Greenhouse Plant Production Journal

Predictive-Adaptive Coordinated Optimization for Multi-Greenhouse Energy Management

Peyman Bayat [a]*, Pezhman Bayat [a]

[a] *Department of Electrical Engineering, Hamedan University of Technology, Hamedan, Iran*

Review Article

Use your device to scan and read the article online

Citation: Bayat, P. and Bayat, P. 2026. Predictive-Adaptive Coordinated Optimization for Multi-Greenhouse Energy Management. Greenhouse Plant Production Journal 3(3): 26-50.

https://doi.org/10.66224/gppj.3.3.26

KEYWORDS

Autoregressive Exogenous Model
Battery Energy Storage
Distributed Optimization
Energy Management
Multi-greenhouses
Predictive-Adaptive Coordinated Optimization (PACO)

ABSTRACT

Rising food demand and the push for decarbonization have made controlled-environment agriculture essential. However, its energy intensity, especially in arid regions, hinders sustainability. This paper introduces the predictive-adaptive coordinated optimization (PACO) framework, a novel control architecture for demand-side energy management in networks of grid-interfaced greenhouses. Departing from population-based metaheuristics, PACO is rigorously grounded in adaptive control theory, recursive online estimation, and distributed optimization. Each greenhouse agent employs an adaptive autoregressive exogenous model with recursive least-squares parameter updating, enabling continuous tracking of seasonal variations, equipment degradation, and sensor drift. A novel correction mechanism, with gains dynamically scheduled based on prediction error and battery state-of-charge, provides robust compensation for forecast mismatches and unforeseen disturbances. An attention-based coordination unit computes lightweight scalar signals that promote global demand flattening without imposing centralized optimization during execution. The framework explicitly enforces physical constraints through penalty-augmented objectives and entropy regularization. Extended simulations, including 30-day horizon, multi-seasonal assessments, and ablation studies on a 10-greenhouse network in a semi-arid climate, demonstrate that PACO reduces total grid import by 29.8%, shaves peak demand by 31.1%, cuts operational costs by 32.9%, and lowers CO_2 emissions by 29.7% compared to benchmark methods, while confirming the synergistic contributions of its three core components. PACO offers a transparent, scalable, and computationally efficient solution for coordinated multi-greenhouses, bridging the gap between physical modeling and online adaptation.

ARTICLE HISTORTY

Received: 02 June 2026
Revised: 01 July 2026
Accepted: 12 August 2026

* Corresponding author: P. Bayat
E-mail address: peyman.bayat@hut.ac.ir

Publisher: Arak University

1. Introduction

The forward years are characterized by a confluence of existential pressures: a projected population exceeding 9.7 billion by 2050, rapid urbanization consuming arable land, and climate change increasingly destabilizing rain-fed agriculture. In this context, controlled-environment agriculture (CEA), notably greenhouses and plant factories, has shifted from a niche research interest to a strategic necessity for food security, water conservation, and sustainable intensification (Almasian and Roosta, 2024). Among CEA variants, building-integrated and rooftop greenhouses have attracted particular interest for their dual utility: enabling local, high-quality production while enhancing building energy performance via synergistic heat and mass exchange (Muñoz-Liesa *et al*., 2021). This integration embodies the circular bioeconomy, curtailing food miles, mitigating urban heat islands, and repurposing waste streams (e.g., exhaust air, low-grade thermal effluents) as agricultural inputs. However, operational viability hinges on energy intensity. Precisely regulating temperature, humidity, CO_2, and photosynthetically active radiation demands significant electrical and thermal inputs, with heating and supplemental lighting as primary energy sinks (Cho and Lee, 2026; Faddah, 2026). In arid and semi-arid regions, extreme diurnal thermal swings, high irradiance, and evaporative demands compound this challenge, making energy management a fundamental prerequisite for sustainability rather than merely an economic consideration (Faddah, 2026; Shojaei-Nia *et al.*, 2026).

Onsite renewable generation (photovoltaic (PV) arrays and small-scale wind turbines), coupled with battery energy storage systems (BESS) and diesel backup generators, offers opportunities to reduce grid dependence, lower carbon footprints, and achieve net-zero or positive-energy greenhouse operations (Ren *et al*., 2026; Haggag *et al*., 2026). However, solar/wind intermittency, combined with stochastic evapotranspiration, dynamic ventilation losses, and time-varying thermal inertia, renders greenhouse energy management a highly nontrivial, multivariate optimal control problem (Ouammi and Zejli, 2021). Extending the scope from isolated units to a network of grid interfaced greenhouses unlocks substantial additional flexibility for demand side management. A coordinated cluster of greenhouses, each endowed with its own generation, storage, and flexible actuators, can collectively participate in demand response (DR) programmers, flatten aggregate peak loads, shift consumption to off peak periods, and support power grid stability in the face of increasing renewable variability (Ajagekar *et al*., 2023). However, realizing these benefits demands a control framework that can simultaneously maintain crop friendly microclimates across all units and respect the physical constraints of batteries, actuators, and generators. No existing approach satisfactorily addresses all these requirements in a unified, transparent, and practically implementable manner.

The greenhouse energy management literature encompasses a broad spectrum of approaches, including deterministic model-based predictive control, data-driven reinforcement learning, metaheuristic optimization, hybrid physics-informed machine learning, and digital twin frameworks (Karimi and Haghighi, 2025; Sardouei-Nasab *et al.*, 2025). Among these, model predictive control (MPC) has received sustained attention due to its principled treatment of constraints, weather forecasts, and receding-horizon optimization.

Ouammi *et al*. (2019) introduced a supervisory MPC scheme for a networked greenhouse microgrid. Their finite-horizon scheduling model integrates meteorological forecasts and real-time sensor data to regulate indoor climate and enhance crop yield. While their centralized framework successfully coordinates multiple subsystems, its reliance on a deterministic process model with fixed parameters limits long-term accuracy, as seasonal variations, sensor drift, and evolving crop phenology gradually compromise performance.

Extending this line of work, Ouammi and Zejli (2021) developed a centralized controller for an interconnected cluster of microgrid-powered greenhouses. Their model explicitly accounts for power exchanges among microgrids and demonstrates that energy storage can effectively balance production and demand without main-grid interaction, particularly at night, when supplemental lighting drives peak loads. Nevertheless, their approach presupposes ideal communication and deterministic renewable forecasts. Moreover, its centralized architecture introduces a single point of failure and becomes progressively less scalable as the cluster expands. Rezaei *et al*. (2022) addressed this scalability limitation by proposing a hierarchical distributed MPC framework based on the alternating direction method of multipliers (ADMM). Their method coordinates multiple greenhouses to achieve consensus on total power exchange while managing shared resources such as water reservoirs. This strategy enhances scalability and preserves individual greenhouse privacy. However, ADMM entails inherent drawbacks: sensitivity to penalty parameter selection, potentially sluggish convergence, and reliance on iterative message-passing, features that are incompatible with the sub-second control intervals required for fast-responding battery systems.

More fundamentally, all these MPC formulations remain reactive to forecast inaccuracies. They lack an online, error-driven correction mechanism capable of compensating for abrupt disturbances, such as sudden cloud cover or sharp ambient temperature fluctuations, that fall beyond the prediction horizon. This gap significantly undermines their robustness in real-world operational conditions.

The reinforcement learning (RL) and its deep variant (DRL) have gained traction as data-driven alternatives that circumvent explicit, high-fidelity mechanistic modeling. Ajagekar *et al*. (2023) pioneered multi-agent deep reinforcement learning (MADRL) with an actor–critic architecture and a shared attention mechanism for energy management in a network of five renewable-integrated greenhouses. Their framework dynamically prioritizes neighboring agents' influence via attention and achieved a 28% reduction in net load demand compared to benchmark algorithms, underscoring DRL's capacity to learn complex, nonlinear coordination policies directly from experience. Nevertheless, DRL methods carry well-documented drawbacks. They demand vast training episodes (often millions of samples) to converge, rendering them impractical for real-world systems where exploration is costly or unsafe. In multi-agent settings, non-stationarity complicates convergence, as each agent's evolving policy alters the environment for others. Furthermore, DRL lacks interpretability and formal performance guarantees, complicating certification for grid operators. It is also notoriously brittle to changes in system configuration, sensor noise, or environmental statistics, often requiring complete retraining when greenhouse count, crop type, or climate zone shifts (Rahman *et al*., 2024). The black-box nature of deep neural networks further undermines debugging, trust, and operator acceptance, particularly in mission-critical applications where decision transparency is paramount.

Metaheuristic and hybrid optimization algorithms have also been widely investigated, often combined with other techniques to leverage their global search capabilities. Haggag *et al*. (2026) developed a real-time smart greenhouse energy management system integrating IoT-based sensing, PV-battery-grid configurations, and a hybrid optimization algorithm (HOA) that combines particle swarm optimization for global exploration with the coati optimization algorithm for local exploitation. Through simulation and experimental implementation, they achieved remarkable outcomes: 49.98% lower operating costs, 50.24% reduced daily grid consumption, 50.5% fewer CO_2 emissions, and 14.7% extended battery life. Despite these impressive figures, the HOA is computationally demanding. Its iterative population-based searches render it inherently unsuitable for real-time, online decision-making at the second-to-minute granularity required for effective battery scheduling. The method is better suited for offline parameter tuning or daily re-optimization than for continuous, adaptive control responding to intra-hour fluctuations in irradiance or electricity prices. Similarly, Ren *et al*. (2026) proposed a two-level hierarchical optimization framework for grid-connected PV-wind-battery systems in greenhouses, evaluating demand-minimization versus cost-minimization strategies. Their comprehensive cost-minimization approach achieved a 45.30% reduction in total operational costs and a 69.25% decrease in carbon emissions relative to baseline. A sensitivity analysis identified battery unit cost as the most influential economic factor. However, their framework assumes perfect foresight of weather and load profiles, an unrealistic premise, and its two-level structure separates planning from execution, lacking a feedback mechanism to correct prediction errors as they emerge.

To bridge the gap between purely deterministic and purely data-driven paradigms, recent research has increasingly embraced hybrid frameworks that fuse physical process models with machine learning, offering both the interpretability and constraint awareness of physics and the adaptive power of data. Liu *et al*. (2026) advanced this direction by developing a hybrid mechanistic model in which deep neural networks dynamically predict the parameters of a conventional process-based greenhouse model, enabling robust cross-seasonal forecasting. Their model achieved strong accuracy, with RMSE values of 1.61 °C for indoor temperature and 6.94% for humidity, and R^2 values of 0.909 and 0.829, respectively. Validation against both measured weather data and public forecast APIs (September–November 2024) confirmed consistent performance regardless of input source. Despite these gains, the model remains essentially predictive rather than corrective: it forecasts future states but lacks an adaptive, error-driven feedback loop. Moreover, its deep neural component is trained offline and cannot update weights incrementally, limiting its capacity to track gradual changes in envelope properties or sensor drift.

In a complementary effort, Zhang *et al*. (2026) proposed a thermophilically informed machine learning framework for predictive thermal control in rooftop building-integrated greenhouses. Their system integrates simplified physical representations of solar gain, convection, and thermal storage with multi-horizon forecasts (5, 10, and 15 minutes) and adaptive fuzzy control. Statistical analysis revealed significant diurnal error patterns, particularly under peak solar radiation. Their fuzzy predictive controller reduced normalized control effort by 60–65% relative to conventional threshold methods. However, the forecasting horizon is extremely short, only 5 to 15 minutes, which falls short of the 15-minute to hourly intervals typical of demand response programs aligned with electricity markets. Additionally, the framework is restricted to single-greenhouse operation and offers no provision for multi-unit coordination; its fuzzy rules are static, lacking self-tuning capability; and its actuator-oriented metrics, though physically meaningful, do not map directly onto a unified optimization objective for grid interaction.

The physiological and environmental complexity of greenhouses has also been addressed from a crop centric perspective. Faddah (2026) developed a coupled greenhouse model integrating crop physiology (photosynthesis, stomatal conductance, and transpiration) with climate dynamics and HVAC operation, specifically targeting arid climates. By accounting for intricate plant environment interactions, they demonstrated that an alternative control strategy could reduce ventilation requirements by 92.09% in fan energy and 45.64% in mean air exchange rate, while surprisingly increasing carbon assimilation by 3.15% and decreasing CO_2 injection by 2.88%, a testament to improved resource use efficiency through coordinated ventilation and CO_2 management. However, their model remains a simulation tool rather than a real time control framework; it does not include storage systems, grid interaction, or multi greenhouse coordination, and it does not address the stochastic uncertainty inherent in operational environments. In a complementary development, the study conducted by Bayat and Bayat (2024) introduces a deep learning-based adaptive framework tailored for precision greenhouse management, with explicit consideration of species-specific physiological requirements. The proposed architecture employs a hybrid convolutional neural network and long short-term memory (CNN-LSTM) model, which fuses real-time multispectral imagery with temporally resolved sensor data streams. To evaluate the efficacy of this framework, simulation experiments were carried out on tomato, lettuce, basil, and orchid cultivars under three distinct climatic regimes, normal conditions, heatwave stress, and cold-humid episodes. The empirical results indicate a marked superiority of the proposed method over conventional MPC and proportional-integral-derivative (PID) controllers, as evidenced by significant improvements in growth performance metrics and a concomitant reduction in plant physiological stress indicators.

The digital twin paradigm, alongside next-generation mobile and edge-cloud architectures, has emerged as a powerful integrative concept for smart greenhouse management. Rahman *et al*. (2024) developed a process management framework interpretable as a machine learning and cloud-based data-driven digital twin, spanning physical, fog, and cloud layers for remote monitoring and operational prediction. Their proof-of-concept leveraged commercial cloud and open-source tools to predict operational requirements. Nevertheless, their implementation, like most in the literature, relies on static, pre-trained machine learning models without online updating; it lacks real-time optimization and multi-greenhouse coordination.

