	Neural gain tuner	Change (%)
Δf_1 peak undershoot (Hz)	-1.5637×10^{-2}	-1.5676×10^{-2}	+0.3
Δf_1 peak overshoot (Hz)	1.8482×10^{-2}	1.6847×10^{-2}	−8.8
Δf_1 settling time (s)	14.34	11.70	−18.4
Δf_2 peak undershoot (Hz)	-1.1734×10^{-2}	-1.0909×10^{-2}	−7.0
Δf_2 peak overshoot (Hz)	1.8020×10^{-2}	1.9784×10^{-2}	+9.8
Δf_2 settling time (s)	12.50	9.75	−22.0
ΔP_{tie} peak deviation (pu)	-3.7382×10^{-3}	-3.9078×10^{-3}	+4.5
ΔP_{tie} settling time (s)	14.31	15.23	+6.4

The pattern in Table 7 is consistent and worth stating explicitly. The first undershoot is essentially unchanged, differing by 0.3%. That excursion occurs before the turbine can contribute meaningfully, so it is set by the equivalent inertia and the rate limit, and no choice of gain can alter it (Saleem & Saha, 2024). The improvement appears in the recovery, where frequency settling shortens by 18 and 22% in the two areas. The tie-line, however, settles 6.4% later, so the gain in frequency regulation is obtained partly at the expense of the interchange, and the scheme should not be described as uniformly superior.

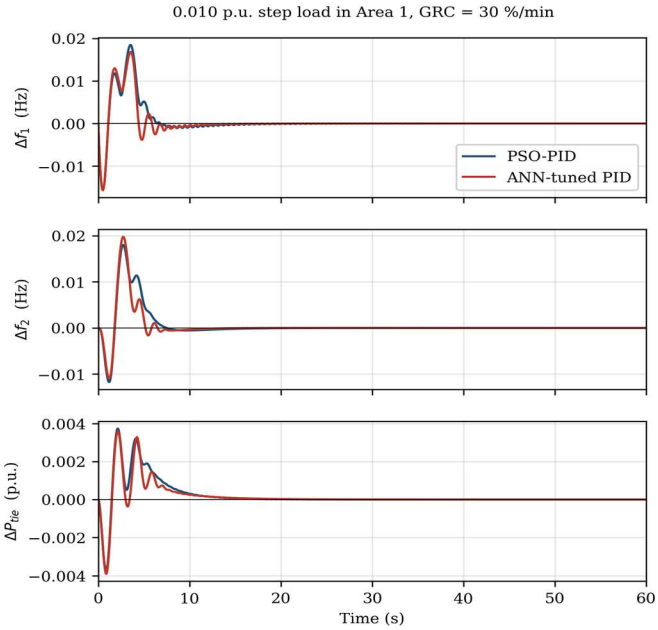


Figure 5 Frequency and tie-line responses to a 0.010 pu step load in Area 1

Figure 6 shows how the gains move during the transient. All three depart substantially from their baseline values and return toward them as the error decays, confirming that the network is exercising its authority rather than remaining at the benchmark configuration it started from. The bounded form of Equation (15) keeps every gain within a fixed interval throughout.

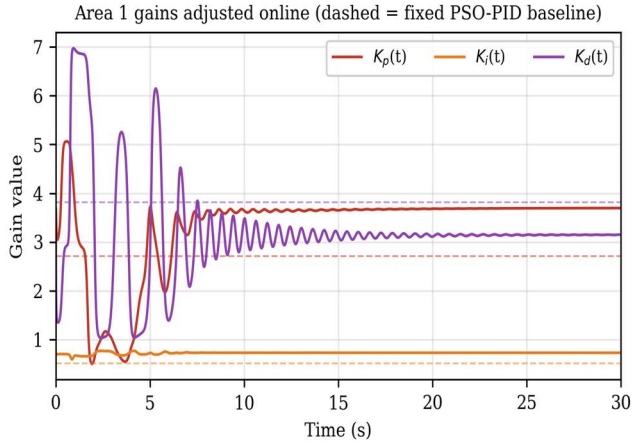


Figure 6 Area 1 gains adjusted online, with dashed lines showing the fixed benchmark values

E. Adaptability to Unseen Disturbances

Performance on the six held-out scenarios is given in Table 8. Within the trained amplitude range the neural controller retains a clear advantage, improving by between 12.8 and 49.1% on four of the six cases. It is 5.0% worse on the combined cloud and load disturbance, which is the only held-out case that combines two disturbance types. On the extrapolation case H6, a load step 40% larger than the largest used in tuning, the two controllers are indistinguishable, differing by 0.1%. The same results are plotted in Figure 7 on a logarithmic scale, which keeps the small within-range cases visible beside the much larger objective of H6.

