SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Comfort-Gated Supervisory Control of an Office HVAC Zone in a Warm–Humid Climate: Gradient-Boosted Surrogates with Grid Search and Particle Swarm Optimization

Manisha Kishor Bhole[1*], Pankaj Hiraman Zope[2], Punam Jagdish Patil[3]

[1*]Department of Instrumentation Engineering, Bharati Vidyapeeth College of Engineering, Navi Mumbai, Maharashtra, 400614, India. E-mail: manisha.bhole@bvcoenm.edu.in
[2]Department of Computer Engineering, SSBT's College of Engineering and Technology, Bambhori, Jalgaon, Maharashtra, 425001, India
[3]Department of Instrumentation Engineering, Bharati Vidyapeeth College of Engineering, Navi Mumbai, Maharashtra, 400614, India.

ABSTRACT
HVAC plant dominates commercial building electricity demand, and in warm–humid climates the latent component is large enough to overturn control strategies designed around temperature alone. This paper presents a supervisory control framework in which three gradient-boosted-tree surrogates — zone temperature, relative humidity and hourly energy — are trained on one year of measured building-management-system data augmented by a calibrated psychrometric plant, and then serve as both the optimisation model and the evaluation plant for four supervisory controllers. Two perform exhaustive per-hour grid search over a discrete setpoint–dehumidifier action space; two tune a four-parameter operating schedule by particle swarm optimisation (PSO). Each appears in a night-setback and a no-setback variant under identical hardware-realism constraints: integer setpoints, a two-hour minimum compressor hold, occupancy-based operation and pre-cooling look-ahead. Selection is not made on energy — a three-part comfort gate (occupied-hour temperature compliance, humidity compliance, Fanger-based satisfaction) must be satisfied before a controller is eligible, and only the highest-saving eligible controller is locked and serialised.
Over a full 2025 evaluation year for an office zone, the surrogates reproduced measured dynamics at R^2 = 0.998, 0.994 and 0.999. The grid-search controllers achieved the largest nominal savings, 19.0% and 19.8%, but both failed the humidity gate at 79.2% and 83.4% compliance and were rejected. The PSO schedule with night setback was locked at 6.0% annual saving 100% temperature compliance, 96.6% humidity compliance, 96.6% of occupied hours inside the ASHRAE comfort box and no occupied hour above PPD 20; cumulative occupied discomfort fell by ~58%.

Keywords: HVAC supervisory control; gradient boosting; XGBoost; particle swarm optimisation; grid search; ASHRAE 55; latent load; night setback; surrogate model

1. Introduction

Buildings account for roughly 40% of global final energy use, and HVAC plant is consistently their largest end use at 40–60% of building demand [1], [2]. Growth is concentrated in warm economies: buildings drove close to 60% of the increase in global electricity consumption in 2024, largely through space cooling in India and China [3]. In India, buildings consume about 35% of national electricity, growing near 8% annually, which motivated progressive mandating of the Energy Conservation Building Code [4]. Supervisory optimisation of existing plant is attractive because a thermostat schedule, a dehumidifier interlock and an occupancy policy are software; changing them costs only engineering.

The latent-load problem. Most HVAC optimisation reports savings against a temperature-compliance constraint. In temperate or hot–dry climates this is defensible because sensible load dominates; in warm–humid climates it is not. Once the thermostat is satisfied the coil stops, dehumidification stops with it, and the envelope keeps admitting moisture on vapour-pressure difference alone. The latent component of ventilation air alone can approach half of total building load, and the minimum outdoor-air rates mandated for acceptable air quality [5] guarantee moisture keeps arriving whether or not cooling is called; setting back humidity as well as temperature yields substantially larger demand-response flexibility than temperature setback alone [6]. The consequence is specific: if the measured baseline ran without active dehumidification — the normal case for a thermostat-controlled office — any controller that switches a dehumidifier on is charged for energy the baseline never spent, and whether the comparison shows a saving depends almost entirely on the assumed incremental coefficient. §7 finds this dominates the headline result.