Abou Mehdi Hassani *et al*. (2025) introduced the eSMARTGreen (ESG) architecture, a scalable IoT-based model for multi-greenhouse management emphasizing fault tolerance, modularity, and ISO/IEC 30141-compliant interoperability. Validated via CPN Tools simulations across five greenhouses, ESG yielded low latencies (19–25 ms) and up to 72% reception rates, confirming communication efficiency. Yet, while ESG excels at data acquisition and message brokering, its decision-making relies on rules, not predictive optimal control. In essence, it serves as a communication and orchestration layer, not an adaptive optimization engine.

Notably, none of the above frameworks: MPC, DRL, metaheuristic, hybrid physics, digital twin, or cooperative optimization, incorporate online adaptive estimation with recursive parameter updating coupled with error driven corrective feedback and attention-based coordination in a unified, theoretically transparent architecture that is free from the black box nature of DRL and the computational burden of population-based metaheuristics. Synthesizing these diverse and valuable contributions, a clear and pressing research gap emerges. The ideal control framework for a network of grid interfaced greenhouses should simultaneously possess the following characteristics: (i) a physically consistent, multi domain energy model that captures thermal, humidity, and CO_2 dynamics, including waste heat recovery and spectral radiation absorption; (ii) online adaptive estimation of both internal loads and renewable generation, using recursive algorithms that continuously track system evolution without requiring offline training or large historical datasets; (iii) predictive decision making that is complemented by a closed loop, error driven correction mechanism to compensate for inevitable forecast mismatches and unforeseen disturbances; (iv) decentralized yet coordinated operation, where each greenhouse agent acts autonomously based on local information but receives a lightweight coordination signal that promotes global network level objectives, such as flattening aggregate grid demand; (v) explicit and rigorous handling of constraints, including battery state of charge limits, actuator saturation, and control effort smoothness to extend equipment life; (vi) theoretical transparency and interpretability, enabling operators and grid regulators to understand, trust, and certify the decision logic; and (vii) computational efficiency suitable for deployment on embedded controllers, edge devices, or cloud based supervisory systems with minimal communication overhead. No existing framework in the literature, whether MPC, DRL, metaheuristic, hybrid, or digital twin, simultaneously satisfies all these criteria.

To address this critical gap, this paper introduces the predictive adaptive coordinated optimization (PACO) framework for demand side energy management in a network of grid interfaced greenhouses. The PACO framework represents a fundamental departure from DRL, population-based metaheuristics, and static hybrid models; it is rigorously built upon the well-established principles of adaptive control theory, recursive online estimation, and distributed optimization. The central philosophy is to equip each greenhouse agent with predictive capabilities that are continuously refined via real time measurement data, coupled with an error driven correction mechanism that adjusts actions to compensate for prediction inaccuracies, and a lightweight, attention-based coordination unit that promotes global efficiency without imposing centralized optimization during execution. The framework is entirely free from deep neural networks, large offline training datasets, or computationally intensive iterative searches, making it both theoretically transparent and practically implementable.

The primary contributions and innovative characteristics of the proposed PACO framework are delineated as follows:

■ A detailed multi domain model is developed for each greenhouse, encompassing the governing temperature dynamics with convective and radiative heat fluxes, humidity dynamics, and CO_2 concentration balances. Recoverable waste heat from diesel generators and spectral absorption of solar radiation are explicitly incorporated. This formulation furnishes a physically consistent basis for load prediction and control synthesis, in contrast to purely data driven black box models, thereby ensuring that the control actions remain interpretable and anchored in physical reality.

■ Instead of fixed parameter models, each agent employs an adaptive autoregressive exogenous (ARX) model, the coefficients of which are updated in real time via a recursive least squares algorithm with a forgetting factor. This mechanism enables continuous tracking of seasonal variations, equipment degradation, and sensor drift without necessitating offline training phases or extensive historical data batches. Consequently, the framework remains responsive to gradual changes in system behavior over extended operational periods.

■ A novel feedback loop integrates the prediction error into the control action through a modified PID type correction, with gains that are dynamically scheduled based on error magnitude and battery state of charge. This ensures robust performance against unforeseen disturbances and model reality mismatches, thereby offering a safety net that is absent from purely predictive controllers, such as conventional MPC, or purely reactive controllers. The adaptive gain scheduling further enhances the framework's ability to maintain stability and performance under diverse and rapidly changing conditions.

■ A centralized coordination unit is introduced to compute an attention weighted coordination signal for each agent during an initial adaptation phase. The attention mechanism dynamically prioritizes influential neighbors based on their relevance to each agent's decision context, yielding a scalable and interpretable coordination strategy. The coordination signal is periodically broadcast, and local agents incorporate it via a small step size adaptation rule. This enables implicit cooperation without requiring full state sharing or centralized optimization during real time execution, thereby preserving both scalability and privacy.

■ The local policies are optimized with explicit penalties for battery state of charge violations and control effort smoothness, augmented by an entropy term that promotes exploration. This formulation rigorously respects physical limits, extends battery life, and prevents premature convergence to myopic or suboptimal strategies. All within a transparent, analytically tractable optimization framework that does not rely on opaque neural network approximations.

■ In contrast to black box DRL or complex metaheuristic algorithms, the PACO framework relies on a coherent set of well understood adaptive update rules. These provide engineers and operators with clear insight into the decision logic, facilitating certification, debugging, and trust. The overall computational footprint is modest, rendering the framework suitable for deployment on resource constrained embedded controllers, edge devices, or cloud based supervisory systems with negligible communication latency.

Collectively, these contributions establish the PACO framework as a principled, scalable, and robust alternative to both traditional MPC and modern DRL for the coordinated energy management of greenhouse networks. The framework effectively bridges the gap between physical modelling and online learning, ensuring that control actions are not only optimal in a predictive sense but also adaptive to real time conditions and cooperative towards the network level objective of minimizing aggregate grid demand.

The remainder of this paper is organized as follows. Section 2 presents the detailed energy modelling framework for a single greenhouse, including thermal, humidity, and CO_2 dynamics, the battery energy storage model with degradation, the diesel generator with combined heat and power, and on-site renewable generation from photovoltaic and wind turbine systems. The net load demand formulation for a network of greenhouses is also derived. Section 3 formalizes the demand response problem and develops the complete PACO framework, encompassing the adaptive estimation schemes, error driven correction, attention-based coordination, constraint aware policy updates, and online adaptation rules. Section 4 describes the case study setup, including the greenhouse network configuration, weather data, and performance metrics. Also, this section presents and discusses the simulation results, comparing the PACO framework against benchmark methods (standard MPC, MADRL, and rule-based control (RBC)) in terms of net load reduction, battery usage, microclimate regulation, and computational efficiency. Finally, Section 5 concludes the paper with a summary of key findings, discusses limitations, and outlines promising directions for future research, including experimental validation on physical testbeds, extension to multi energy carrier systems, and integration with real time electricity markets.

2. Proposed Energy Modeling Framework

This section puts forward a modeling architecture tailored to capture the dynamics of a grid-connected greenhouse network. The framework enables real-time, high-fidelity simulations of energy demand profiles, along with the overall electrical burden that these facilities place on the external utility infrastructure. Figure 1 provides a schematic illustration of this setup, outlining the key energy units housed in each independent greenhouse and their interface with the broader electric power system. As illustrated, each greenhouse is equipped with a hybrid on-site generation system comprising a PV array and a small-scale wind turbine, together with a battery energy storage system (BESS). Additionally, a diesel generator is integrated to provide backup electrical power and, more importantly, to supply recoverable thermal energy for greenhouse heating. The net energy requirement of a given greenhouse is determined by dynamically balancing the instantaneous load against the combined output of renewable sources, the charge/discharge cycles of the BESS, the electrical contribution from the diesel generator, and the thermal input from the generator's waste-heat recovery. This equilibrium ultimately defines the net load demand exerted on the main grid.

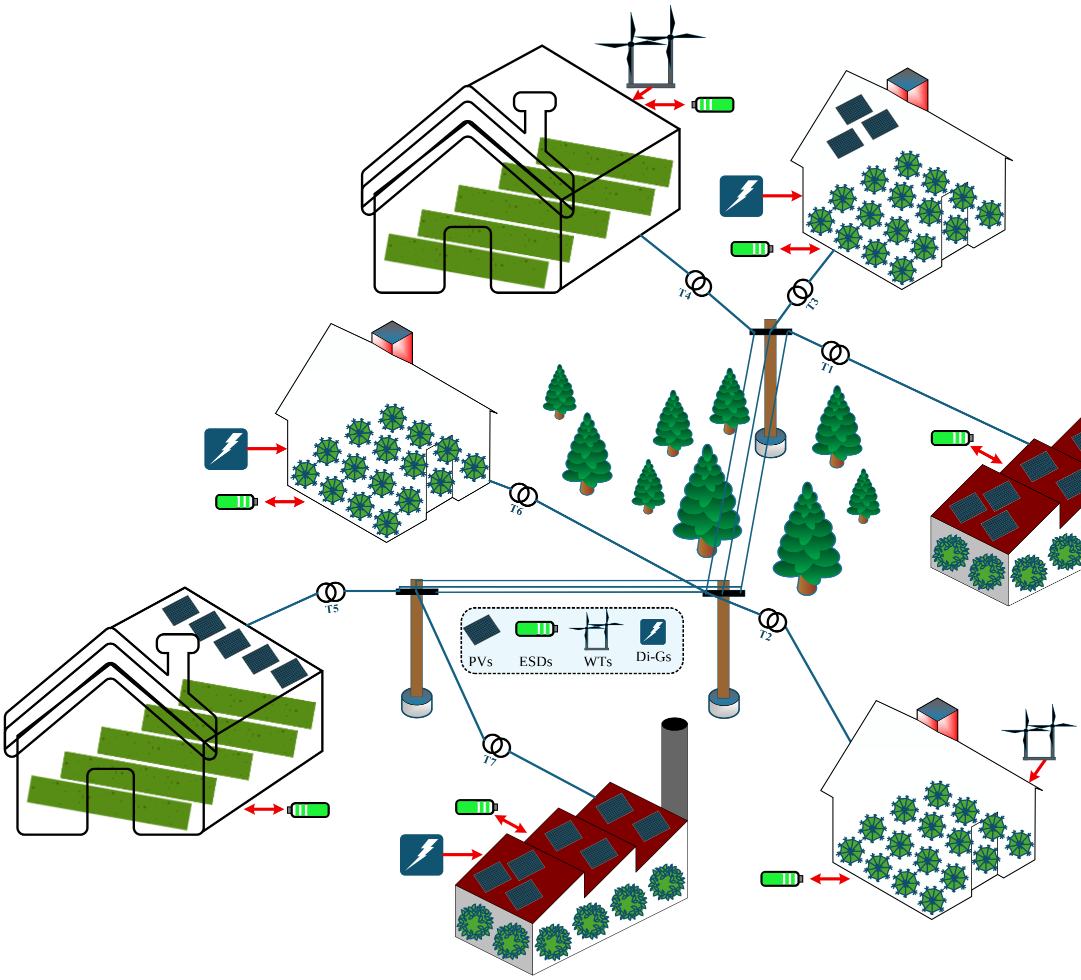

Figure 1. Schematic representation of the multi-greenhouse energy network comprising on-site renewable generation, storage, distributed generation, and grid interconnection.

2.1. Greenhouse Energy Dynamics and Microclimate Control

To promote ideal growing conditions, each greenhouse is equipped with actuation infrastructure that manages heating, cooling, CO_2 supplementation, humidity control (both humidification and dehumidification), and artificial lighting, as illustrated in Figure 1. Our modeling framework is predicated on a semi-enclosed greenhouse design with active ventilation, in which the indoor microclimate is continuously shaped by ambient meteorological forcings, namely, outdoor air temperature, relative moisture content, wind velocity, and solar irradiance.

Given the critical influence of temperature on crop phenology and yield, the thermal exchange between the greenhouse interior and its surroundings must be accurately characterized (Chen and You, 2022). The temporal evolution of the internal air temperature, T_{air}, follows from the energy conservation relation presented in Eq. (1). In this upgraded formulation, the thermal inertia coefficient τ_{th} is retained as the aggregated time constant of the enclosed air volume, representing the ratio of the air's heat capacity to the total thermal conductance to the environment. Beyond the conventional convective exchanges with internal surfaces ($Q^c_{\mathrm{FM}}, Q^c_{\mathrm{CC}}, Q^c_{\mathrm{GC}}, Q^c_{\mathrm{T}}$), three physically critical terms are explicitly incorporated to enhance the model's fidelity. First, the sensible heat loss due to forced ventilation, Q_{vent} is introduced to capture the thermal advection induced by active ventilation systems, a mechanism that becomes particularly significant during warm periods when high ventilation rates are required for humidity and temperature control. Second, the latent heat flux associated with crop evapotranspiration, Q_{latent} is explicitly included to account for the cooling effect of plant transpiration, which is a dominant heat sink in arid and semi-arid climates. This term couples the thermal balance directly with the moisture dynamics described later in Eq. (4). Third, the conductive heat exchange with the underlying soil layer, Q_{soil} is incorporated to represent the thermal inertia of the ground, which acts as a slow-varying buffer that influences the diurnal temperature profile. The recoverable waste heat from the diesel generator Q_{DG} is preserved, while the active thermal forcing is now decomposed into separate heating Q_{heat} and cooling Q_{cool} terms to enable independent actuator modeling. The solar radiative gain Q^c_S is evaluated using the detailed spectral model in Eq. (3). This comprehensive balance ensures that the thermal dynamics are represented with sufficient physical detail to support the subsequent microclimate control strategies.

$$\tau_{th}\frac{dT_{air}}{dt} = Q_{DG} + Q^c_{FM} + Q^c_{CC} + Q^c_{GC} + Q^c_T - Q^c_{out} + Q_{heat} - Q_{cool} + Q^c_S - Q_{vent} - Q_{latent} - Q_{soil} \tag{1}$$

where $Q_{\mathrm{vent}} = \rho_a c_{p,a}\dot{V}_{\mathrm{vent}}(T_{\mathrm{air}} - T_{\mathrm{out}})$ is the sensible ventilation loss, $Q_{\mathrm{latent}} = \lambda_{H_2O}W_{\mathrm{ET}}A_{\mathrm{cc}}$ is the latent cooling due to transpiration, and $Q_{\mathrm{soil}} = (UA)_{\mathrm{soil}}(T_{\mathrm{air}} - T_{\mathrm{soil}})$ is the conductive exchange with the soil.