Table 8 Objective on held-out scenarios

Scenario	Fixed-gain PID	Neural gain tuner	Change
H1: 0.020 pu, Area 1	1.9921	1.7365	−12.8%
H2: 0.005 pu, Area 1	0.4347	0.2212	−49.1%
H3: 0.015 pu, Area 2	1.1188	0.6721	−39.9%
H4: 0.005 / 0.010 both areas	0.4198	0.3363	−19.9%
H5: PV cloud plus 0.010 pu load	3.4921	3.6660	+5.0%
H6: 0.035 pu, Area 1 (extrapolation)	17.4328	17.4468	+0.1%
Total	24.8904	24.0789	−3.3%

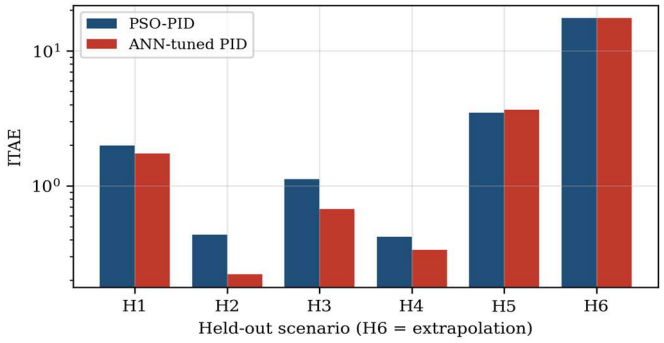


Figure 7 Objective J on the six held-out scenarios, logarithmic scale

The behaviour on H6 is the property that matters most for practical use. Bounding the gain corrections means that an input outside the trained range drives the hyperbolic tangent towards saturation, and the resulting gains sit near the edge of a fixed interval rather than taking arbitrary values (Razmi et al., 2025; Shaharear et al., 2026). The controller therefore reverts towards benchmark behaviour instead of destabilising. An earlier version of this work, trained on a narrower amplitude range, was 176% worse than the benchmark on the same extrapolation case, which shows that this property depends on the training range and not on the bounded form alone.

Figure 8 shows the response to the held-out cloud transient. Both controllers track the slow irradiance change without difficulty, and the difference between them is small, which is consistent with the 5.0% figure in Table 8. A ramped renewable disturbance of this duration is easier for a rate-limited system to follow than a step of equivalent magnitude, since the constraint is less likely to bind (Agüera-Pérez et al., 2024).

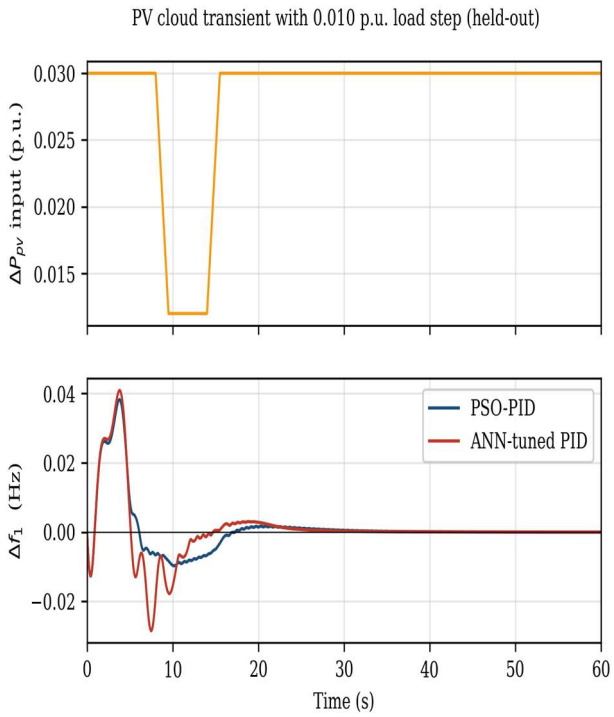


Figure 8 Response to the held-out photovoltaic cloud transient with a simultaneous load step

F. Rate Limit Sensitivity

The generation rate constraint was fixed at 30% per minute for all preceding results. Table 9 shows how both controllers behave when that value changes, with the benchmark re-tuned at each setting and the neural controller left as trained. The 30% per minute row is labelled as the design point because it is the only setting at which the neural tuner was trained, and at every other setting the tuner runs unchanged.