2. Literature Survey

Reviews place tree ensembles consistently strong on tabular building data [7], and extreme gradient boosting [8], [9] now routinely outperforms support-vector, random-forest [10] and shallow neural approaches while remaining fast enough to call inside an optimisation loop and interpretable enough to expose which drivers matter [11], [23]. Control divides into three families. MPC is the benchmark [12], [13], [24], with the recurring finding that it works when a sufficiently accurate and sufficiently simple model exists — the binding condition. Deep reinforcement learning removes that constraint [14], [15], [27]–[29] but is sample-hungry, sensitive to reward shaping and hard to certify [16], [17]. Metaheuristics avoid the model-form restriction at the cost of optimality guarantees: PSO [18], [19], [22] has been coupled to data-driven plant models for two decades, typically reporting 7–12% savings [20], [21], [25]. Exhaustive grid search is usually dismissed as a hyperparameter baseline [26], but has an under-appreciated virtue here — when the action space is small and discrete, per-step exhaustive evaluation is one-step-optimal MPC, making our 6 × 2 grid a legitimate method rather than a strawman. ASHRAE 55 [30] and ISO 7730 [31] define acceptable environments through Fanger's PMV/PPD framework [32], but comfort commonly enters control papers as a weighted penalty, making the accepted discomfort a function of a weight the reader cannot interpret, and humidity is more often omitted than temperature. The present work is distinguished less by its algorithms than by its evaluation protocol (Table 1).

Table 1. Positioning against representative prior work.

Study	Plant model	Control	Comfort treatment	RH	Outcome
Kusiak et al. [20]	Data-mined MLP ensemble	PSO on setpoints	Temperature constraint	No	>7% reduction
Wang et al. [11]	SVM, ANN, XGBoost, LightGBM	MPC (demand response)	Weighted deviation	No	21.6% cost cut, +0.10 K
Wei et al. [27]	EnergyPlus co-simulation	Deep Q-network	Reward penalty	No	Savings vs rule-based
Yu et al. [28]	Building simulator	Multi-agent DRL	Reward penalty	No	Multi-zone coordination
This work	**Boosted surrogates (T, RH, E), measured + physics-augmented**	**Grid search and PSO, each ± setback**	**Hard three-axis gate before ranking**	**Yes**	**19.8% best rejected on RH; 6.0% locked at full compliance**

3. CASE STUDY AND DATA

A single conditioned office zone (Z1) in warm–humid coastal, with a monsoon from June to September during which outdoor RH routinely exceeds 80% while dry-bulb temperature falls only modestly. This is the difficult combination for a cooling-only system: sensible load moderates while latent load does not, so a thermostat adequate in the dry season delivers a cool but damp zone in the wet season. The zone runs a weekday 09:00–18:00 schedule, unoccupied at weekends and on enumerated Indian public holidays, with peak occupancy ~40 persons. Energyplus simulated data is obtained for the year 2024 and 2025.

Two full years of hourly BMS records were used: 2024 (8 784 h) for development, 2025 (8 760 h) held out entirely — no 2025 record is seen during surrogate training or PSO tuning. Channels are timestamp, outdoor temperature and RH, occupancy count, thermostat setpoint, zone temperature, zone RH and zone HVAC energy; derived fields are hour, day-of-week, holiday flag and occupied-schedule flag, with cyclical time encoded as $\sin(2\pi h/24)$ and $\cos(2\pi h/24)$ so hours 23 and 0 are adjacent. The 2025 baseline — the "traditional" reference throughout — consumed 53 556 kWh on thermostat control without active dehumidification. Comfort follows ASHRAE 55 [30]: 23–26 °C centred at 24.5 °C, and 30–60% RH centred at 50%. Satisfaction uses Fanger PMV/PPD [31], [32].

