

Deep Learning-Enabled Adaptive Control Of Power Converters For Next-Generation Renewable Energy Integration

Ayush Kumar¹, Baby Khooshbu¹, Kamini Kumari¹, Ajay Kumar², Ozair Ahmad³

¹Assistant Professor, Department of Electronics and Communication Engineering, Saharsa College of Engineering, Saharsa

²Assistant Professor, Department of Electrical Engineering, Saharsa College of Engineering, Saharsa

³Assistant Professor, Department of Electronics and Communication Engineering, Government Engineering College, Kishanganj.

Abstract

The accelerating penetration of solar photovoltaic (PV) and wind-based distributed generation has exposed the structural limitations of fixed-gain proportional-integral (PI) and classical linear controllers used in grid-interfacing power converters. This paper proposes a deep learning-enabled adaptive control framework that combines a Long Short-Term Memory (LSTM) feature extractor with a Deep Deterministic Policy Gradient (DDPG) reinforcement-learning agent to regulate DC-link voltage, track the maximum power point, and suppress harmonic distortion under rapidly varying irradiance, temperature, and load conditions. The proposed controller continuously learns an optimal duty-ratio/modulation policy from measured voltage, current, irradiance, and temperature states, eliminating the need for manual gain re-tuning. A MATLAB/Simulink-based hardware-representative model of a two-stage grid-connected converter was used to benchmark the proposed method against conventional PI, fuzzy-logic, and artificial neural network (ANN) controllers. Simulation results show that the proposed controller reduces settling time by up to 81.6%, overshoot by up to 87.3%, and total harmonic distortion (THD) by up to 70.7% relative to the conventional PI controller, while improving maximum power point tracking (MPPT) efficiency to 99.6%. These findings demonstrate that deep learning-enabled adaptive control offers a scalable and disturbance-resilient pathway for next-generation renewable energy integration.

Keywords: Deep learning; adaptive control; power converters; renewable energy integration; reinforcement learning; LSTM; DDPG; maximum power point tracking.

1. Introduction

The global transition toward decarbonized electricity generation has made power-electronic converters the primary interface between renewable energy sources (RES) and the utility grid. Unlike conventional synchronous generation, PV arrays and wind turbines exhibit highly non-linear and time-varying power characteristics that are governed by irradiance, temperature, wind speed, and partial-shading effects. Power converters — boost/buck DC-DC stages, voltage-source inverters (VSI), and dual-active-bridge topologies — must therefore continuously adapt their switching behaviour to maintain stable DC-link voltage, extract maximum available power, and inject a low-distortion current into the grid.

Classical control approaches, including proportional-integral (PI) and proportional-integral-derivative (PID) regulators, remain dominant in commercial inverters because of their simplicity and well-understood tuning rules. However, these linear controllers are designed around a single operating point and rely on fixed gains obtained through small-signal linearization. When irradiance or load changes abruptly, the fixed-gain assumption breaks down, producing large overshoots, slow settling, and elevated total harmonic distortion (THD), particularly in weak-grid or high-penetration scenarios (Xing et al., 2024). Fuzzy-logic and sliding-mode controllers partially address non-linearity but still depend on hand-crafted rule bases or switching surfaces that do not generalize well across converter topologies and disturbance profiles.

The stakes of this control problem are rising rapidly. National grid codes increasingly require renewable power plants to provide fault ride-through, reactive-power support, and frequency response services that were traditionally the domain of synchronous generation, placing additional dynamic-performance demands on the converter controller. At the same time, the falling cost of embedded graphics-processing units and neural-network accelerators has made it computationally feasible to deploy inference-only deep-learning models directly on converter control boards, closing the gap between offline model training and real-time onboard execution. These two trends — more demanding grid requirements and more capable embedded hardware — motivate a re-examination of adaptive, learning-based control as a practical rather than purely theoretical option for commercial power converters.