Table 9 Sensitivity to the generation rate constraint

Rate limit	PID (re-tuned) J	Neural tuner J	Constraint active
Removed	0.1660	6.4383	0.0%
100%/min	1.4950	6.7652	4.4%
50%/min	4.0649	7.0021	6.9%
30%/min	10.1459	7.9369	10.4%
20%/min	23.7353	108.4548	13.5%
10%/min	117.7866	608.7499	97.0%

Figure 9 plots the same objectives against the rate limit, together with the fraction of each run for which the constraint is active. Two conclusions follow from this experiment. First, removing the constraint improves the achievable objective from 10.1459 to 0.1660, a factor of roughly sixty, which identifies the rate limit rather than the control law as the dominant limitation on this system. Any effort to improve frequency regulation here competes for the small margin that the constraint leaves. This quantifies a claim that is often made in the literature without measurement (Ramjug-Ballgobin & Ramlukon, 2021; Saini & Ohri, 2023).

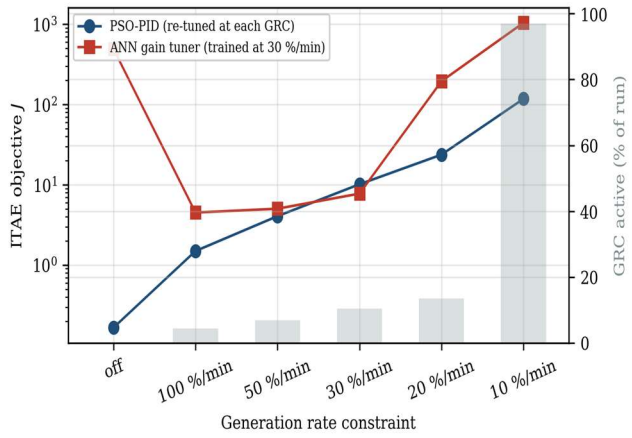


Figure 9 Objective against the generation rate constraint, with the fraction of the run for which the constraint is active

The second conclusion concerns the neural controller. It is superior only at the rate limit for which it was trained, and at every other setting, most sharply at 20 and 10% per minute, it is worse than a re-tuned benchmark. The gain law is therefore specific to the plant it was trained against and does not transfer across changes in the rate constraint. A fair comparison at each setting would require retraining, which was not performed. As presented, Table 9 measures sensitivity to plant mismatch, not relative controller quality across the range.

V. LIMITATIONS

Six limitations are stated so that the scope of the results is not over-read. First, the gain law is trained at a single value of the generation rate constraint and does not transfer to other values, as Section IV.F shows. Second, the improvement is confined to damping and settling, and the first frequency excursion is unchanged, so the scheme does not address the quantity that determines whether under-frequency protection operates. The remaining four limitations concern generalisation and scope. The advantage is not uniform across scenarios, with small losses on combined disturbances offsetting larger gains elsewhere. Outside the trained amplitude range the controller reverts to benchmark behaviour rather than continuing to improve. Training is done offline, so the controller adapts its gains online but does not learn online, and no stability proof is offered for the resulting time-varying closed loop (Awad et al., 2026; Das et al., 2026; Li et al., 2025; Shaharear et al., 2026). Finally, the photovoltaic representation is a first-order lag with a time-varying input rather than a detailed model of the array and its converter.

VI. CONCLUSION

A two-area load frequency control system was modelled with an explicit tie-line synchronising integrator, valve position limits and a generation rate constraint acting on the turbine mechanical output. The constraint was converted from 30% per minute to 5.000×10^{-3} pu per second and shown to be active for 10.4% of the simulated horizon, and removing it improved the achievable objective by a factor of about sixty. A feed-forward network with two inputs and four hidden neurons, adjusting the three controller gains online from the area control error and its filtered derivative, reduced the objective by 23.5% against a benchmark tuned over the same scenarios, with the same optimiser and budget, and with the benchmark included in the initial swarm. Frequency settling improved by 18.4 and 22.0% in the two areas while the first undershoot was unchanged, which is consistent with that excursion being set by inertia and the rate limit rather than by the controller. On held-out disturbances within the trained amplitude range the advantage persisted, and beyond that range the controller reverted to benchmark behaviour rather than destabilising.