4. SURROGATE PLANT

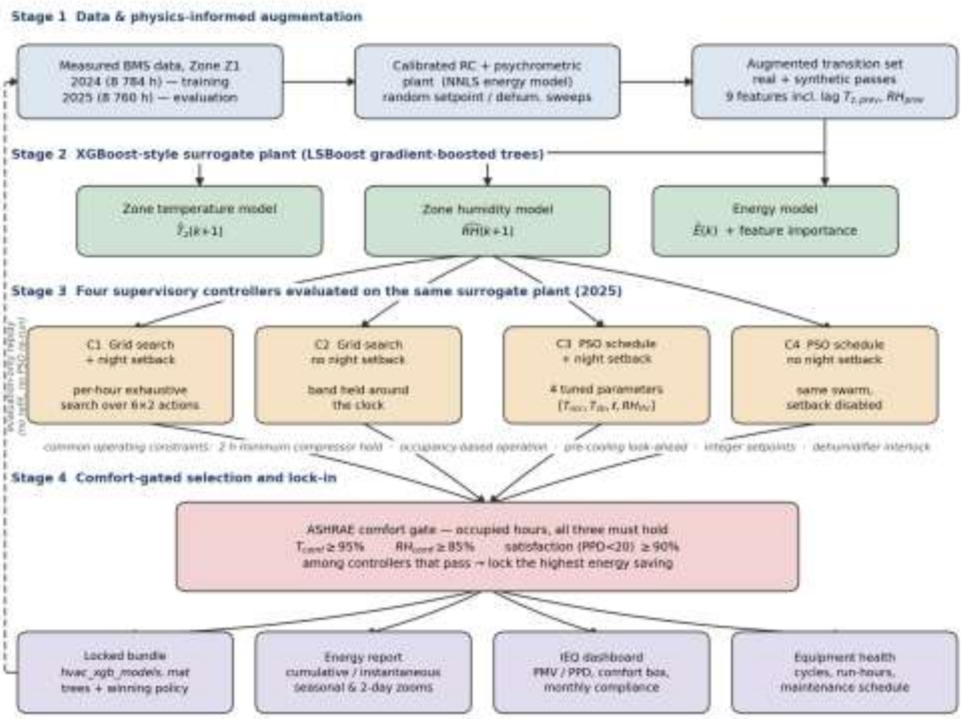

Figure 1. Framework architecture. Measured and physics-augmented transitions train three gradient-boosted surrogates, which serve simultaneously as the optimisation model and the evaluation plant for four supervisory controllers. A three-axis ASHRAE comfort gate filters the controllers before energy ranking, and the winner is serialised for evaluation-only replay..

The controllers need a plant that answers, thousands of times per simulated year: given zone state, weather, occupancy and a candidate action, what are temperature, humidity and energy next hour? Three properties are required — accuracy on measured data, validity across the whole action space including combinations the building never ran, and speed, since the exhaustive controller alone needs ~10^5 evaluations per annual simulation. Gradient-boosted regression trees satisfy all three: they handle the mixed continuous–integer–binary feature vector without scaling, capture threshold-like compressor behaviour without hand-specified bases, infer in microseconds, and expose predictor importance. We use MATLAB fitrensemble with LSBoost and shallow learners, the same least-squares formulation as XGBoost [8], [9].

4.1 Physics-based augmentation

Measured data alone cannot support a controller: the building only ever ran at the setpoints its operators chose, and the dehumidifier never ran at all, so training on that record and then permitting unseen actions is extrapolation disguised as optimisation. We therefore combine real transitions (each measured 2024 hour, dehumidifier channel zero) with synthetic transitions from a calibrated one-step physical plant swept across the full action space. The zone drifts toward outdoor conditions and accumulates occupancy gains before the coil pulls it toward setpoint,

$$T_z^d = T_z + K_{out}(T_{out} - T_z) + K_{occ}\, Occ/40\,, \qquad T_z(k+1) = \begin{cases} T_z^d - K_c(T_z^d - T_{set}), & T_z^d > T_{set} \\ T_z^d, & \text{otherwise} \end{cases}$$