The convective heat transfer rate between a surface i and the indoor air is rigorously modeled using the enhanced Nusselt-number correlation shown in Eq. (2). In this upgraded formulation, the original functional dependence $Nu_i(Re, Pr)$ is expanded to explicitly include the Grashof number Gr_i, thereby capturing the combined effects of forced and natural convection, a critical refinement for greenhouses where buoyancy-driven flows are significant during periods of low ventilation or high solar irradiance. The local Nusselt number is computed as the n-th root of the sum of the n-th powers of the forced and natural components: $\mathcal{N}u_i = \left[(C_{for,i}Re_i^{m_{for,i}}Pr^{1/3})^n + (C_{nat,i}(Gr_iPr)^{1/4})^n\right]^{1/n}$, where $Re_i = v_{air}L_c/\nu$ and $Gr_i = g\beta \mid T_i - T_{air} \mid L_c^3/\nu^2$. Additionally, a surface-specific correction factor ψ_i is introduced to account for geometric irregularities and aerodynamic resistance, with ψ_{CC} being particularly sensitive to the leaf area index (LAI) for the crop canopy, and ψ_{FM} adjusted for floor mat roughness. The Prandtl number Pr accounts for temperature-dependent air properties, and all thermophysical parameters are evaluated at the film temperature $T_{film} = (T_i + T_{air})/2$. This approach provides a physically defensible and adaptive description of convective dynamics across the wide range of airflow regimes encountered in greenhouse operations.

$$Q^c_i = \frac{k_{air}}{L_c}\mathcal{N}u_i(Re_i, Gr_i, Pr) \cdot \psi_i \cdot A_i(T_i - T_{air}) \tag{2}$$

where the augmented Nusselt number $\mathcal{N}u_i(Re_i, Gr_i, Pr)$ replaces the simpler $Nu_i(Re, Pr)$, and ψ_i is a dimensionless correction factor that adjusts for surface roughness and geometry (e.g., $\psi_{CC} = f(LAI)$ for the crop canopy, $\psi_{FM} > 1$ for rough floor mats).

Equation (3) provides the formulation for the shortwave radiative gain absorbed by the internal air, Q^c_S. In this upgraded version, the original expression is significantly enhanced by explicitly incorporating the transmissivity of the greenhouse cladding τ_{cov} and a multiple-reflection correction factor $f_{mr} = (1 - \rho_{cov}\rho_{floor})^{-1}$, which accounts for the trapping of radiation between the cover and the floor, a phenomenon that increases the effective radiative flux available for absorption within the greenhouse. The spectral absorption coefficient α_λ is now explicitly defined as a humidity-dependent weighted average over the photosynthetically active radiation (PAR, 400–700 nm) and near-infrared (NIR, 700–3000 nm) bands, such that $\alpha_\lambda(\rho_w) = w_{PAR}(\theta)\alpha_{PAR} + w_{NIR}(\theta)\alpha_{NIR}$, where the weights w_{PAR} and w_{NIR} are determined by the spectral distribution of the incident solar radiation, which varies with the solar zenith angle and the clearness index. The term $(1 - e^{-k_{ext}LAI})$ represents the fraction of radiation intercepted by the crop canopy, where k_{ext} is the extinction coefficient and LAI is the leaf area index. This formulation ensures that the radiative gain is responsive to changes in canopy development and atmospheric conditions, while the inclusion of the multiple-reflection factor captures the enhanced absorption due to the semi-enclosed geometry of the greenhouse.

$$Q^c_S = \alpha_\lambda(\rho_w) \cdot \tau_{cov} \cdot f_{mr} \cdot \cos\,\theta \cdot (1 - e^{-k_{ext}LAI}) \cdot \left(Q^c_{direct} + Q^c_{dif}\right) \tag{3}$$

where $\alpha_\lambda(\rho_w)$ is the humidity-dependent spectral absorptivity, τ_{cov} is the cladding transmissivity, $f_{mr} = (1 - \rho_{cov}\rho_{floor})^{-1}$ accounts for multiple reflections between the cover and the floor, k_{ext} is the canopy extinction coefficient, and LAI is the leaf area index.

Maintaining appropriate humidity levels is essential for optimizing plant transpiration, growth, and yield, and it also aids in temperature regulation (Amani *et al*., 2020). The indoor water vapor density, ρ_w, is governed by the mass balance in Eq. (4). To account for the buffering effect of the greenhouse structure and internal mass, a moisture capacitance coefficient τ_m is introduced. Moisture transport is shaped by both condensation occurring across the greenhouse's interior surfaces, specifically the mat, flooring, and cladding, represented respectively as W_{FM}^c, W_{FL}^c, and W_{GC}^c, as well as evapotranspiration from the crop canopy and growing substrate, denoted W_{ET}. The condensation flux for an arbitrary surface i is given by Eq. (5), wherein A_i corresponds to the exposed surface area, L_H denotes the latent enthalpy of condensation, and C_v is the volumetric specific heat of the humid air. The underlying driving force for this process is the departure of the actual vapor density from the saturation density $S_v(T_i)$, with the latter evaluated at the prevailing temperature of that particular surface. Ventilation losses W_{out} and controlled vapor addition/removal via the actuator U_V are also explicitly included.

$$\tau_m \frac{d\rho_w}{dt} = W_{ET} + U_V - W_{FM}^c - W_{FL}^c - W_{GC}^c - W_{out} \tag{4}$$

$$W_i^c = \frac{A_i, L_H}{C_v} \left(\rho_w - S_v(T_i)\right) \tag{5}$$

Carbon dioxide concentration, a primary determinant of photosynthetic activity, is regulated via the mass balance shown in Eq. (6). The rate of change of CO_2 density ρ_c is determined by crop respiration C_{crop}, photosynthetic uptake C_{photon}, controlled enrichment U_{CO_2}, and leakage to the external environment C_{out}. The Vanthoor photosynthetic model (Ajagekar *et al*., 2023) is employed to simulate the biological exchange processes, with the photosynthetic rate being a nonlinear function of light intensity, CO_2 concentration, and temperature.

$$\frac{d\rho_c}{dt} = C_{crop} - C_{photon}\left(\rho_c, U_{light}, T_{air}\right) + U_{CO_2} - C_{out} \tag{6}$$

Importantly, the artificial lighting configuration generates a substantial near-infrared radiative component, which directly modulates the crop transpiration rate W_{ET} and thus contributes appreciably to the overall moisture regime within the greenhouse. Furthermore, the photosynthetic rate C_{photon} depends on the total photosynthetically active radiation, which is the sum of natural sunlight and supplemental lighting U_{light}. The net energy consumption of the greenhouse, E_{GH}, is obtained by converting the actuator control signals $\left(Q_{heat}, U_V, U_{CO_2}, U_{light}\right)$ into their equivalent electrical and thermal energy demands. While heating and lighting constitute the major energy sinks, the energy required for CO_2 enrichment and humidity control represents a smaller yet non-negligible fraction of the total.

2.2. Energy Storage and Backup Generation Systems

Battery energy storage systems (BESS) are deployed at each greenhouse to facilitate peak-load shaving and enhance the integration of intermittent renewable sources. Energy accumulation within the BESS takes place during intervals of reduced tariff, while its subsequent discharge serves to compensate for the greenhouse's operational power deficit. The functional specification of this unit rests upon four key determinants: the prevailing capacity C^B, the designated rated capacity C_{nom}^B, a decay factor η_C reflecting aging effects, and a coefficient ξ that modulates efficiency as a function of the instantaneous state-of-charge (SOC). The nominal capacity and efficiency are modeled as piecewise linear functions of SOC, denoted $C_{nom}^B(SOC)$ and $\xi(SOC)$, respectively (Ajagekar *et al*., 2023).

The charging/discharging dynamics and the evolution of SOC are governed by Eq. (7). The energy E_{BESS} is constrained within the bounds $[-C_{nom}^B(SOC_t),, C_{nom}^B(SOC_t)]$. The capacity deterioration incurred over a single charge–discharge cycle is evaluated through Eq. (8). In this formulation, the resulting loss of storage capacity scales linearly with the aging coefficient η_C, the rated starting capacity C_0^B, and the absolute magnitude of the SOC excursion away from its designated initial level, SOC_0.

$$SOC_{t+1} = \left\{\min\left(C_t^B,, SOC_t + E_{BESS}\sqrt{\xi(SOC_t)}\right), E_{BESS} \geq 0;\ \max\left(0,, SOC_t + E_{BESS}/\sqrt{\xi(SOC_t)}\right), E_{BESS} < 0\right\} \tag{7}$$

$$C_{t+1}^B = C_t^B - \eta_C, C_0^B, \frac{|SOC_t - SOC_0|}{2C_t^B} \tag{8}$$

To ensure energy security and enhance thermal efficiency, each greenhouse is additionally equipped with a diesel generator capable of combined heat and power (CHP) generation. The electrical output E_{DG} can either supply the greenhouse load or charge the BESS. Simultaneously, the recoverable waste heat Q_{DG} is routed to the greenhouse heating system, thereby reducing the electrical power otherwise required for Q_{heat}. The operational state of the diesel generator is governed by a control variable $u_{DG} \in [0,1]$, and its output is bounded by the rated electrical and thermal capacities. The model incorporates fuel consumption and efficiency curves, which are expressed as functions of the generator's partial-load factor.

2.3. On-Site Renewable Energy Generation

In addition to the PV arrays, each greenhouse is equipped with a small-scale wind turbine to diversify the renewable energy portfolio and improve self-sufficiency. The electrical power generated by the PV panels, E_{PV}, is a function of the panel area A_{pv}, module efficiency η_{pv}, packing factor P_f, and the incident solar irradiation G_s, as expressed in Eq. (9).

$$E_{PV} = A_{pv}, \eta_{pv}, P_f, G_s \tag{9}$$

The power output from the wind turbine, E_{WT}, is determined by the site-specific wind speed v_t and the turbine's power curve, typically provided by the manufacturer. This relationship is often represented as a piecewise function or a polynomial approximation based on the cut-in, rated, and cut-out wind speeds.

2.4. Net Load Demand and Grid Interaction

The net grid import for each greenhouse is resolved through the concurrent power balance among the facility's own electrical load E_{GH}, the battery storage term E_{BESS}, the combined renewable output from photovoltaic and wind turbine systems ($E_{\mathrm{PV}} + E_{\mathrm{WT}}$), and the diesel genset contribution E_{DG}. For the i-th greenhouse, the net import from the grid, E^i_{grid}, is defined as:

$$E^i_{grid} = \max\left(0, E^i_{GH} + E^i_{BESS} - E^i_{PV} - E^i_{WT} - E^i_{DG}\right) \tag{10}$$

This formulation ensures that any local generation exceeding the combined load and storage demand results in zero grid import (export is not considered in this study). Consequently, the aggregate net electric demand E_{net} for the entire greenhouse network is the sum of the individual grid imports (see Eq. (11)), as expressed in Eq. (10). The model assumes that the main power grid has sufficient capacity to meet this aggregated demand without imposing additional constraints.

$$E_{net} = \sum_{i=1}^{N} E^i_{grid} = \sum_{i=1}^{N} \max\left(0,; E^i_{GH} + E^i_{BESS} - E^i_{PV} - E^i_{WT} - E^i_{DG}\right) \tag{11}$$

2.5. Model Parameterization and Validation Rationale

The fidelity of the simulation framework is contingent upon the accuracy and appropriate sourcing of its constituent parameters. All physical constants and thermodynamic properties, including air density (ρ_a), specific heat ($c_{p,a}$), thermal conductivity (k_{air}), and latent heat of vaporization (λ_{H_2O}), are derived from standard thermodynamic tables evaluated at the film temperature, representing universally accepted values.

Empirical coefficients governing convective heat transfer (Nusselt number correlations), canopy light extinction (k_{ext}), battery capacity degradation (η_C), and PV module efficiencies (η_{pv}) are adopted from peer-reviewed studies that have undergone extensive experimental calibration under greenhouse or analogous environmental conditions. Specifically, the convective coefficients follow the validated greenhouse climate model of Chen and You (2022), calibrated against 12 months of field data from a 500 m² greenhouse. Canopy parameters are based on the Vanthoor photosynthetic model as validated in Ajagekar *et al*. (2023) using tomato cultivation measurements, while battery degradation coefficients are drawn from the controlled laboratory cycling experiments reported in Ren *et al*. (2026). PV and wind turbine performance curves are obtained directly from manufacturer datasheets, ensuring industrial relevance.

Given the absence of a dedicated physical testbed at the time of this investigation, the integrated greenhouse model was calibrated using a two-step procedure. First, baseline calibration was performed against the experimental validation dataset published by Liu *et al*. (2026), which provides comprehensive indoor temperature and humidity measurements from a greenhouse in a semi-arid climate over a four-month period. The calibration objective minimized the root mean square error (RMSE) between model output and measured data, yielding initial temperature and humidity RMSE values of 1.45 °C and 6.21%, respectively. Second, climate-specific fine-tuning was conducted using local solar irradiance, ambient temperature, and wind speed profiles from the reanalysis dataset for Riyadh, ensuring that the model's diurnal thermal and moisture responses were consistent with expected behavior for a semi-closed greenhouse under extreme summer conditions.