Three directions follow from this work. First, the gain law could be made amplitude invariant by normalising the network input by a running estimate of the error magnitude, which would remove the dependence on the trained disturbance range. Second, training across several rate limits would test whether a single gain law can cover a range of plant conditions. Third, a stability analysis of the resulting time-varying closed loop is required before any hardware trial, since the present results establish performance by simulation alone.

Author Contributions

R. A. Alao: Conceptualization, Methodology, Software, Formal analysis, Supervision. **I. A. Aleshinloye:** Methodology, Software, Validation, Writing – original draft. **F. W. Shittu:** Methodology, Software, Investigation, Writing – original draft. **I. O. Kolawole**: Writing – original draft, Writing – review & editing. **A. O. Otuoze:** Conceptualization, Methodology, Supervision

ACKNOWLEDGEMENT

The authors are grateful to the Department of Electrical and Electronics Engineering, University of Ilorin, Ilorin, Nigeria, for the support during the compilation of this work. We specially thank NJTD for organizing the free research training workshop that gave the authors valuable hands-on experience in publishing. We are also deeply grateful to the reviewers for their constructive and valuable comments, which significantly improved the quality of this article.

REFERENCES

Agüera-Pérez, A., Espinosa-Gavira, M. J., Palomares-Salas, J. C., González-de-la-Rosa, J. J., Sierra-Fernández, J. M., & Florencias-Oliveros, O. (2024). Meteorological contexts in the analysis of cloud-induced photovoltaic transients: A review. *Renewable and Sustainable Energy Reviews, 202*, 114689. https://doi.org/10.1016/j.rser.2024.114689.

Åström, K. J., & Hägglund, T. (2006). *Advanced PID control*. ISA - The Instrumentation, Systems, and Automation Society.

Awad, M. A., Attia, M. A., & EL-Ebiary, A. H. (2026). Intelligent adaptive frequency regulation of interconnected power networks under renewable uncertainty and time delays. *Scientific Reports, 16*(1), 19933. https://doi.org/10.1038/s41598-026-59656-8.

Ayman, M., Attia, M. A., & Asim, A. M. (2026). An adaptive model predictive control approach for robust load frequency control under renewable energy disturbances. *Scientific Reports, 16*(1), 2607. https://doi.org/10.1038/s41598-025-33986-5.

Bakeer, A., Magdy, G., Chub, A., & Bevrani, H. (2022). A sophisticated modeling approach for photovoltaic systems in load frequency control. *International Jour nal of Electrical Power & Energy Systems, 134*, 107330. https://doi.org/10.1016/j.ijepes.2021.107330.

Bevrani, H. (2014). *Robust power system frequency control* (2nd ed.). Springer. https://doi.org/10.1007/978-3-319-07278-4.

Biswas, D. K., Debbarma, S., & Singh, P. P. (2025). LFC of multi-area power system using a robust full-order terminal sliding mode controller considering communication delay. *Electrical Engineering, 107*(4), 4421–4435. https://doi.org/10.1007/s00202-024-02753-6.

Can, Ö., Ayaş, M. Ş., & Şahin, A. K. (2026). Hippopotamus optimization–tuned sigmoid PID controller for load frequency control of a two-area thermal power system with renewable energy sources. *Scientific Reports, 16*(1), 11763. https://doi.org/10.1038/s41598-026-41620-1.

Clerc, M., & Kennedy, J. (2002). The particle swarm - explosion, stability, and convergence in a multidimensional complex space. *IEEE Transactions on Evolutionary Computation, 6*(1), 58–73. https://doi.org/10.1109/4235.985692.

Das, S. K., Swain, S. C., Nayak, B., Dash, R., Osumanu, F., Mohapatra, B., Al Bustanji, S., & Rubanenko, O. (2026). Adaptive long short term memory-PI based control for load frequency regulation in electric vehicle integrated two area power systems. *Scientific Reports, 16*(1), 20611. https://doi.org/10.1038/s41598-026-49840-1.

Dash, R., Reddy, K. J., Mohapatra, B., Bajaj, M., & Zaitsev, I. (2025). An approach for load frequency control enhancement in two-area hydro-wind power systems using LSTM + GA-PID controller with augmented Lagrangian methods. *Scientific Reports, 15*(1), 1307. https://doi.org/10.1038/s41598-025-85639-2.

Elgerd, O. I. (1982). *Electric energy systems theory: An introduction* (2nd ed.). McGraw-Hill.

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.