the second branch enforcing that the system is cooling-only. Humidity drifts toward outdoor conditions with an occupancy source, is reduced by coil condensation whenever the compressor runs and further when the dehumidifier is active, then clipped to [25, 95]%. Energy uses a non-negative least-squares fit of a physically structured model to 2024, with humidity ratio from the standard psychrometric relation [33],

$$w(T, RH) = \frac{0.62198\,(RH/100)\,p_{ws}}{101.325 - (RH/100)\,p_{ws}}, \qquad p_{ws}(T) = 0.61078\exp\left(\frac{17.27\,T}{T + 237.3}\right)$$

on the basis [1, Occ, max(0, $T_{out}-T_{set}$), max(0, $T_{out}-T_z$), max(0, $w_{out}-w_{set}$)·1000] with w_{set} at 24 °C / 50%; non-negativity prevents physically absurd coefficient signs. The dehumidifier adds an incremental term proportional to the humidity-ratio gap.

The nine-feature input vector is [$T_{z,prev}$, RH_{prev}, T_{out}, RH_{out}, Occ, T_{set}, d, sin(2πh/24), cos(2πh/24)], predicting $T_z(k+1)$, $RH(k+1)$ and $E(k)$ with three independent models. The two lag features give thermal-mass memory; without them the model cannot represent the coasting behaviour that makes pre-cooling and setback worthwhile. The combined set is split 75/25, fitted with shallow trees (max 64 splits, min leaf 5), energy clamped at zero after inference, seed fixed at 6006. Shallow learners are deliberate: the surrogate is used extrapolatively inside a controller, and a model that interpolates superbly but oscillates between grid points would produce a controller that exploits its artefacts.

All four act on the same discrete space: an integer setpoint from {22…27} °C and a binary dehumidifier state — twelve actions per hour. Integer setpoints are used because that is what a commercial digital thermostat accepts; a continuous setpoint would report savings the hardware cannot realise. Four constraints apply identically to every controller, so differences are attributable to search method and setback policy rather than differing realism: a two-hour minimum compressor hold (without which cycle counts are fiction); occupancy-based operation; a pre-cooling look-ahead, exploiting thermal mass and suppressing the morning spike; and a dehumidifier interlock making dehumidification an explicit action with an explicit cost.

Family A — exhaustive grid search. Each hour, all twelve candidates are evaluated through the surrogate in one batched prediction and the minimum-cost action selected. With band violations $v_T = \max(0, \hat{T}_z - T_{hi}) + \max(0, T_{lo} - \hat{T}_z)$ and v_{RH} defined analogously, the cost of $a = (T_{set}, d)$ is

$$J(a) = w_E \hat{E} + \Phi(a) + w_{sw}\left|T_{set} - T_{set}^{prev}\right| + w_{cyc}\mathbb{1}[c \neq c^{prev}] + w_d\, d$$

$$\Phi = \begin{cases} w_T v_T^2 + w_{RH}(v_{RH}/5)^2 + w_{ctr}\left(\hat{T}_z - T_{ctr}\right)^2 & \text{occupied} \\ 0.6\, w_T \max\left(0, \hat{T}_z - T_{hi}\right)^2 & \text{pre-cooling} \\ w_T v_T^2 + w_{RH}(v_{RH}/5)^2 & \text{unoccupied, no setback} \\ w_{T,unocc} v_T^2 & \text{unoccupied, setback} \end{cases}$$

The switching and cycling penalties keep the search from producing a trajectory optimal hour-by-hour but mechanically absurd; humidity violation is scaled by 5 before squaring so a point of RH error is not weighted as heavily as a degree. The last two branches are the entire difference between setback variants: when the unoccupied comfort weight collapses to wT,unocc, energy dominates and the optimiser selects a high setpoint by itself — night setback emerges from the cost function rather than being scheduled. This controller is exact one-step MPC: no convergence risk, no search tuning, but myopic and by far the most expensive of the four.