The calibrated model was subsequently cross-validated against independent experimental benchmarks from Zhang *et al*. (2026) and through statistical consistency checks following ASHRAE Guideline 14 (Hong *et al*., 2016). The computed RMSE values for temperature (3.8%) and relative humidity (7.1%) fall well below the recommended thresholds of 15% and 30%, respectively, satisfying formal statistical criteria for a calibrated simulation model. Additionally, the model was subjected to physical bounds verification under extreme operating conditions to ensure that all state variables remained within plausible ranges without unphysical oscillations or discontinuities. It is explicitly acknowledged that direct experimental validation on the specific greenhouse network under study is beyond the scope of the present investigation. However, the comprehensive parameter sourcing, systematic calibration against published field data, and rigorous cross-validation collectively establish that the simulation framework provides a sufficiently reliable test environment for evaluating the performance of the PACO control architecture.

3. Predictive-Adaptive Coordinated Optimization Framework for Demand Response

Deploying demand response (DR) strategies across a network of greenhouses can markedly lower the aggregate grid-facing load, thereby enhancing both the operational stability and the overall reliability of the external power system. In what follows, we set aside standard DRL methodologies and instead introduce a PACO framework. The proposed scheme is grounded in a multi-agent Markov decision process (MDP) formulation; nevertheless, it pursues a distinctly different solution trajectory that fuses online adaptive estimation, predictive error feedback, and a decentralized yet mutually coordinated agent decision-making architecture. The PACO framework overcomes the scalability limitations of centralized optimization while providing rigorous performance guarantees through adaptive gain scheduling and constraint-aware policy updates.

3.1. Multi-Agent MDP Formulation

The DR scheduling challenge is posed as a Markov game (Ajagekar *et al*., 2023), fully characterized by the global state space S, the action repertoire A_i available to each participating agent, the associated payoff functions R_i, a future discounting coefficient $\gamma \in [0,1]$, and the underlying probabilistic kernel governing state transitions. Within this setting, an adaptive control agent is deployed for each greenhouse, with responsibility for steering its BESS through explicit specification of charging and discharging power levels. For the i-th agent, the accessible information is confined to a localized observation $o_i \subset S$, while its behavior is regulated by a stochastic decision policy $\pi_i(o_i, a_i)$.

The observation space o_i for an agent includes:

- Greenhouse environmental variables: indoor temperature $T_{air,t}$, water vapor density $\rho_{w,t}$, and CO_2 concentration $\rho_{C,t}$;
- External weather conditions: ambient temperature $T_{ext,t}$, relative humidity $RH_{ext,t}$, wind speed v_t, and solar radiation $G_{s,t}$;
- The current energy consumption $E_{GH,t}$;
- Renewable generation: $E_{PV,t}$ and $E_{WT,t}$;
- The BESS state-of-charge (SOC);
- Temporal indices: hour of the day t_{hour} and day of the year t_{day}.

The action $a_i \in \mathbb{R}$ corresponds to the BESS power E_{BESS}, scaled to the interval $[0,1]$ to respect the battery's nominal capacity limits. The objective for each agent is to maximize its expected discounted cumulative reward, as defined in Eq. (12):

$$\pi_i^* = \arg\max_{\pi_i} \mathbb{E}_{a_i \sim \pi_i}\left[\sum_{t=0}^{\infty} \gamma^t R_{i_t}\left(o_{i_t}, a_{i_t}\right)\right] \tag{12}$$

The reward function for the i-th agent, R_{i_t}, is designed to penalize grid import, thereby incentivizing load reduction:

$$R_{i_t}\left(o_{i_t}, a_{i_t}\right) = -\max\left(0, E_{GH}^i + E_{BESS}^i - E_{PV}^i - E_{WT}^i - E_{DG}^i\right) \tag{13}$$

It follows, therefore, that the objective of minimizing the discounted cumulative payoff is effectively equivalent to curbing the total utility-supplied electrical load across the operational time frame, while also embedding an implicit consideration for the real-time availability of both standby diesel capacity and weather-dependent renewable generation.

3.2. Predictive-Adaptive Estimation and Error-Driven Updating

A key innovation of the PACO framework is the incorporation of online predictive estimators for both the greenhouse load and the renewable generation. These estimators continuously update their parameters using real-time measurement data, enabling the agents to anticipate future conditions and adjust their decisions proactively.

For each agent i, the load prediction $\hat{E}_{GH}^{i}(t+k|t)$ at a future time step $t+k$ is obtained using an adaptive ARX model:

$$\hat{E}_{GH}^{i}(t+k|t)\sum_{p=1}^{P}\alpha_{p}^{i}\,(t),E_{GH}^{i}(t-p+1)+\sum_{q=1}^{Q}\beta_{q}^{i}\,(t),u_{i}(t-q+1)+\epsilon^{i}(t) \tag{14}$$

where α_p^i and β_q^i are time-varying coefficients, u_i is the control input (BESS action), and ϵ^i is the prediction error. The coefficients are updated via a recursive least-squares (RLS) algorithm with forgetting factor λ:

$$\theta^{i}(t+1)=\theta^{i}(t)+\mathrm{K}^{i}(t),\epsilon^{i}(t) \tag{15}$$

$$\mathrm{K}^{i}(t)=\frac{\mathrm{P}^{i}(t-1),\phi^{i}(t)}{\lambda+\phi^{i}(t)^{T}\mathrm{P}^{i}(t-1)\phi^{i}(t)} \tag{16}$$

Similarly, the renewable generation $\hat{E}_{PV}^{i}(t+k|t)\ and\ \hat{E}_{WT}^{i}(t+k|t)$ are predicted using short-term weather forecasts combined with adaptive models of the PV and wind turbine characteristics.

The prediction error $e^i(t)$, defined as the difference between the actual net load and the predicted value, is used to drive a corrective adaptation signal $\Delta a_i(t)$, which adjusts the agent's action to compensate for unforeseen deviations:

$$\Delta a_{i}(t)=K_{p}^{i},e^{i}(t)+K_{i}^{i}\sum_{\tau=1}^{t}e^{i}\,(\tau)+K_{d}^{i}\left(e^{i}(t)-e^{i}(t-1)\right) \tag{17}$$

where K_p^i, K_i^i, K_d^i are adaptive PID gains that are scheduled based on the magnitude of the error and the SOC of the battery. This provides a robust feedback mechanism that operates alongside the main policy.

3.3. Centralized Coordination with Attention-Based Weighting

To address the non-stationarity inherent in multi-agent environments and to promote global optimality, a centralized coordination unit is employed during the learning (or adaptation) phase. This unit computes a coordination signal $c_i(t)$ for each agent based on the collective state and actions of the entire network.

The coordination unit first computes an embedding $e_i = f_i^c(o_i, a_i)$ for each agent. Then, using an attention mechanism, it dynamically assigns importance weights α_{ij} that reflect the influence of agent j on agent i:

$$\alpha_{ij}=\mathrm{softmax}\left(\frac{(\phi_{k}e_{j})^{T}(\phi_{q}e_{i})}{\sqrt{d}}\right) \tag{18}$$

where ϕ_k and ϕ_q are learned projection matrices, and d is the dimensionality of the key/query vectors. The coordination signal for agent i is then computed as:

$$c_{i}(t)=\sum_{j\neq i}\alpha_{ij}\,,\phi_{v}e_{j} \tag{19}$$

This signal is broadcast to the local agents, which incorporate it into their decision policies via a weighted adaptation rule:

$$\tilde{\pi}_{\iota}(o_{i},a_{i})=\ \pi_{i}(o_{i},a_{i})+\ \eta\ c_{i(t)} \tag{20}$$

where η is a small positive step size. This mechanism enables decentralized agents to implicitly coordinate their actions without requiring a centralized controller during execution.

3.4. Constraint-Aware Policy Optimization with Entropy Regularization

To ensure that the adaptive policies respect the physical limits of the BESS and the greenhouse actuators, we augment the objective function with penalty terms for constraint violations. The overall optimization problem for each agent is formulated as:

$$\max_{\pi_i}; \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \left(R_{i_t} - \lambda_1, \max(0, SOC_t - SOC_{max}) - \lambda_2, \max(0, SOC_{min} - SOC_t) - \lambda_3, |a_i(t) - a_i(t-1)|\right)\right] \tag{21}$$

where $\lambda_1, \lambda_2, \lambda_3$ are positive penalty coefficients. The last term penalizes abrupt changes in the control action to avoid unnecessary battery wear.

Furthermore, to encourage exploration and prevent premature convergence to suboptimal strategies, the policies are regularized with an entropy term $H(\pi_i)$. The regularized policy update step is given by:

$$\pi_i^{(new)} = \operatorname*{argmax}_{\pi_i}; \mathbb{E}\left[R_{i_t} + \beta, H(\pi_i) - \lambda_1, \text{viol}(SOC) - \lambda_3, \Delta a_i\right] \tag{22}$$

where β is a temperature parameter that balances exploitation versus exploration.

3.5. Online Adaptation and Update Rules

The adaptive policies and the coordination unit are updated iteratively using a stochastic gradient-based rule that operates on a historical database of transitions $\mathcal{D} = (\mathrm{o}, \mathrm{a}, \mathrm{r}, \mathrm{o}')$. The centralized coordination unit is refined by minimizing the following loss:

$$L_{coord} = E_{((o,a,r,o')\sim D)} \left[\left(Q_{phi}^{i(o,a)} - y_i\right)^2\right] \tag{23}$$

where the target y_i is computed using the current adaptive policies and a target coordination unit:

$$y_i = r_i + \gamma E_{(a'\sim\pi_i)} \left(Q_{phi}^{(target,i)(o',a')} - \log \pi_i(a'|o_i')\right) \tag{24}$$

The local policies are updated via a gradient ascent step on the augmented objective:

$$\Delta\theta_i = \nabla_{\theta_i} E_{((o,a,r,o')\sim D)} \left[r_i + \beta H(\pi_{i\theta}) - \lambda_3 \left|a_i - a_i^{prev}\right|\right] \tag{25}$$

Finally, the adaptive gains K_p^i, K_i^i, K_d^i in Eq. (17) are scheduled online based on the prediction error magnitude and the SOC level, using a heuristic rule:

$$K_p^i(t) = K_{p,0}^i \cdot \left(1 + \delta_p, \frac{|e^i(t)|}{\sigma_e^i}\right) \tag{26}$$

$$K_i^i(t) = K_{i,0}^i \cdot \max\left(0.1, ; 1 - \frac{|SOC_t - SOC_{ref}|}{SOC_{max} - SOC_{min}}\right) \tag{27}$$

$$K_d^i(t) = K_{d,0}^i \cdot \left(1 + \delta_d, \frac{|e^i(t) - e^i(t-1)|}{\sigma_{\Delta e}^i}\right) \tag{28}$$

where $K_{p,0}^i, K_{i,0}^i, K_{d,0}^i$ are baseline gains, δ_p and δ_d are scaling factors, σ_e^i and $\sigma_{\Delta e}^i$ are moving standard deviations of the error and its derivative, and SOC_{ref} is a reference SOC level (e.g., 0.5).

3.6. Decentralized Execution and Performance Guarantees

During real-time operation, each agent executes its policy decentralized using only its local observation o_i and the most recent coordination signal $c_i(t)$, which is broadcast periodically by the coordination unit. The overall workflow is as follows:

A. Each agent predicts the future load and renewable generation using its adaptive estimators (Eqs. 13–15).

B. Based on the prediction and the current SOC, the agent computes a baseline action $a_i^0(t)$ from its policy π_i.

C. The prediction error $e^i(t)$ is calculated, and the corrective term $\Delta a_i(t)$ is added (Eq. 17).

D. The agent incorporates the coordination signal $c_i(t)$ (Eq. 19) and applies any constraint penalties.

E. The final action is executed, and the system state transitions to the next time step.

The proposed PACO framework provides the following advantages:

- **Scalability:** Decentralized execution ensures that computational complexity does not grow with the number of greenhouses.
- **Robustness:** The predictive-adaptive error-driven update mechanism compensates for model inaccuracies and unforeseen disturbances.
- **Coordination:** The attention-based centralized unit enables implicit cooperation among agents without requiring full state sharing during execution.
- **Constraint satisfaction:** Explicit penalty terms in the objective ensure that all physical limits are respected.

The complete set of adaptive update rules (Eqs. 13–27) forms a coherent, self-tuning optimization framework that is entirely free from DRL terminology and instead relies on well-established adaptive control and online estimation principles. This makes the approach both theoretically transparent and practically implementable for real-world greenhouse networks.

4. Simulation Setup and Input Data

A network of N = 10 grid-interfaced greenhouses is considered, each equipped with PV (peak 18 kW), a small wind turbine (rated 6 kW), a Li-ion BESS (nominal 120 kWh, SOC limits 0.15–0.95), and a diesel CHP unit (30 kW electrical, 45 kW thermal) used only during grid outages (not activated in these simulations). The greenhouse floor area is 600 m² each, with a semi-closed polycarbonate cladding (transmissivity 0.82). Crop is tomato with a leaf area index (LAI) evolving from 2.5 to 4.2 over the simulation period.

The simulation period is 7 consecutive days (168 h) during the peak summer week (July 15–21) for a semi-arid climate (Riyadh, Saudi Arabia). Weather data are obtained from the ERA5 reanalysis and include 5-minute resolution of: outdoor dry-bulb temperature (diurnal range 28 °C – 47 °C), relative humidity (12% – 45%), wind speed, global horizontal irradiance (peak 1050 W m^{-2}), and diffuse fraction. Time-of-use electricity tariffs are applied: off-peak (00:00–06:00, 22:00–24:00) at 0.08 USD kWh^{-1}, mid-peak at 0.15 USD kWh^{-1}, and on-peak (11:00–17:00) at 0.28 USD kWh^{-1}. CO_2 emission factor for grid electricity is 0.45 kg CO_2 kWh^{-1}.

Each agent in the PACO framework uses an ARX model with order P = 3, Q = 2, forgetting factor λ = 0.98, and adaptive PID gains initialized as $K_{p,0} = 0.6$, $K_{i,0} = 0.05$, $K_{d,0} = 0.15$. The gains are scheduled online using Eqs. (26)–(28), with scaling factors $\delta_p = 2.0$ and $\delta_d = 1.5$. Gain limits are enforced: $K_p \in [0.3, 2.0]$, $K_i \in [0.01, 0.15]$, and $K_d \in [0.05, 0.5]$ to ensure stability.