Recent advances in deep learning and deep reinforcement learning (DRL) offer a data-driven alternative that learns a control policy directly from interaction with the converter environment, without requiring an explicit analytical plant model. Reinforcement-learning agents such as Deep Deterministic Policy Gradient (DDPG) and Twin-Delayed DDPG have already demonstrated strong performance in voltage-source-converter parameter tuning (Xing et al., 2024) and grid-connected inverter control with LCL filters (Rajamalliah et al., 2024). Independently, recurrent architectures such as Long Short-Term Memory (LSTM) networks have proven effective at forecasting irradiance-driven power fluctuations and improving maximum power point tracking (MPPT) accuracy under rapidly varying conditions (Younas et al., 2024; Roy et al., 2024). This paper integrates these two complementary capabilities into a single deep learning-enabled adaptive controller. An LSTM block extracts temporal features from the measured state trajectory (voltage, current, irradiance, temperature), and a DDPG agent maps these features to a continuous control action — the converter duty ratio or modulation index — guided by a reward function that jointly penalizes voltage ripple, THD, and MPPT tracking error. The specific contributions of this work are: (i) a unified LSTM-DDPG control architecture applicable to both DC–DC and DC–AC converter stages; (ii) a composite reward formulation that balances dynamic response, power quality, and steady-state accuracy; and (iii) a comparative simulation study against PI, fuzzy, and ANN controllers under standardized irradiance-step and load-step disturbances. The remainder of the paper is organized as follows: Section 2 reviews related work; Section 3 presents the proposed system architecture; Section 4 details the deep learning controller design; Section 5 reports simulation results; Section 6 discusses implications and limitations; and Section 7 concludes the paper.

2. Related Work

Data-driven control of power converters has progressed rapidly over the last two years. Xing et al. (2024) proposed a deep-reinforcement-learning framework for tuning voltage-source-converter control parameters in renewable generation systems, showing improved stability margins compared with fixed-gain benchmark models. Rajamalliah et al. (2024) applied DRL directly to a grid-connected inverter with an LCL filter, demonstrating robust current-control performance for renewable solar applications without requiring re-tuning across operating points. Pei et al. (2024) introduced a model-free DRL algorithm with enhanced grid visibility for voltage regulation in distribution networks using smart PV inverters, highlighting the value of learning-based approaches when accurate network models are unavailable.

On the converter-side MPPT problem, Younas et al. (2024) developed a stacked LSTM-based MPPT controller for a 100 kWp dual-stage grid-tied PV system and reported superior tracking speed relative to feed-forward neural networks and perturb-and-observe methods. Roy et al. (2024) validated an LSTM-based MPPT approach using real-


world irradiance and temperature data, confirming that recurrent architectures capture the temporal dependencies of solar power fluctuations more effectively than static neural networks. Complementary optimization-based approaches include the PSO-ANFIS MPPT controller for fuel-cell-fed DC-DC converters (Hussaian Basha, 2024) and the super-twisting sliding-mode/grey-wolf-optimizer hybrid MPPT strategy of Deghfel et al. (2024), both of which improved tracking efficiency but remained sensitive to controller-specific parameter selection. Idris et al. (2024) provided a comparative evaluation of classical (perturb-and-observe, incremental conductance) and AI-based (fuzzy-logic) MPPT techniques, concluding that intelligent controllers offer consistently lower steady-state ripple across variable irradiance profiles.

Collectively, this literature indicates two parallel trends: reinforcement learning is well suited to closed-loop converter regulation and disturbance rejection, while recurrent deep-learning models are well suited to short-horizon forecasting and MPPT. However, most existing studies address these two capabilities separately. The present work distinguishes itself by fusing an LSTM temporal-feature extractor directly into the state representation of a DDPG control agent, allowing a single learned policy to simultaneously perform predictive MPPT and adaptive voltage/current regulation for a grid-connected renewable power converter.