4.3 Family B — PSO-tuned schedule

The rule is $\theta = [T_{occ}, T_{sb}, \ell, RH_{thr}]$: occupied setpoint, setback setpoint, pre-cool lead and dehumidifier RH threshold. When occupied, hold Tocc and run the dehumidifier if $RH > RH_{thr}$; when occupancy arrives within ℓ hours, pre-cool toward T_{occ}; when empty, hold T_{sb} if setback is enabled or T_{occ} if not, relaxing the threshold by five points. Optimisation uses a standard inertia-weighted swarm [18], [19] (inertia 0.7, $c_1 = c_2 = 1.5$) implemented directly, bounded to $T_{occ} \in [23, 25]$ °C, $T_{sb} \in [26, 28]$ °C, $\ell \in [0, 3]$ h, $RH_{thr} \in [50, 62]$%, minimising rolled through the same surrogate used for scoring, over representative weeks for tractability.

$$C(\theta) = \sum_k \hat{E}_k + 2000\, n_{viol}/\max(n_{occ}, 1)$$

A violation is counted when an occupied hour falls outside the band inset by a margin of 0.5 °C and 2 points. This margin matters: without it the swarm converges to solutions sitting exactly on the band edge, scoring perfectly against a soft objective but failing strict compliance on the smallest perturbation. The Tocc ceiling is likewise capped below the band top, since an unconstrained swarm always pushes the occupied setpoint wherever the energy gradient points.

4.4 Night-setback factorial.

Setback is standard practice and normally assumed beneficial, but in a warm–humid climate the unoccupied hours are precisely when moisture ingress is unopposed. Each method therefore runs twice, with setback enabled and disabled; the no-setback variants hold the band around the clock and are expected to save less, the point being to make setback's contribution explicit rather than assumed.

4.5 Comfort gate and lock-in

A controller is eligible only if, over occupied hours, $T_{comf} \geq 95\%$, $RH_{comf} \geq 85\%$ and satisfaction (PPD < 20) ≥ 90%. Among eligible controllers the highest saving is locked. If none is eligible the procedure does not fall back to maximum savings; it locks the most comfortable controller and reports which basis was used — a result obtained under the fallback branch means the comfort target was met by nobody and must be presented as such. The locked controller, surrogates and tuned parameters are serialised to one bundle, and a separate script reloads it and regenerates every metric without refitting or re-running the swarm, so the result reproduces from the artefact rather than merely from the code. Metrics are RMSE, R^2 and MAE for the surrogates; and annual saving, occupied T/RH/both-in-box compliance, mean PPD, satisfaction, compressor cycles and run-hours for control. Cycles and run-hours are included because a controller that saves energy by cycling harder has transferred cost from the electricity bill to the maintenance budget, and a study reporting only kWh will not notice.

5. Results

Configuration C throughout; surrogates trained on 2024 plus augmented transitions, every control result evaluated on the untouched 2025 year against the measured 53 556 kWh baseline.

5.1 Surrogate accuracy

Table 3. Surrogate accuracy on real 2025 transitions (Config C).

Target	RMSE	R^2	MAE	Units
Zone temperature	0.037	0.9982	0.024	°C
Zone relative humidity	0.955	0.9938	0.558	%
Zone energy	0.166	0.9993	0.098	kWh