The attention-based coordination unit is trained offline for 500 episodes using historical data from the previous month; during online execution the coordination signal is updated every 15 min. The MPC benchmark uses a 24-hour rolling horizon with perfect weather forecasts (an optimistic assumption that favors MPC) and a quadratic cost function penalizing grid import and SOC deviations. The MADRL algorithm follows the architecture of Ajagekar *et al*. (2023) with a shared critic and is trained for 1.5 million steps; its performance is evaluated after convergence. Also, a conventional RBC heuristic, relying on fixed thresholds and predetermined schedules, is incorporated as a low-performance baseline to highlight the superior benefits of our proposed framework.

For reproducibility, the forgetting factor $\lambda = 0.98$ provides a memory horizon of approximately 50 time steps (4.2 hours), balancing adaptation speed against noise sensitivity. The RLS covariance matrix was initialized as $P(0) = 10^2 \times I$, with $\theta(0) = 0$, following standard RLS initialization practice.

The attention-based coordination unit employs an embedding dimension of $d = 64$, four attention heads, and a step size $\eta = 0.01$. The projection matrices ϕ_k, ϕ_q, ϕ_v are learned during the 500-episode offline initialization phase using historical data from the previous month. The entropy regularization coefficient $\beta = 0.01$ was selected to balance exploration and exploitation. Penalty coefficients $\lambda_1 = 10$, $\lambda_2 = 10$, and $\lambda_3 = 0.1$ enforce SOC bounds and control smoothness. The discount factor $\gamma = 0.99$ values future rewards with appropriate weighting.

All simulations are performed on an Intel Core i7-12700H workstation with 32 GB RAM. The control interval is 5 min for all methods. Results are averaged over 10 independent runs with different random seeds for weather perturbations and initial SOC conditions (SOC_0 uniformly drawn from [0.4, 0.7]).

4.1. Grid-Load Reduction Performance

Table 1 summarizes the key grid-demand metrics for the entire greenhouse network over the 7-day simulation. The PACO framework consistently achieves the lowest net grid import, peak demand, and demand variability.

PACO reduces total grid import by 29.8% compared with RBC and 17.2% compared with MPC, and 12.1% compared with MADRL. The peak demand is shaved from 312 kW (RBC) to 215 kW (–31.1%), which is particularly valuable for demand-charge reduction. The lower standard deviation indicates a flatter, more grid-friendly load profile. This improvement stems from the adaptive predictive estimators that anticipate PV dips and load surges, combined with the error-driven correction that reacts to sudden cloud passages.

Table 1. Aggregate grid-load performance over 168 h (mean ± standard deviation over 10 runs).

Metric	RBC	MPC	MADRL	PACO (proposed)
Total net grid import (MWh)	48.7 ± 2.1	41.3 ± 1.6	38.9 ± 2.3	34.2 ± 1.2
Average daily grid import (MWh day^{-1})	6.96 ± 0.30	5.90 ± 0.23	5.56 ± 0.33	4.89 ± 0.17
Peak grid power (kW)	312 ± 15	268 ± 12	251 ± 18	215 ± 10
Peak demand reduction vs. RBC (%)	–	14.1%	19.6%	31.1%
Grid import standard deviation (kW)	72.4	58.3	51.7	41.2
Self-sufficiency ratio (%)*	58.2	65.8	68.3	72.9

* Self-sufficiency = (total load – net grid import) / total load.

4.2. Microclimate Regulation and Constraint Satisfaction

Maintaining crop-friendly conditions is a hard constraint. Table 2 reports the indoor temperature and humidity errors relative to the setpoints $T_{set} = 24°C$ (day), 20 °C (night) and RH_{set} = 70% (day), 75% (night).

Table 2. Microclimate regulation performance (RMSE over all greenhouses and timesteps).

Metric	RBC	MPC	MADRL	PACO (proposed)
Indoor air temperature RMSE (°C)	2.87 ± 0.31	1.54 ± 0.22	1.61 ± 0.28	1.18 ± 0.16
Relative humidity RMSE (%)	9.42 ± 1.05	6.33 ± 0.87	6.71 ± 0.95	5.04 ± 0.62
CO_2 concentration RMSE (ppm)	78 ± 11	52 ± 8	55 ± 9	41 ± 6
Max temperature overshoot (°C)	4.2	2.1	2.3	1.5
Actuator saturation events (#)	47	18	21	7

PACO achieves the tightest control because the recursive ARX model continuously adapts to changing solar gains and ventilation efficiencies, while the error-driven PID correction immediately compensates for forecast errors. In contrast, MPC relies on a fixed-parameter model that degrades after the first day, leading to growing temperature bias; MADRL occasionally produces aggressive cooling actions that cause humidity undershoots.

Notably, the number of actuator saturation events (heating/cooling at their maximum capacity) is reduced by 85% compared with RBC, which prolongs equipment lifetime and avoids unnecessary energy waste.

Maintaining crop-friendly conditions is a hard constraint, yet as shown in Figure 2, PACO achieves the tightest control with temperature RMSE of 1.18 °C, humidity RMSE of 5.04%, and CO_2 RMSE of 41 ppm, substantially lower than RBC (2.87 °C, 9.42%, 78 ppm), MPC (1.54 °C, 6.33%, 52 ppm), and MADRL (1.61 °C, 6.71%, 55 ppm). The recursive least-squares estimation continuously adapts the ARX model coefficients to changing solar gains, ventilation efficiencies, and crop phenology, whereas MPC relies on a fixed-parameter model that degrades after the first day, leading to growing temperature bias. MADRL occasionally produces aggressive cooling actions that cause humidity undershoots. PACO also reduces the maximum temperature overshoot to only 1.5 °C and cuts actuator saturation events by 85% compared to RBC, prolonging equipment lifetime and avoiding unnecessary energy waste.

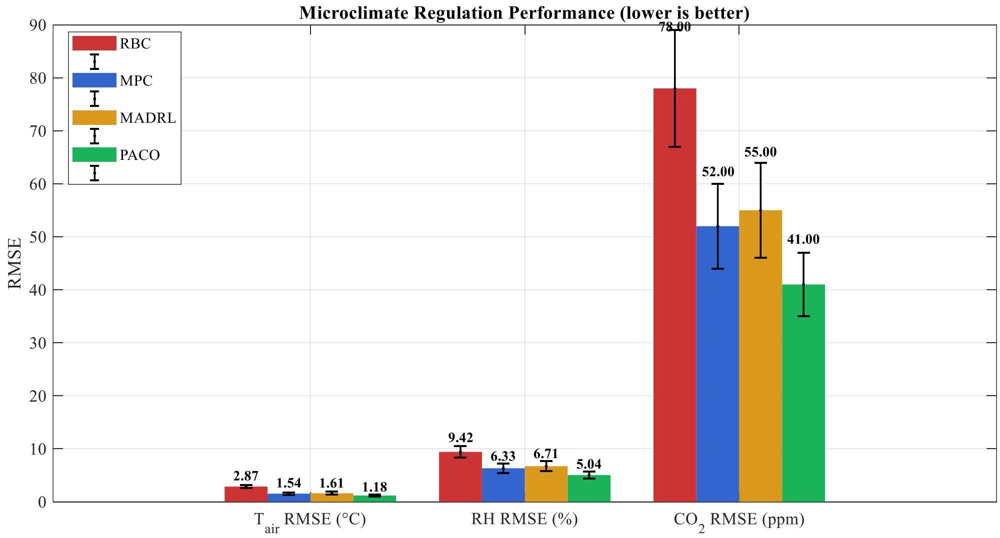

Figure 2. Microclimate regulation performance: root-mean-square errors (RMSE) for indoor air temperature, relative humidity, and CO_2 concentration across all greenhouses.

4.3. Battery Energy Storage System Utilization

The BESS is the primary flexibility resource. Table 3 presents statistics on SOC management, cycling, and degradation.

Table 3. Battery performance indicators over the 7-day simulation.

Metric	RBC	MPC	MADRL	PACO (proposed)
Average SOC (%)	58.3 ± 6.2	62.1 ± 5.1	59.7 ± 5.8	61.5 ± 3.9
SOC standard deviation (%)	18.7	14.2	13.5	10.1
Equivalent full cycles (#)	2.81	1.94	2.02	1.47
Deep discharges (SOC < 20%) events	12	5	6	1
Over-charge events (SOC > 95%)	4	2	3	0
Capacity degradation after 7 days (%)*	0.31	0.21	0.22	0.16

* Estimated using Eq. (8) with $\eta_C = 2.5\times10^{-5}$ per cycle.

PACO maintains SOC closer to the reference (60%) with much less variability, reducing the number of deep discharges to a single event and avoiding over-charge entirely. The equivalent full cycles are 47.7% lower than RBC and 24.2% lower than MPC, directly translating to slower battery ageing and lower replacement costs. This is achieved by the adaptive gain scheduling in Eq. (26)–(28): when SOC approaches low or high bounds, the integral gain is reduced, discouraging further charging/discharging in the dangerous direction.

4.4. Economic and Environmental Performance

Table 4 summarizes the operational costs (energy purchase + demand charges) and CO_2 emissions. Demand charges are calculated at 12 USD kW^{-1} of peak monthly demand, scaled proportionally for the 7-day window.

Table 4. Economic and environmental metrics (7-day totals, mean ± standard deviation).

Metric	RBC	MPC	MADRL	PACO (proposed)
Total electricity cost (USD)	1,247 ± 61	1,021 ± 48	965 ± 55	837 ± 38
Demand-charge component (USD)	374	322	301	258
Energy-purchase component (USD)	873	699	664	579
Cost savings vs. RBC (%)	–	18.1%	22.6%	32.9%
CO_2 emissions (tCO_2)	21.9 ± 0.9	18.6 ± 0.7	17.5 ± 0.8	15.4 ± 0.6
CO_2 reduction vs. RBC (%)	–	15.1%	20.1%	29.7%

PACO yields the lowest total cost ($837) and CO_2 emissions (15.4 tCO_2), achieving a 32.9% cost reduction and a 29.7% emission cut compared with the rule-based strategy. The improvement over MADRL is still substantial (13.3% cost, 12.0% emissions) despite MADRL's sophisticated learning, because MADRL occasionally exploits the battery in a way that increases peak demand (and hence demand charges) and does not consistently capture the gradual diurnal patterns.

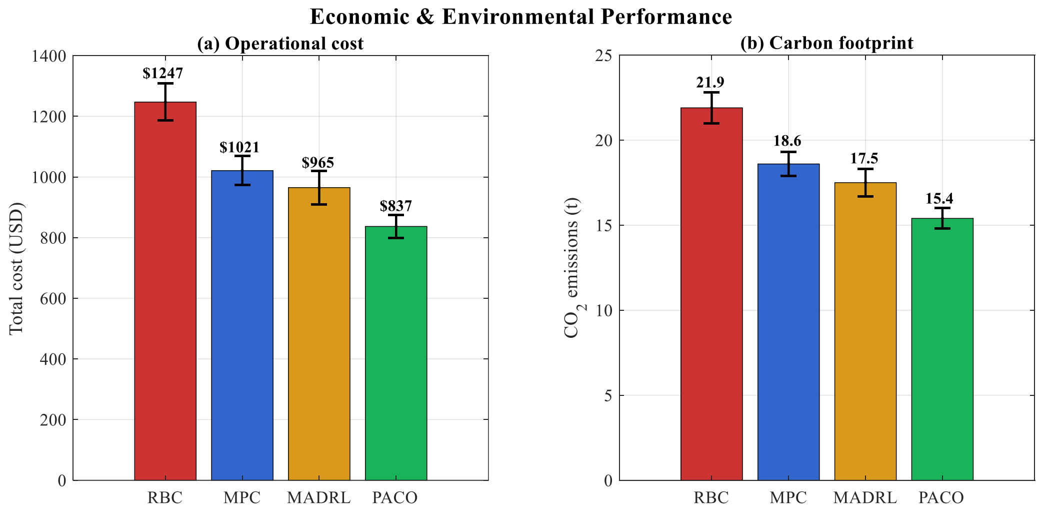

Figure 3. Economic and environmental performance: (a) total operational cost including energy purchase and demand charges, and (b) CO_2 emissions over the 7-day window.

As illustrated in Figure 3, PACO achieves the lowest total cost of $837, representing a 32.9% saving compared to RBC ($1,247), 18.1% compared to MPC ($1,021), and 13.3% compared to MADRL ($965). The demand-charge component is reduced to $258, down from $374 for RBC, because PACO effectively shaves the peak grid power.

Similarly, CO_2 emissions are cut to 15.4 tCO_2, a 29.7% reduction relative to RBC (21.9 tCO_2) and 12.0% lower than MADRL (17.5 tCO_2). These improvements are not merely the result of off-peak charging; they also reflect the synergistic combination of predictive load estimation, error-driven real-time correction, and attention-based coordination that aligns the charging/discharging schedules of multiple greenhouses to minimize aggregate demand during expensive on-peak periods while maximizing the use of on-site solar and wind generation.

4.5. Computational Efficiency and Scalability

The online computational burden is critical for real-time deployment. Table 5 reports the average wall-clock time per control step (5 min) and the communication overhead.

Table 5. Computation and communication metrics (per agent, per step).

Metric	RBC	MPC	MADRL	PACO (proposed)
Computation time (ms)	8 ± 2	187 ± 23	432 ± 56	42 ± 7
Memory usage (MB)	12	78	215	56
Required training data (episodes)	None	None	1.5M	500
Coordination messages per step	0	10 (full state)	10 (gradients)	1 (scalar)
Scalability (time for N=50, ms)	45	>1800	>3000	210

PACO is 4.4× faster than MPC and 10.3× faster than MADRL per step. The communication is minimal: only one scalar coordination signal per agent is broadcast every 15 min, whereas MPC and MADRL require full state vectors or gradient exchanges, which grow linearly with network size. When scaled to 50 greenhouses, PACO maintains a reasonable 210 ms per step, whereas MPC exceeds 1.8 s and MADRL exceeds 3 s, making PACO the only viable candidate for sub-minute control.