3. Proposed System Architecture

The proposed system, illustrated in Figure 1, consists of a renewable energy source (PV array or wind turbine with rectifier front-end), a power converter stage (boost DC-DC converter cascaded with a voltage-source converter), an LCL output filter, and the point of common coupling with the utility grid or local load. Voltage, current, irradiance, and temperature sensors continuously feed a measurement vector to the deep learning-enabled adaptive controller, which computes the duty ratio (for the DC-DC stage) and the modulation index/reference current (for the grid-side VSC) at each control interval. Unlike a conventional cascaded PI structure, the controller in this work is a single learned policy that receives the full state vector and outputs the complete control action, with the internal LSTM layer implicitly capturing the dynamic coupling between the DC-link and the AC-side current loop. The controller is trained offline in a simulated environment using episodes that replay realistic irradiance and load profiles, and the trained policy is then deployed for real-time inference, where forward-pass computation is fast enough to meet the control-loop timing requirements of a typical 10–20 kHz switching converter.

Figure 1. Proposed deep learning-enabled adaptive control architecture for renewable energy power converters.



4. Deep Learning-Based Adaptive Controller Design

4.1 State, Action, and Reward Formulation

The controller state at time step t is defined as $s_t = [v(t), i(t), P(t), G(t), T(t), \Delta v(t)]$, comprising the measured DC-link voltage, converter current, instantaneous power, solar irradiance, module temperature, and the voltage deviation from reference. The continuous action at corresponds to the duty ratio of the DC-DC converter (bounded within $[0.1, 0.9]$) and the modulation index of the grid-side inverter. The reward function is formulated as a weighted composite:

$$r(t) = -w_1 \cdot |V_{ref} - V(t)| - w_2 \cdot THD(t) - w_3 \cdot |P_{mpp}(t) - P(t)| - w_4 \cdot \Delta D(t)^2$$


where w_1 – w_4 are tuning weights that balance voltage regulation, harmonic suppression, MPPT tracking error, and control-effort smoothness (ΔD is the change in duty ratio between consecutive steps, included to discourage chattering). This composite formulation follows the general design philosophy adopted in recent DRL-based converter and inverter controllers (Rajamalliah et al., 2024; Pei et al., 2024), adapted here to jointly cover DC-side and AC-side objectives.

4.2 Network Architecture and Training

The controller comprises two functional blocks. First, a two-layer LSTM network (64 and 32 hidden units) processes a sliding window of the last 20 state samples to extract temporal features that capture ramp-rate and trend information in irradiance and load. Second, these features are passed to a DDPG agent consisting of an actor network (three fully connected layers of 256-128-64 neurons with ReLU activations and a tanh output layer) and a critic network of matching depth that estimates the action-value function. The actor and critic are trained using the Adam optimizer with a learning rate of 1×10^{-4} (actor) and 1×10^{-3} (critic), a discount factor $\gamma = 0.99$, soft target-update rate $\tau = 0.005$, and a replay buffer of 10^6 transitions. Exploration noise is introduced using an Ornstein-Uhlenbeck process during training and annealed to zero prior to deployment. Training was carried out over 500 episodes, each simulating 1 second of converter operation with randomized irradiance-step and load-step disturbances, consistent with training protocols reported for DRL-based converter controllers in recent literature (Xing et al., 2024; Younas et al., 2024). Figure 2 shows the resulting reward-convergence and actor-critic loss curves, confirming stable policy convergence after approximately 350 episodes.

Figure 2. Training performance of the proposed deep reinforcement learning controller: (a) reward convergence; (b) actor-critic network loss.

Figure 2. Training performance of the proposed deep reinforcement learning controller



5. Simulation Setup and Results

The proposed controller was benchmarked in a MATLAB/Simulink model of a 5 kW two-stage grid-connected PV converter (boost stage + three-phase VSC) with a 400 V DC-link reference and a switching frequency of 20 kHz. Four controllers were compared under identical irradiance-step ($1000 \rightarrow 600 \rightarrow 1000 \text{ W/m}^2$) and resistive load-step disturbances: (i) a conventional PI controller tuned using the Ziegler–Nichols method, (ii) a fuzzy-logic controller with a 49-rule base, (iii) a feed-forward ANN controller trained by supervised learning on PI-generated trajectories, and (iv) the proposed LSTM-DDPG adaptive controller.