El-Hameed, M. A., Saeed, M., Kabbani, A., & Abd El-Hay, E. (2025). Efficient load frequency controller for a power system comprising renewable resources based on deep reinforcement learning. *Scientific Reports, 15*(1), 18379. https://doi.org/10.1038/s41598-025-03310-2.

El-Rifaie, A. M., Abid, S., Ginidi, A. R., & Shaheen, A. M. (2025). Fractional order PID controller based-neural network algorithm for LFC in multi-area power systems. *Engineering Reports, 7*(2), e70028. https://doi.org/10.1002/eng2.70028.

Gonzalez-Salazar, M. A., Kirsten, T., & Prchlik, L. (2018). Review of the operational flexibility and emissions of gas- and coal-fired power plants in a future with growing renewables. *Renewable and Sustainable Energy Reviews, 82*, 1497–1513. https://doi.org/10.1016/j.rser.2017.05.278.

Gulzar, M. M., Sibtain, D., Alqahtani, M., Alismail, F., & Khalid, M. (2025). Load frequency control progress: A comprehensive review on recent development and challenges of modern power systems. *Energy Strategy Reviews, 57*, 101604. https://doi.org/10.1016/j.esr.2024.101604.

Halim, A. H., Ismail, I., & Das, S. (2021). Performance assessment of the metaheuristic optimization algorithms: An exhaustive review. *Artificial Intelligence Review, 54*(3), 2323–2409. https://doi.org/10.1007/s10462-020-09906-6.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks, 2*(5), 359–366. https://doi.org/10.1016/0893-6080(89)90020-8.

Kalyan, C. N. S., Goud, B. S., Reddy, C. R., Kumar, M. K., Bajaj, M., El-Naggar, M. F., & Kamel, S. (2022). Performance assessment of open-loop and closed-loop generation rate constraint models for optimal LFC of the three-area reheat thermal system. *Frontiers in Energy Research, 10*, 920651. https://doi.org/10.3389/fenrg.2022.920651.

Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of ICNN'95 - International Conference on Neural Networks, 4*, 1942–1948. https://doi.org/10.1109/ICNN.1995.488968.

Khalil, A. E., Boghdady, T. A., Alham, M. H., & Ibrahim, D. K. (2023). Enhancing the conventional controllers for load frequency control of isolated microgrids using proposed multi-objective formulation via artificial rabbits optimization algorithm. *IEEE Access, 11*, 3472–3493. https://doi.org/10.1109/ACCESS.2023.3234043.

Kumar, S., Gupta, A., & Bindal, R. K. (2024). Load-frequency and voltage control for power quality enhancement in a SPV/wind utility-tied system using GA and PSO optimization. *Results in Control and Optimization, 16*, 100442. https://doi.org/10.1016/j.rico.2024.100442.

Kundur, P. (1994). *Power system stability and control*. McGraw-Hill.

Li, Y., Gao, S., Chen, X., Fan, D., & Zhang, M. (2025). Load frequency control of power systems based on deep reinforcement learning with leader–follower consensus control for state of charge. *Processes, 13*(11), 3669. https://doi.org/10.3390/pr13113669.

Oladipo, S., Sun, Y., & Adegoke, S. A. (2025). Application of soft computing techniques to load frequency control system of electric power systems: A brief survey. *Cogent Engineering, 12*(1), 2572297. https://doi.org/10.1080/23311916.2025.2572297.

Paliwal, N., Srivastava, L., & Pandit, M. (2022). Application of grey wolf optimization algorithm for load frequency control in multi-source single area power system. *Evolutionary Intelligence, 15*(1), 563–584. https://doi.org/10.1007/s12065-020-00530-5.

Park, Y.-K., Moon, S.-W., & Kim, T.-S. (2021). Advanced control to improve the ramp-rate of a gas turbine: Optimization of control schedule. *Energies, 14*(23), 8024. https://doi.org/10.3390/en14238024.

Patra, K. C., & Patnaik, A. (2026). Automatic generation control and load frequency control with renewable energy integration in power system using digital deadbeat controller. *Discover Electronics, 3*(1), 39. https://doi.org/10.1007/s44291-026-00193-z.

Qu, Z., Younis, W., Liu, X., Junejo, A. K., Almutairi, S. Z., & Wang, P. (2024). Optimized PID controller for load

frequency control in multi-source and dual-area power systems using PSO and GA algorithms. *IEEE Access, 12*, 186658–186678. https://doi.org/10.1109/ACCESS.2024.3445165.