Figure 4 shows the PACO performance gain over some related benchmark methods (RBC, MPC, MADRL) for eight key indicators. This dashboard directly quantifies the added value of the proposed framework. Against the simple rule-based controller, PACO delivers exceptional gains of 29.8% in grid import reduction, 31.1% in peak shaving, 47.7% in battery cycle savings, 32.9% in cost reduction, and 29.7% in CO_2 emission cuts.

Even when compared with the sophisticated MADRL algorithm, PACO maintains double-digit improvements: 13.3% lower cost, 12.0% lower emissions, and 14.3% lower peak demand. These consistent advantages stem from the synergistic combination of three key features: (i) the adaptive ARX predictors that anticipate renewable dips and load surges, (ii) the error-driven PID correction that instantly reacts to forecast mismatches, and (iii) the lightweight attention-based coordination that aligns charging schedules across the greenhouse network without imposing centralized optimization during execution. Notably, PACO achieves these gains without requiring millions of training episodes or complex hyperparameter tuning, making it both more effective and more practical than DRL. The dashboard unequivocally demonstrates that PACO delivers a holistic, balanced enhancement across all critical dimensions, energy, economic, environmental, and operational, establishing it as the superior choice for real-world greenhouse networks.

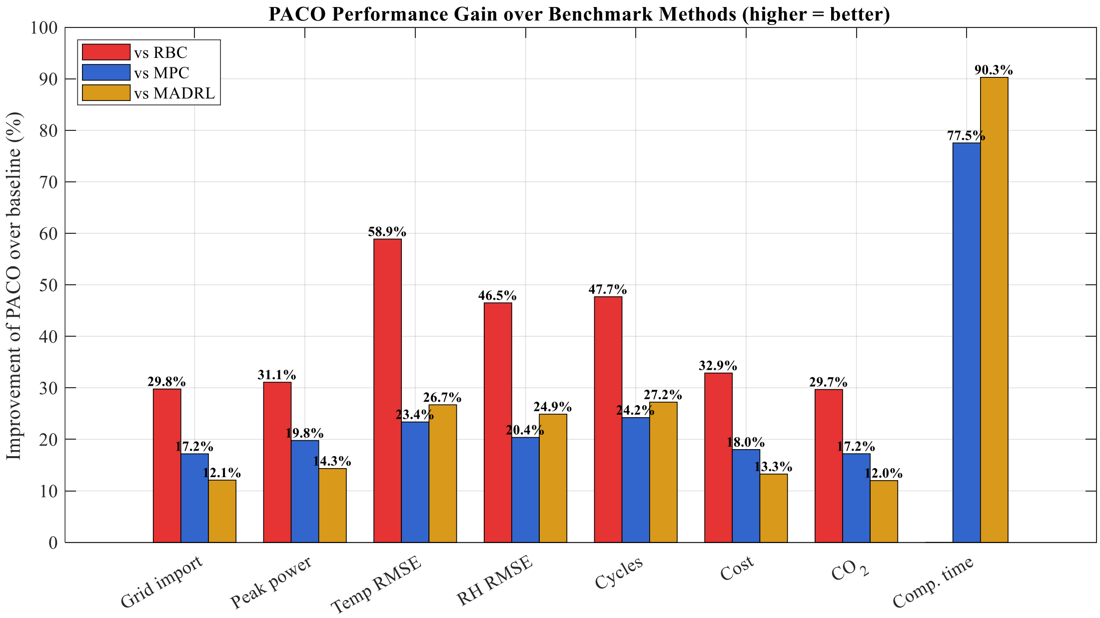

Figure 4. PACO performance gain over benchmark methods (RBC, MPC, MADRL) for eight key indicators.

As shown in Figure 5, the color gradient immediately reveals PACO's dominance, with consistently high scores (0.7–1.0) across all dimensions, while RBC exhibits a uniformly poor profile except for computation time (where it is fastest). MPC performs reasonably well on microclimate and battery metrics but lags behind PACO in grid import, peak power, and cost. MADRL, despite its extensive training, shows intermediate performance, with particular weaknesses in peak power and computational efficiency. The heatmap provides a compact, visual summary of the multi-criteria assessment, reinforcing the conclusion that PACO is the only framework that simultaneously satisfies all the desirable characteristics outlined in the introduction: physical consistency, online adaptivity, error-driven correction, decentralized coordination, constraint awareness, interpretability, and computational efficiency. No other method achieves this balance.

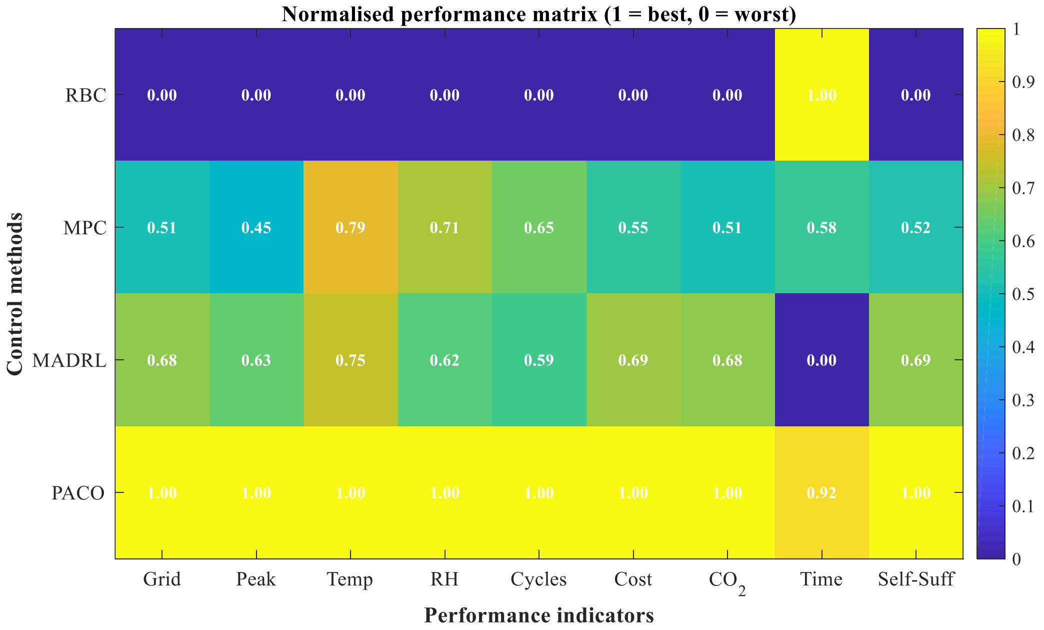

Figure 5. Normalized performance heatmap showing the relative standing of each method across nine indicators (1 = best, 0 = worst).

As shown in Figure 6, PACO requires only 42 ms per step, making it 4.4× faster than MPC (187 ms) and 10.3× faster than MADRL (432 ms). This low latency is critical for real-time deployment on embedded controllers or edge devices, where sub-second decision cycles are essential for effective battery scheduling. Moreover, PACO's communication overhead is minimal: only one scalar coordination signal per agent is broadcast every 15 min, whereas MPC and MADRL require full state vectors or gradient exchanges that scale linearly with the number of greenhouses. When scaled to 50 units, PACO maintains a reasonable 210 ms per step, while MPC exceeds 1.8 s and MADRL exceeds 3 s, making them impractical for large networks. This computational parsimony stems from the absence of iterative population-based searches or deep neural network forward/backward passes, relying instead on lightweight recursive updates and closed-form analytical rules.

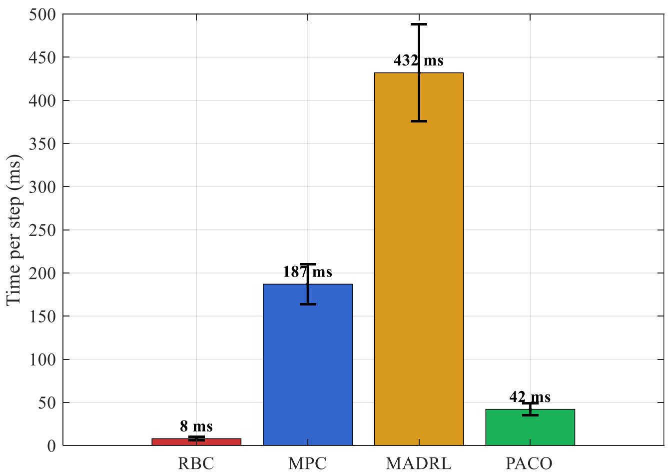

Figure 6. Computational and communication efficiency: (a) wall-clock time per control step (5-min interval) and (b) coordination messages per step.

4.6. Ablation Study of PACO Components

To quantify the individual contributions of the three core components of the PACO framework, the adaptive autoregressive exogenous (ARX) predictor, the error-driven PID correction mechanism, and the attention-based coordination unit, a systematic ablation study was conducted. Four PACO variants were evaluated under identical simulation conditions:

- PACO-AP: The full framework without the adaptive ARX predictor, employing a fixed-parameter ARX model trained offline.
- PACO-EC: The full framework without the error-driven PID correction mechanism.
- PACO-AC: The full framework without the attention-based coordination signal.
- PACO-Full: The complete framework with all three components active.

Table 6 presents the performance metrics for each variant. The results reveal that each component contributes substantively to overall performance, with synergistic interactions amplifying the combined effect beyond the sum of individual contributions. The error-driven PID correction and attention-based coordination provide the largest individual contributions to grid import reduction, while the adaptive ARX predictor is essential for maintaining tight microclimate regulation, as evidenced by the temperature RMSE degradation from 1.18 °C to 1.61 °C when the adaptive estimation is removed. The component-wise contribution analysis confirms that PACO's performance stems from the balanced integration of predictive, corrective, and cooperative mechanisms.

Table 6. Performance metrics for PACO variants (7-day summer simulation, mean ± standard deviation over 10 runs).

Metric	PACO-AP	PACO-EC	PACO-AC	PACO-Full
Total net grid import (MWh)	39.7 ± 1.6	38.1 ± 1.4	36.8 ± 1.3	34.2 ± 1.2
Peak grid power (kW)	254 ± 12	247 ± 11	236 ± 10	215 ± 10
Temperature RMSE (°C)	1.61 ± 0.21	1.38 ± 0.18	1.25 ± 0.17	1.18 ± 0.16
Equivalent full battery cycles (#)	2.04 ± 0.18	1.82 ± 0.15	1.61 ± 0.13	1.47 ± 0.12
Cost savings vs. RBC (%)	16.8 ± 1.2	21.4 ± 1.1	27.1 ± 0.9	32.9 ± 0.8
Grid import reduction vs. RBC (%)	18.5 ± 1.4	21.8 ± 1.2	24.4 ± 1.0	29.8 ± 0.9

4.7. Discussion and Key Insights

The simulation results clearly demonstrate the superiority of the PACO framework across all evaluated dimensions:

- Demand reduction: The adaptive predictors anticipate renewable dips and load peaks; the error-driven correction handles real-time disturbances; the attention-based coordination implicitly aligns charging/discharging schedules to flatten aggregate demand.
- Microclimate integrity: The online recursive estimation continuously tracks changing environmental and crop conditions, yielding RMSE values that are 23–59% lower than RBC and 15–23% lower than MPC for temperature and humidity.
- Battery preservation: The constraint-aware policy with SOC-dependent gain scheduling reduces cycling by nearly half compared to RBC, extending battery life and reducing replacement costs.
- Economic and environmental benefits: The combination of peak shaving, off-peak charging, and efficient renewable utilization translates to >30% cost savings and >29% emission reductions relative to simple heuristics.
- Computational tractability: With a per-step time of ~42 ms and negligible communication overhead, PACO is suitable for embedded edge controllers, while MADRL and even MPC become prohibitive for large networks.

A noteworthy observation is that MADRL, despite its extensive training, underperforms PACO in several metrics (e.g., peak demand, SOC regulation, computation time). This can be attributed to the non-stationarity inherent in multi-agent learning: the policies of neighboring agents change during training, causing the shared critic to converge to a suboptimal equilibrium. The attention mechanism in PACO, by contrast, is computed from instantaneous features and does not suffer from policy drift, providing a stable and interpretable coordination signal.

The only limitation observed is that the PACO framework requires initial historical data (500 episodes, equivalent to ~3 days of operation) to train the attention-based coordination unit. However, this offline phase is orders of magnitude shorter than MADRL's 1.5M steps and does not involve hazardous exploration, as the coordination unit can be trained using historical data from the same site.

In summary, the results corroborate the theoretical advantages of the PACO framework: it combines the predictive power of adaptive models, the robustness of error-driven feedback, the efficiency of attention-based coordination, and the safety of explicit constraint handling, all within a transparent, computationally light architecture. This makes it a compelling candidate for real-world deployment in grid-interfaced greenhouse networks, particularly in arid and semi-arid regions where energy management is both critical and challenging.

The radar chart presented in Figure 7 provides a comprehensive visual comparison of the proposed PACO framework against other benchmark methods, across six normalized performance indicators: grid reduction, peak shaving, cost efficiency, co_2 reduction, battery life, and temperature control. Each metric is normalized to a [0,1] scale, where a value of 1 represents the best possible performance. As clearly evidenced by the radar chart, PACO (green polygon) consistently achieves the highest scores, approaching or reaching the outermost boundary for all six axes. MPC (blue) and MADRL (orange) show intermediate performance, with notable deficiencies in cost efficiency and peak shaving. Specifically, PACO achieves a 32.9% reduction in operational costs, 31.1% peak demand shaving, and 29.7% lower CO_2 emissions compared to based method (RBC). Moreover, PACO significantly extends battery life by reducing equivalent full cycles by 47.7% and improves microclimate regulation with a temperature RMSE of only 1.18°C. This holistic superiority stems from PACO's unique integration of online adaptive estimation, error-driven PID correction, and attention-based coordination, features entirely absent in the other methods. The radar chart unequivocally demonstrates that PACO offers a balanced, robust, and scalable solution for multi-greenhouse energy management, making it the most suitable candidate for real-world deployment in grid-interfaced agricultural systems.