Table 1 summarizes the steady-state and dynamic performance metrics averaged over ten randomized disturbance trials. The proposed controller achieves the lowest settling time (12.6 ms), lowest overshoot (1.3%), lowest THD (1.42%), highest MPPT efficiency (99.6%), and lowest steady-state error (0.18%) among all four controllers, confirming the benefit of combining temporal feature learning with a continuous-action reinforcement-learning policy.

Table 1. Comparative dynamic and steady-state performance of the evaluated controllers

Controller Type	Settling Time (ms)	Overshoot (%)	THD (%)	MPPT Efficiency (%)	Steady-State Error (%)
Conventional PI	68.4	10.2	4.85	96.3	1.42
Fuzzy Logic Controller	45.7	6.8	3.60	97.5	0.95
ANN-Based Controller	31.2	4.1	2.95	98.4	0.61
Proposed DL-Based Adaptive Controller (LSTM + DDPG)	12.6	1.3	1.42	99.6	0.18

Figure 3 illustrates the DC-link voltage response of the conventional PI controller and the proposed DL-based adaptive controller under a sudden irradiance drop ($1000 \rightarrow 600 \text{ W/m}^2$) at $t = 0.30 \text{ s}$ and a subsequent irradiance rise ($600 \rightarrow 1000 \text{ W/m}^2$) at $t = 0.65 \text{ s}$. The PI controller exhibits a pronounced voltage dip/overshoot of approximately $\pm 50 \text{ V}$ and requires close to 70 ms to re-settle, whereas the proposed controller confines the deviation to within $\pm 12 \text{ V}$ and restores the reference voltage in under 15 ms, demonstrating its superior disturbance-rejection capability under realistic irradiance transients.

Figure 3. Dynamic DC-link voltage response under sudden irradiance variation: conventional PI controller versus the proposed deep learning-based adaptive controller.



Beyond the transient metrics reported in Table 1, the proposed controller also displayed consistent performance across partial-shading and unbalanced-load scenarios not shown in the main comparison, suggesting that the learned policy generalizes beyond the specific disturbance patterns encountered during training. This generalization capability is consistent with the improved robustness reported for LSTM-augmented MPPT controllers under variable real-world irradiance data (Roy et al., 2024).

5.1 Computational Considerations

A practical concern for any learning-based controller is whether its inference latency is compatible with the switching period of the converter. In the present study, the trained actor network — a lightweight three-layer feed-forward path following the LSTM feature extractor — required an average forward-pass time of approximately 38 μ s on a desktop-class processor, well within the 50 μ s control period corresponding to a 20 kHz switching frequency. This is consistent with recent benchmarking of DRL-based converter controllers, which report inference times in the tens of microseconds range for actor networks of comparable size (Xing et al., 2024). For embedded deployment on a digital signal processor or field-programmable gate array, further latency reduction can be achieved through network pruning, quantization, or fixed-point implementation, none of which materially altered tracking accuracy in preliminary offline tests conducted for this study. Training itself remains an offline, one-time computational cost; the 500-episode training run completed in under two hours on a single graphics-processing-unit workstation, after which the learned policy is frozen for field deployment and periodically retrained only if the operating envelope changes substantially.

6. Discussion

The results confirm that fusing a temporal feature extractor (LSTM) with a continuous-action reinforcement-learning agent (DDPG) yields measurable improvements over both classical and static-neural-network controllers. The performance gap is most pronounced during transient disturbances, where fixed-gain PI control cannot adapt its effective bandwidth, while the proposed controller implicitly adjusts its response based on the learned temporal context of the disturbance. These findings align with recent reports that DRL-based converter and inverter controllers outperform fixed-gain baselines in both stability margin and dynamic response (Xing et al., 2024; Rajamallaiah et al., 2024).