Rahnama, H. S. A. (2020). Designing a PD-(1+PI) controller for LFC of an entirely renewable microgrid using PSO-TVAC. *International Journal on Technical and Physical Problems of Engineering, 12*(45), 19–27.

Ramjug-Ballgobin, R., & Ramlukon, C. (2021). A hybrid metaheuristic optimisation technique for load frequency control. *SN Applied Sciences, 3*(5), 547. https://doi.org/10.1007/s42452-021-04482-y.

Razmi, P., Rahimi, T., Sabahi, K., Gheisarnejad, M., & Khooban, M.-H. (2025). Adaptive fuzzy gain scheduling PID controller for frequency regulation in modern power system. *IET Renewable Power Generation, 19*(1), e12569. https://doi.org/10.1049/rpg2.12569.

Ross, S., Gordon, G. J., & Bagnell, J. A. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. *Proceedings of Machine Learning Research, 15*, 627–635.

Safari, A., Babaei, F., & Farrokhifar, M. (2021). A load frequency control using a PSO-based ANN for micro-grids in the presence of electric vehicles. *International Journal of Ambient Energy, 42*(6), 688–700. https://doi.org/10.1080/01430750.2018.1563811.

Sahib, M. A., & Ahmed, B. S. (2016). A new multiobjective performance criterion used in PID tuning optimization algorithms. *Journal of Advanced Research, 7*(1), 125–134. https://doi.org/10.1016/j.jare.2015.03.004.

Saini, N., & Ohri, J. (2023). Optimal load frequency control of a multi-area power system with dead band effect and generation rate constraints. *Majlesi Journal of Electrical Engineering, 17*(1). https://doi.org/10.30486/mjee.2023.1970197.0.

Saleem, M. I., & Saha, S. (2024). Assessment of frequency stability and required inertial support for power grids with high penetration of renewable energy sources. *Electric Power Systems Research, 229*, 110184. https://doi.org/10.1016/j.epsr.2024.110184.

Shaharear, M. A., Zhou, C., Shaikh, S., & Faruk, M. M. H. (2026). Load frequency control in multi-area power systems using incremental proportional–integral–derivative and model-free adaptive control. *Applied System Innovation, 9*(3), 59. https://doi.org/10.3390/asi9030059.

Shahzad, M. I., Gulzar, M. M., Shahzad, A., Arishi, A., & Murtaza, A. F. (2026). From classical to AI-driven load frequency control: Addressing smart grid challenges with renewable energy sources and EVs integration. *Renewable and Sustainable Energy Reviews, 226*, 116207. https://doi.org/10.1016/j.rser.2025.116207.

Shi, Y., & Eberhart, R. (1998). A modified particle swarm optimizer. *Proceedings of the IEEE International Conference on Evolutionary Computation*, 69–73. https://doi.org/10.1109/ICEC.1998.699146.

Soliman, A. M. A., Bahaa Eldin, M., & Mehanna, M. A. (2022). Application of WOA tuned type-2 FLC for LFC of two area power system with RFB and solar park considering TCPS in *interline*. *IEEE Access, 10*, 112007–112018. https://doi.org/10.1109/ACCESS.2022.3215530.

Tolba, M. S., Omera, A., Zaery, M., & Abido, M. A. (2026). An adaptive neural network-driven PID control framework for load frequency regulation in renewable-integrated power system. *Arabian Journal for Science and Engineering, 51*(8), 11325–11344. https://doi.org/10.1007/s13369-025-10989-y.

Wang, Z., & Liu, Y. (2021). Adaptive terminal sliding mode based load frequency control for multi-area interconnected power systems with PV and energy storage. *IEEE Access, 9*, 120185–120192. https://doi.org/10.1109/ACCESS.2021.3109141.

Wärtsilä. (2024). *Wärtsilä 34SG gas engine product guide*. Wärtsilä Corporation.

Yakout, A. H., Dashtdar, M., AboRas, K. M., Ghadi, Y. Y., Elzawawy, A., Yousef, A., & Kotb, H. (2024). Neural network-based adaptive PID controller design for over-frequency control in microgrid using honey badger algorithm. *IEEE Access, 12*, 27989–28005. https://doi.org/10.1109/ACCESS.2024.3367288.

Younis, W., Yameen, M. Z., Junejo, A. K., Abido, M. A., & Khalid, M. (2025). Robust load frequency control in renewable integrated multi-area grids using hybrid SA and QIO tuned PIDF controller. *Scientific Reports, 15*(1), 36848. https://doi.org/10.1038/s41598-025-20698-z.