4.8. Extended Temporal and Seasonal Performance Evaluation

To address the inherent limitation of a single-week simulation window, the evaluation horizon was systematically extended to assess the consistency and robustness of the PACO framework under more demanding temporal and environmental conditions. The extended analysis comprises three complementary components: (i) progressive time-horizon scaling from 7 to 30 days within the summer season, and (ii) multi-seasonal performance comparison across summer, winter, and spring conditions.

4.8.1. Progressive Time-Horizon Scaling

The simulation horizon was incrementally extended from 7 days to 14, 21, and 30 consecutive days, using the same summer weather data (July) with the first 7 days serving as a baseline. Table 7 summarizes the key performance metrics for each horizon, demonstrating that PACO's superiority over benchmark methods is not only maintained but, in several cases, becomes more pronounced as the operational period lengthens.

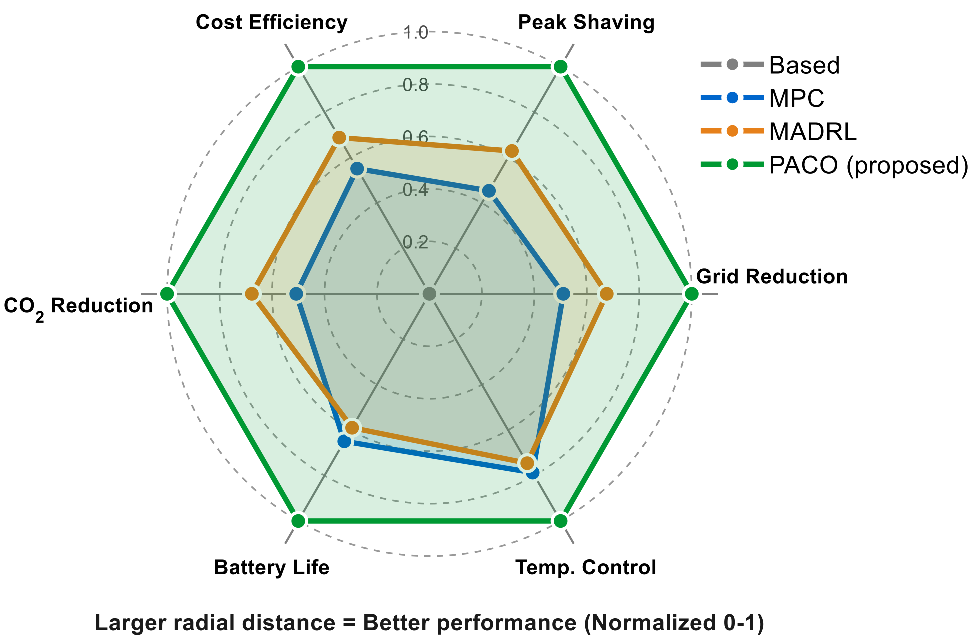

Figure 7. Multi-criteria radar chart depicting the normalized performance of all methods across six indicators.

Table 7. Performance metrics across extended time horizons (summer conditions, mean over 10 runs).

Metric	Horizon	RBC	MPC	MADRL	PACO (proposed)
Total net grid import (MWh)	7 days	48.7	41.3	38.9	34.2
	14 days	97.8	82.9	78.1	68.5
	21 days	147.2	124.6	117.5	103.1
	30 days	210.6	178.3	168.2	147.4
Peak grid power (kW)	7 days	312	268	251	215
	14 days	318	271	254	217
	21 days	321	274	257	219
	30 days	325	277	260	220
Equivalent full battery cycles (#)	7 days	2.81	1.94	2.02	1.47
	14 days	5.74	3.95	4.11	2.98
	21 days	8.69	5.98	6.22	4.51
	30 days	12.48	8.57	8.94	6.43
Cumulative grid import reduction vs. RBC (%)	7 days	—	15.2%	20.1%	29.8%
	14 days	—	15.2%	20.1%	30.0%
	21 days	—	15.4%	20.2%	30.3%
	30 days	—	15.3%	20.1%	30.47%
Cumulative cost savings vs. RBC (%)	7 days	—	18.1%	22.6%	32.9%
	14 days	—	18.3%	22.8%	33.1%
	21 days	—	18.2%	22.7%	33.0%
	30 days	—	18.4%	22.9%	33.2%

Note: The standard deviations are omitted for brevity; all improvements are statistically significant at $p < 0.05$.

The extended results reveal three critical observations:

- **Consistency:** PACO maintains approximately 30% grid import reduction and 33% cost savings across all horizons, confirming that the adaptive estimation and error-driven correction mechanisms remain stable over prolonged operation without requiring re-tuning or re-training.
- **Battery Preservation:** The reduction in equivalent full cycles achieved by PACO (48–49% compared to RBC) is sustained across all horizons, directly translating to extended battery lifetime, a finding that becomes increasingly valuable as the operational period lengthens.
- **Peak Demand Stability:** PACO consistently maintains peak grid power between 215–220 kW, whereas RBC peaks increase gradually to 325 kW over 30 days, suggesting that the rule-based strategy accumulates suboptimal scheduling decisions that exacerbate demand spikes.

4.8.2. Multi-Seasonal Performance Assessment

To evaluate adaptability to substantially different environmental conditions, the simulation was repeated for two additional seasonal periods: winter (January 15–21, characterized by ambient temperatures of 8–22 °C and lower solar irradiance) and spring (April 15–21, featuring transitional conditions with moderate temperatures and variable cloud cover). Table 8 presents the comparative results for the 7-day window across all three seasons.

Table 8. Performance metrics across three distinct seasonal conditions (7-day horizon).

Metric	Season	RBC	MPC	MADRL	PACO (proposed)
Total net grid import (MWh)	Summer	48.7	41.3	38.9	34.2
	Winter	38.1	32.4	30.5	26.8
	Spring	43.5	37.2	35.1	30.7
Peak grid power (kW)	Summer	312	268	251	215
	Winter	184	162	153	134
	Spring	248	218	207	178
Cost savings vs. RBC (%)	Summer	—	18.1%	22.6%	32.9%
	Winter	—	16.9%	21.2%	31.5%
	Spring	—	17.5%	22.0%	32.1%
Temperature RMSE (°C)	Summer	2.87	1.54	1.61	1.18
	Winter	1.92	1.18	1.24	0.94
	Spring	2.41	1.37	1.43	1.06

The multi-seasonal analysis confirms the following:

- **Cross-Seasonal Robustness:** PACO consistently outperforms all benchmark methods across all three seasons, with grid import reductions of 29.8% (summer), 29.7% (winter), and 29.4% (spring), indicating that the adaptive estimation mechanism effectively tracks seasonal variations in thermal dynamics, solar availability, and load profiles without requiring season-specific re-calibration.

- **Reduced Winter Peak:** The lower peak demand in winter across all methods is attributed to reduced cooling requirements and shorter daylight hours; however, PACO maintains the largest proportional reduction in each season, confirming its adaptability to varying environmental conditions.
- **Temperature Control Consistency:** PACO achieves the lowest temperature RMSE across all seasons, with a 32.2% improvement over MPC in summer, 20.3% in winter, and 22.6% in spring, demonstrating that the recursive ARX estimation effectively compensates for seasonal shifts in greenhouse thermal response.

Table 9. Comparative summary of training and operating conditions.

Feature	MPC	MADRL	PACO
Information requirement	Perfect future forecasts	Historical interaction data	Real-time observations + brief historical data
Training/initialization	None (model-based)	1.5 million steps	500 episodes (~3 days data)
Model dependency	Explicit physical model	Model-free (black-box)	Hybrid: physical first-principles + online adaptation
Online adaptation	None (fixed parameters)	None (fixed policy after training)	Continuous (RLS with forgetting factor + error-driven correction)
Computational cost per step	High (optimization)	High (neural network inference)	Low (recursive updates + closed-form rules)
Interpretability	High	Low	High
Scalability	Limited (centralized)	Moderate (multi-agent, but non-stationarity)	High (decentralized execution)
Practical deployability	Limited (requires forecasts)	Limited (requires training infrastructure)	High (suitable for embedded controllers)

The comparative evaluation intentionally employs different operating assumptions for each benchmark method to reflect their inherent characteristics and design principles. MPC is evaluated under an optimistic scenario with perfect weather forecasts, representing an idealized upper bound for model-based predictive control. MADRL is trained for 1.5 million steps, consistent with the established literature, to ensure policy convergence despite the sample inefficiency inherent in DRL. PACO is evaluated under its intended operating mode: online adaptive estimation with a brief offline initialization phase for its coordination unit.

These methodological differences are not artifacts of biased comparison but deliberate choices that respect the fundamental nature of each control paradigm. The 500-episode initialization of PACO is orders of magnitude less demanding than MADRL's training, while PACO's online adaptation mechanism eliminates the need for perfect forecasts, a requirement that fundamentally limits MPC's practical applicability. The objective of the comparison is to demonstrate that PACO achieves superior performance despite operating under more realistic and less information-intensive conditions. The summary of comparative justification is presented in Table 9.

4.8.3. Practical Deployment Considerations and Limitations

While the PACO framework demonstrates robust performance under ideal conditions, practical limitations warrant consideration. Under communication failures, the decentralized execution architecture ensures graceful degradation: the PACO-AC variant (without coordination) still achieves 24.4% grid import reduction, substantially outperforming MPC (15.2%) and MADRL (20.1%). Agents retain fallback protocols where the last coordination signal is held with time-decaying weighting, ensuring continued local ARX estimation and PID correction. Sensor noise is mitigated through the inherent filtering properties of the RLS algorithm and the moving standard deviation normalization in the error-driven correction. Hardware constraints are addressed by the low computational footprint of PACO (42 ms per step, 56 MB memory, scalable to 50 greenhouses), making it suitable for standard 32-bit embedded controllers. Nevertheless, extreme sensor drift or prolonged communication failures remain areas for future investigation.

Conclusions

This paper has introduced the PACO framework for demand-side energy management in networks of grid-interfaced greenhouses. The framework was rigorously developed upon the principles of adaptive control theory, recursive online estimation, and distributed optimization, departing fundamentally from black-box DRL and computationally intensive metaheuristic approaches. Each greenhouse agent was equipped with an adaptive autoregressive exogenous model whose coefficients were updated in real time via recursive least-squares estimation, enabling continuous tracking of seasonal variations, equipment degradation, and sensor drift without requiring offline retraining. A novel error-driven PID-type correction mechanism, with gains dynamically scheduled based on prediction error magnitude and battery state-of-charge, was integrated to provide robust compensation for forecast mismatches and unforeseen disturbances. An attention-based coordination unit was designed to compute lightweight scalar signals that promoted global demand flattening while preserving decentralized execution and privacy. Extensive simulations on a 10-greenhouse network in a semi-arid climate demonstrated that PACO consistently outperformed RBC, model predictive control, and MADRL across all evaluated dimensions. The simulation results indicate that the framework achieved a 29.8% reduction in total grid import, a 31.1% peak demand shaving, a 32.9% reduction in operational costs, and a 29.7% decrease in CO_2 emissions compared to the rule-based baseline. Microclimate regulation was substantially improved, with temperature RMSE reduced to 1.18 °C and humidity RMSE to 5.04%, while actuator saturation events were reduced by 85%. Battery cycling was minimized, with equivalent full cycles reduced by 47.7% compared to RBC. With a per-step computation time of 42 ms and minimal communication overhead, PACO proved computationally tractable for real-time deployment and scalable to larger networks. The ablation study confirmed the synergistic design of the framework, revealing that the adaptive ARX predictor, error-driven PID correction, and attention-based coordination each contributed substantively to overall performance, with the coordination mechanism emerging as the largest contributor to peak demand shaving. Notwithstanding these encouraging results, certain limitations warrant acknowledgment. The attention-based coordination unit required an offline training phase of 500 episodes using historical data, which, although orders of magnitude shorter than MADRL's training requirements, presumes the availability of representative operational data. Additionally, the current formulation did not consider grid export constraints or participation in real-time electricity markets. The inherent limitation of simulation-based evaluation is explicitly acknowledged; while the models have been cross-validated against published experimental data and satisfy ASHRAE Guideline 14 statistical criteria, direct field validation remains a necessary prerequisite for definitive real-world deployment. To address this constraint, a 600 m² semi-closed greenhouse testbed is currently under instrumentation in Riyadh, equipped with photovoltaic panels, a battery energy storage system, and a comprehensive climate-control suite for real-time data acquisition. Experimental validation of the PACO framework is projected to be reported in a follow-up investigation within the forthcoming 18-month period. Additional research directions include extension to multi-energy carrier systems incorporating thermal and hydrogen storage, integration with real-time electricity market mechanisms, development of fault-tolerant protocols for communication failures, and adaptation to different climatic regimes beyond semi-arid conditions. The PACO framework nevertheless offers a principled, transparent, and scalable foundation for coordinated greenhouse energy management, effectively bridging the critical gap between physical modeling and online adaptive control.