Nevertheless, several practical considerations remain. First, the controller's performance depends on the fidelity and diversity of the training environment; a policy trained purely in simulation may require domain-adaptation or fine-tuning before field

deployment, a concern also raised in real-world DRL voltage-control studies. Second, the computational footprint of the LSTM-DDPG inference pipeline, while modest, must be validated on the specific digital signal processor or field-programmable gate array intended for deployment to guarantee that inference completes within the converter's switching period. Third, safety-constrained variants of the reward function or action-clipping mechanisms are advisable to prevent the learned policy from exploring unsafe duty-ratio regions during the early stages of training. Addressing these considerations through hardware-in-the-loop validation is identified as a priority for future work.

7. Conclusion and Future Scope

This paper presented a deep learning-enabled adaptive control framework for power converters supporting next-generation renewable energy integration. By combining an LSTM-based temporal feature extractor with a DDPG reinforcement-learning agent, the proposed controller learns a unified policy for DC-link voltage regulation, maximum power point tracking, and harmonic suppression, without relying on fixed linear gains. Simulation results on a representative two-stage grid-connected PV converter showed substantial improvements over conventional PI, fuzzy-logic, and ANN controllers, including an 81.6% reduction in settling time, an 87.3% reduction in overshoot, and a 70.7% reduction in THD, alongside a 99.6% MPPT efficiency. These results support the viability of deep learning-enabled adaptive control as a scalable solution for high-penetration renewable power systems.

Future work will extend the proposed framework to multi-converter microgrid coordination, incorporate safety-constrained reinforcement learning to guarantee bounded control actions during online learning, and validate the controller on a hardware-in-the-loop and physical prototype platform to quantify real-time computational latency and robustness under sensor noise and communication delay.

References

1. Deghfel, N., Badoud, A. E., Al-Ahmadi, A. A., Bajaj, M., Zaitsev, I., & Ghoneim, S. S. M. (2024). Improving maximum power point tracking efficiency in solar photovoltaic systems using super-twisting algorithm and grey wolf optimizer. *IET Renewable Power Generation*, 18(15), 3329–3354. <https://doi.org/10.1049/rpg2.13138>
2. Hussaian Basha, H. (2024). A novel development of wide voltage supply DC–DC converter for fuel stack application with PSO-ANFIS MPPT controller. *Scientific Reports*, 14(1), Article 18826. <https://doi.org/10.1038/s41598-024-69802-9>
3. Idris, D. K. M., Birasa, R., Martiş, C. S., & Philippe, D. R. (2024). Comparative analysis of classical and AI-based MPPT techniques in stand-alone PV systems: Perturb and observe, incremental conductance, and fuzzy logic. In *Proceedings of the 2024 IEEE International Conference and Exposition on Electric and Power Engineering (EPEI)* (pp. 139–144). IEEE. <https://doi.org/10.1109/EPEI63510.2024.10758020>
4. Pei, Y., Ye, K., Zhao, J., Yao, Y., Su, T., & Ding, F. (2024). Enhanced-visibility model-free deep reinforcement learning algorithm for voltage control in realistic distribution systems using smart inverters. *Applied Energy*, 359, Article 122716.
5. Rajamallaiyah, A., Karri, S. P. K., Alghaythi, M. L., & Alshammari, M. S. (2024). Deep reinforcement learning based control of a grid connected inverter with LCL-filter for renewable solar applications. *IEEE Access*, 12, 22278–22295. <https://doi.org/10.1109/ACCESS.2024.3364058>
6. Roy, B., Adhikari, S., Datta, S., Devi, K. J., Devi, A. D., & Ustun, T. S. (2024). Harnessing deep learning for enhanced MPPT in solar PV systems: An LSTM approach using real-world data. *Electricity*, 5(4), 843–860. <https://doi.org/10.3390/electricity5040042>

7. Xing, Y., Zhang, G., Wang, B., et al. (2024). Modeling and deep reinforcement learning based control parameter tuning for voltage source converter in a renewable energy generation system. *Journal of Electrical Engineering & Technology*, 19, 947–959. <https://doi.org/10.1007/s42835-024-01820-z>
8. Younas, U., Kulaksiz, A. A., & Ali, Z. (2024). Deep learning stack LSTM based MPPT control of dual stage 100 kWp grid-tied solar PV system. *IEEE Access*, 12, 77555–77574.