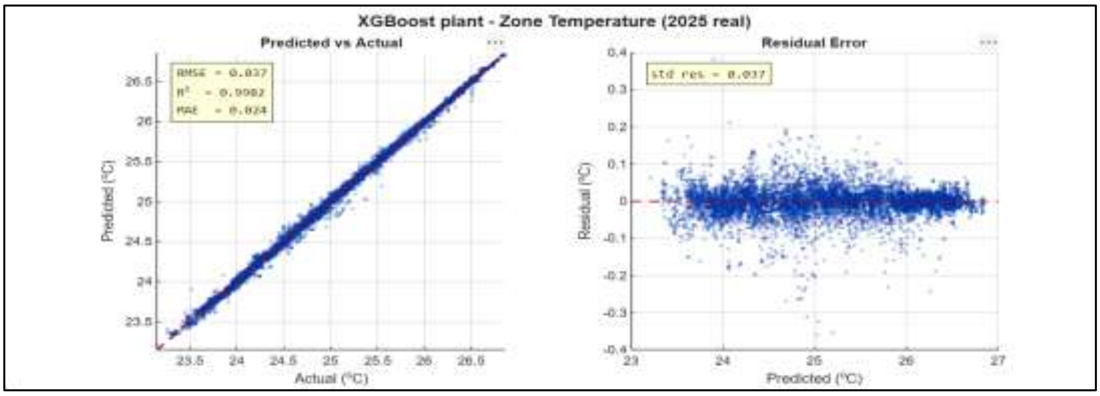

Figure 2. Zone temperature surrogate on real 2025 transitions: predicted versus actual (left), residuals (right).

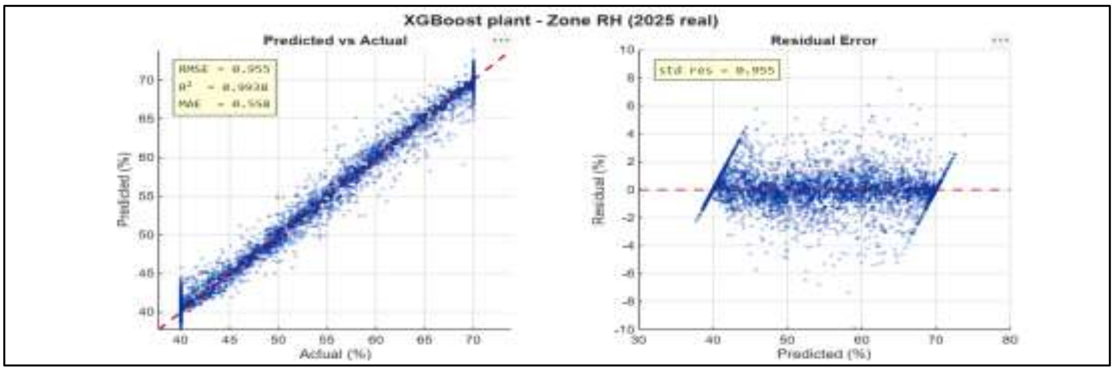

Figure 3. Zone relative humidity surrogate on real 2025 transitions.

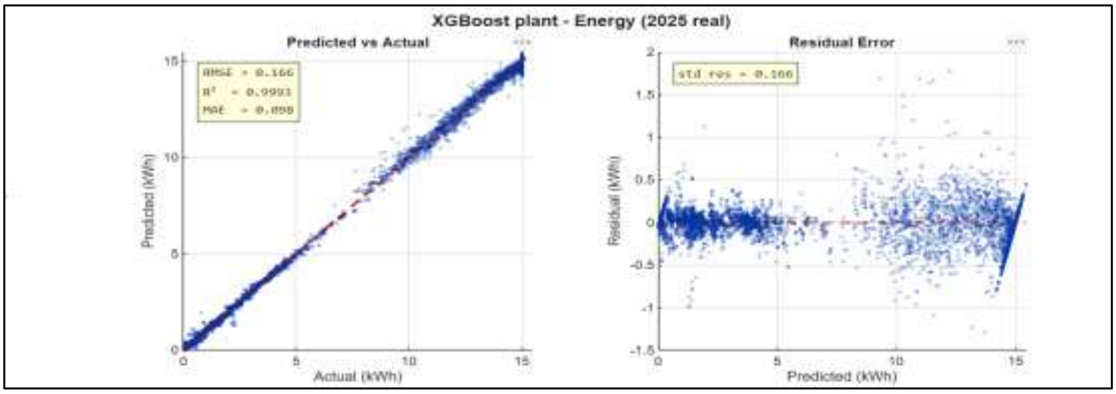

Figure 4. Zone energy surrogate on real 2025 transitions.