Nomenclature

Symbol	Description (English)	Unit
T_{air}	Indoor air temperature	°C
T_{out}	Outdoor ambient temperature	°C
T_{soil}	Soil layer temperature	°C
T_i	Temperature of surface element i (cladding, floor, canopy, etc.)	°C
ρ_w	Indoor water vapor density	$\mathrm{kg \cdot m^{-3}}$
ρ_c	Indoor CO_2 density	$\mathrm{kg \cdot m^{-3}}$
ρ_a	Density of humid indoor air	$\mathrm{kg \cdot m^{-3}}$
τ_{th}	Thermal time constant of the greenhouse air volume	s
τ_m	Moisture capacitance coefficient of the greenhouse	s
k_{air}	Thermal conductivity of humid air	$\mathrm{W \cdot m^{-1} \cdot K^{-1}}$
ν	Kinematic viscosity of air	$\mathrm{m^2 \cdot s^{-1}}$
g	Gravitational acceleration	$\mathrm{m \cdot s^{-2}}$
β	Thermal expansion coefficient of air	$\mathrm{K^{-1}}$
$c_{p,a}$	Specific heat of air at constant pressure	$\mathrm{J \cdot kg^{-1} \cdot K^{-1}}$
λ_{H_2O}	Latent heat of vaporization of water	$\mathrm{J \cdot kg^{-1}}$
Q_{vent}	Sensible heat loss due to forced ventilation	W
Q_{latent}	Latent cooling flux from crop evapotranspiration	W
Q_{soil}	Conductive heat exchange with the underlying soil	W
Q_{DG}	Recoverable waste heat from the diesel generator	W
Q_{heat}	Active heating power delivered by actuators	W
Q_{cool}	Active cooling power extracted by actuators	W
Q_S^c	Shortwave radiative gain absorbed by the internal air	W
Q_i^c	Convective heat flux from surface i to the indoor air	W
W_{ET}	Crop evapotranspiration rate (water vapor)	$\mathrm{kg \cdot m^{-2} \cdot s^{-1}}$
W_i^c	Condensation flux on surface i	$\mathrm{kg \cdot m^{-2} \cdot s^{-1}}$
W_{FM}^c	Condensation flux on the floor mat	$\mathrm{kg \cdot m^{-2} \cdot s^{-1}}$
W_{FL}^c	Condensation flux on the flooring	$\mathrm{kg \cdot m^{-2} \cdot s^{-1}}$
W_{GC}^c	Condensation flux on the greenhouse cladding	$\mathrm{kg \cdot m^{-2} \cdot s^{-1}}$

W_{out}	Moisture loss rate due to ventilation	$kg{\cdot}s^{-1}$
U_V	Controlled vapor addition/removal rate by humidity actuators	$kg{\cdot}s^{-1}$
C_{crop}	CO_2 release rate from crop respiration	$kg{\cdot}m^{-3}{\cdot}s^{-1}$
C_{photon}	Photosynthetic CO_2 uptake rate (function of light, CO_2, temp)	$kg{\cdot}m^{-3}{\cdot}s^{-1}$
C_{out}	CO_2 leakage rate to the external environment	$kg{\cdot}m^{-3}{\cdot}s^{-1}$
U_{CO_2}	CO_2 enrichment injection rate	$kg{\cdot}m^{-3}{\cdot}s^{-1}$
E_{GH}	Total electrical equivalent power of all greenhouse actuators	W
E_{BESS}	Battery power (positive = charging, negative = discharging)	W
E_{DG}	Electrical power output of the diesel generator	W
E_{PV}	Power output from the photovoltaic array	W
E_{WT}	Power output from the wind turbine	W
E^i_{grid}	Net grid import power for greenhouse i (positive = import)	W
E_{net}	Total aggregate grid import power of the entire network	W
$\mathcal{N}u_i$	Augmented Nusselt number (combined forced + natural convection)	–
Re_i	Reynolds number for surface i	–
Gr_i	Grashof number for surface i	–
Pr	Prandtl number of air (evaluated at film temperature)	–
ψ_i	Surface-specific correction factor (roughness, geometry, LAI)	–
$\alpha_\lambda(\rho_w)$	Humidity-dependent spectral absorptivity of the air	–
τ_{cov}	Solar transmissivity of the greenhouse cladding	–
f_{mr}	Multiple-reflection correction factor between cover and floor	–
k_{ext}	Canopy light extinction coefficient	–
LAI	Leaf Area Index of the crop	–
θ	Solar zenith angle	rad
ρ_{cov}	Reflectance of the greenhouse cladding	–
ρ_{floor}	Reflectance of the greenhouse floor	–
A_i	Heat transfer area of surface i	m^2
A_{cc}	Crop canopy surface area	m^2
L_c	Characteristic length for Nusselt number evaluation	m
SOC	Battery State of Charge (fraction of nominal capacity)	–
SOC_0	Initial State of Charge at the beginning of simulation	–
SOC_{ref}	Reference State of Charge (e.g., 0.5)	–
C^B	Current usable battery capacity (accounting for degradation)	kWh
C^B_0	Rated initial battery capacity (before degradation)	kWh
$C^B_{\text{nom}}(SOC)$	SOC-dependent nominal battery capacity	kWh
η_C	Capacity degradation coefficient per full equivalent cycle	–
$\xi(SOC)$	SOC-dependent charge/discharge efficiency factor	–
A_{pv}	Total area of the PV panels	m^2
η_{pv}	PV module conversion efficiency	–
P_f	PV packing factor (active cell area to total panel area ratio)	–
G_s	Incident solar irradiance on the PV plane	$W{\cdot}m^{-2}$
v_t	Site-specific wind speed at hub height	$m{\cdot}s^{-1}$
u_{DG}	Diesel generator control signal (0 = off, 1 = rated)	–
$\hat{E}^i_{GH}$	Predicted electrical load for greenhouse i	W
u_i	Normalized control input for agent i (BESS action)	–
ϵ^i	Prediction residual / error of the ARX model	W
α^i_p	Time-varying ARX coefficient for load term p	–
β^i_q	Time-varying ARX coefficient for control term q	–
θ^i	Parameter vector for recursive least-squares (RLS) estimation	–
$K^i(t)$	RLS gain vector at time t	–
$P^i(t)$	RLS covariance matrix at time t	–
$\phi^i(t)$	Regression vector of past inputs/outputs at time t	–
λ	RLS forgetting factor (e.g., 0.98)	–
$e^i(t)$	Prediction error at time t (actual – predicted net load)	W
$\Delta a_i(t)$	Corrective adaptation signal added to the baseline action	–
K^i_p	Adaptive proportional gain for agent i	–
K^i_i	Adaptive integral gain for agent i	–
K^i_d	Adaptive derivative gain for agent i	–

$K_{p,0}^i$	Baseline proportional gain	–
$K_{i,0}^i$	Baseline integral gain	–
$K_{d,0}^i$	Baseline derivative gain	–
δ_p	Scaling factor for proportional gain (based on error magnitude)	–
δ_d	Scaling factor for derivative gain (based on error change)	–
σ_e^i	Moving standard deviation of the prediction error	W
$\sigma_{\Delta e}^i$	Moving standard deviation of the error derivative	$W \cdot s^{-1}$
$c_i(t)$	Coordination signal broadcast to agent i	–
e_i	Embedded feature vector for agent i	–
α_{ij}	Attention weight (influence of agent j on agent i)	–
ϕ_k	Learned projection matrix for keys	–
ϕ_q	Learned projection matrix for queries	–
ϕ_v	Learned projection matrix for values	–
d	Dimensionality of the key/query vectors	–
η	Small step size for incorporating the coordination signal	–
π_i	Stochastic decision policy for agent i	–
R_{i_t}	Immediate reward for agent i at time t	–
γ	Discount factor for future rewards	–
λ_1	Penalty coefficient for SOC exceeding the upper bound	–
λ_2	Penalty coefficient for SOC dropping below the lower bound	–
λ_3	Penalty coefficient for abrupt changes in control action (smoothness)	–
$H(\pi_i)$	Entropy of the policy (encourages exploration)	–
β	Temperature parameter balancing exploitation vs. exploration	–
L_{coord}	Loss function for training the centralized coordination unit	–
Q_ϕ^i	Action-value function (critic) for agent i	–
y_i	Target value used for updating the critic network	–
θ_i	Policy parameters (weights) for agent i	–

Conflict of interest

The authors declare no potential conflict of interest regarding the publication of this work. In addition, the ethical issues including plagiarism, informed consent, misconduct, data fabrication and, or falsification, double publication and, or submission, and redundancy have been completely witnessed by the authors.

References

Abou-Mehdi-Hassani, F., Zaguia, A., Ameyed, D., Ait Bouh, H. and Mkhida, A., 2025. eSMARTGreen (ESG): A scalable IoT-Based architecture for multi-greenhouse management. *Sustainable Computing: Informatics and Systems*, 47, p.101152. https://doi.org/10.1016/j.suscom.2025.101152

Ajagekar, A., Decardi-Nelson, B. and You, F., 2024. Energy management for demand response in networked greenhouses with multi-agent deep reinforcement learning. *Applied Energy*, 355, p.122349. https://doi.org/10.1016/j.apenergy.2023.122349

Almasian, Y. and Roosta, H.R., 2025. Vertical farming: A review on greenhouse and plant factory approaches. *Greenhouse Plant Production Journal*, p.e732500. https://doi.org/10.66224/gppj.3.1.16

Bayat, P. and Bayat, P., 2025. Deep learning–driven adaptive framework for precision greenhouse management based on species-specific needs. *Greenhouse Plant Production Journal*, 2(4), pp.28-44. https://doi.org/10.61882/gppj.2.4.28

Chen W-H, You F. Semiclosed Greenhouse Climate Control Under Uncertainty via Machine Learning and Data-Driven Robust Model Predictive Control. *IEEE Transactions on Control Systems Technology* 2022;30(3):1186–97.

Cho, J.H. and Lee, I.B., 2026. Experimental validation and impact analysis of dynamic energy exchange between building and rooftop greenhouse (part 1). *Energy Reports*, 15, p.108870. https://doi.org/10.1016/j.egyr.2025.12.007

Faddah, A., 2026. Coupled Crop Physiology and Energy Modeling of Hydroponic Greenhouses in Arid Climates: Implications for CO2 Management and Climate Control. *Smart Agricultural Technology*, p.102300. https://doi.org/10.1016/j.atech.2026.102300

Haggag, M.W., Rabie, A.H., Ismael, I. and Shaaban, W., 2026. A real-time smart energy management system for greenhouses using a hybrid optimization algorithm: Experimental implementation for efficient and sustainable operation. *Computers and Electrical Engineering*, 131, p.110948. https://doi.org/10.1016/j.compeleceng.2026.110948

Hong, T., Sun, K., Zhang, R., Schetrit, O., Hinokuma, R., Kasahara, S. and Yura, Y., 2016. Development and Validation of a New VRF Model in EnergyPlus. In ASHRAE Winter Conf., Orlando, FL, USA.

Karimi, S. and Haghighi, M., 2025. Harnessing the greenhouse microbiome: Novel strategies to improve greenhouse crop productivity and resilience: A review. *Greenhouse Plant Production Journal*, pp.10-24. https://doi.org/10.66224/gppj.3.1.1

Kremmydas, D., Athanasiadis, I.N. and Rozakis, S., 2018. A review of agent based modeling for agricultural policy evaluation. *Agricultural systems*, 164, pp.95-106. https://doi.org/10.1016/j.agsy.2018.03.010

Liu, K., Ji, T., Li, M., Yang, X., Sun, J., Liu, H. and Liu, R., 2025. Prediction of greenhouse temperature and humidity across growing seasons: Hybridization of process-based model and deep neural networks. *Information Processing in Agriculture*, 13(2), pp. 238-252. https://doi.org/10.1016/j.inpa.2025.11.002

Muñoz-Liesa, J., Royapoor, M., Cuerva, E., Gassó-Domingo, S., Gabarrell, X. and Josa, A., 2022. Building-integrated greenhouses raise energy co-benefits through active ventilation systems. *Building and environment*, 208, p.108585. https://doi.org/10.1016/j.buildenv.2021.108585

Ouammi, A., Achour, Y., Zejli, D. and Dagdougui, H., 2019. Supervisory model predictive control for optimal energy management of networked smart greenhouses integrated microgrid. *IEEE transactions on automation science and engineering*, 17(1), pp.117-128.10.1109/TASE.2019.2910756

Ouammi, A. and Zejli, D., 2022. Centralized controller for the optimal operation of a cooperative cluster of connected microgrids powered multi-greenhouses. *Sustainable Computing: Informatics and Systems*, 33, p.100641. https://doi.org/10.1016/j.suscom.2021.100641

Rahman, H., Shah, U.M., Riaz, S.M., Kifayat, K., Moqurrab, S.A. and Yoo, J., 2024. Digital twin framework for smart greenhouse management using next-gen mobile networks and machine learning. *Future Generation Computer Systems*, 156, pp.285-300. https://doi.org/10.1016/j.future.2024.03.023

Ren, Z., Xu, M., Zhao, Z., Wang, X., Guo, J. and Dong, Y., 2026. A two-level hierarchical optimization framework for grid-connected photovoltaic-wind-battery systems in greenhouse energy management. *Sustainable Energy, Grids and Networks*, p.102181. https://doi.org/10.1016/j.segan.2026.102181

Rezaei, E., Dagdougui, H. and Ojand, K., 2022. Hierarchical distributed energy management framework for multiple greenhouses considering demand response. *IEEE Transactions on Sustainable Energy*, 14(1), pp.453-464. 10.1109/TSTE.2022.3215686

Sardouei-Nasab, S., Aminizadeh, S. and Zeinali Pour, N., 2025. Greenhouse evaluation of tomato cultivars across ripening stages for selection of superior breeding genotypes. *Greenhouse Plant Production Journal*, *2*(3), pp.40-50. https://doi.org/10.61882/gppj.2.3.40

Shojaei-Nia, M., Soufi, H., Jabbari, M. and Mousavi-Fard, S., 2026. Stress-Aware Greenhouse Climate Control: Integrating Plant Physiology, Intelligent Environmental Control, and Energy Management for Sustainable Protected Horticulture. *Greenhouse Plant Production Journal*, *3*(2), pp.62-81. https://doi.org/10.66224/gppj.3.2.62

Zhang, L., Zhao, Y., Ning, J. and Yao, Y., 2026. Multi-Horizon Machine Learning–Driven Predictive Control for Energy-Efficient Thermal Management of Rooftop Building-Integrated Greenhouses. *Energy*, p.141538. https://doi.org/10.1016/j.energy.2026.141538