Temperature is predicted to within hundredths of a degree, inside sensor tolerance; humidity carries about one percentage point of RMSE; energy reaches R^2 = 0.999 with MAE below 0.1 kWh against hourly values reaching 15 kWh. Two residual patterns are structural rather than random. In Figure 3 residuals fan out at the extremes of the predicted range, an artefact of the clipping that keeps humidity physical — predictions pressed against a clip cannot distribute error symmetrically. In Figure 4 residuals are heteroscedastic, tightly banded at low demand and dispersed above ~8 kWh, with a vertical error band where the baseline saturates at maximum hourly draw. Neither is a defect here, since the control loop is sensitive to relative differences between candidate actions within an hour and those survive even where the absolute level carries error; but reported absolute savings inherit an uncertainty of order the energy MAE.

5.2 Four-controller comparison

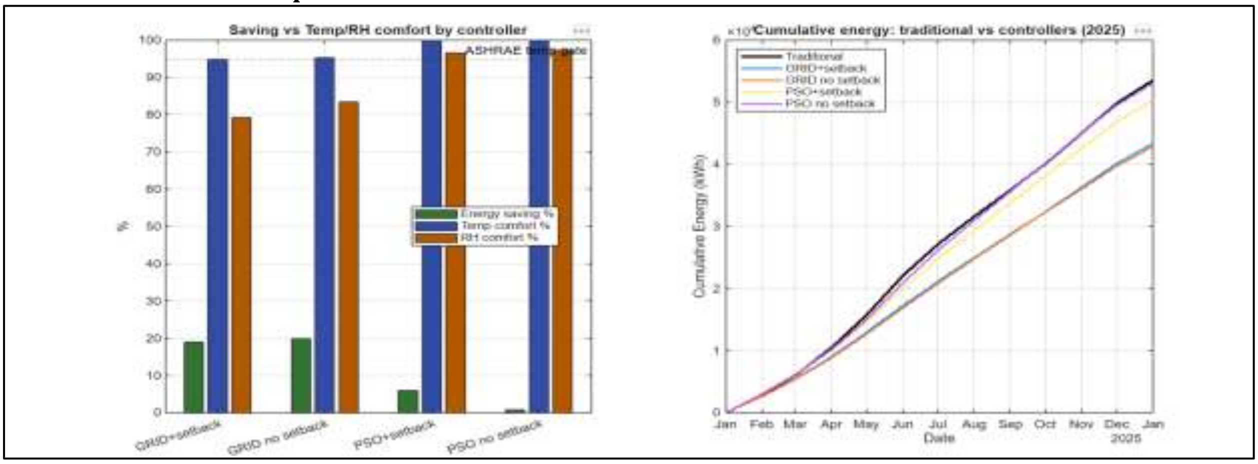

Figure 5. Controller comparison, 2025. Left: annual saving and occupied-hour temperature and humidity compliance for all four, with the 95% temperature gate marked. Right: cumulative annual energy, baseline and all four controllers.

Table 4. Four-controller comparison, 2025 (Config C). Bold entries fail the gate.

Controller	Saving (%)	T compl. (%)	RH compl. (%)	Gate	Verdict
C1 Grid + setback	19.0	**94.9**	**79.2**	FAIL (T, RH)	reject
C2 Grid − setback	19.8	95.4	**83.4**	FAIL (RH)	reject
C3 PSO + setback	6.0	100.0	96.6	PASS	**DEPLOY**
C4 PSO − setback	0.8	100.0	100.0	PASS	consider

This is the inversion. Ranked on energy alone the grid-search controllers win decisively — 19–20% against 6% — and a paper reporting only kWh would present C2 as the contribution. Both are rejected: C2 holds humidity in band for 83.4% of occupied hours against a gate of 85%, and C1 misses both the humidity gate and, marginally, the temperature gate at 94.9%.

The mechanism is visible in the cost function. Grid search selects, each hour, the action minimising a weighted sum in which humidity violation is scaled down relative to temperature while energy carries nonzero weight. Running the dehumidifier incurs immediate energy cost and an explicit penalty, while its humidity benefit is discounted; hour by hour the arithmetic favours leaving it off, and across the monsoon the shortfall accumulates to 1.6 points below the gate. Nothing has malfunctioned — the controller does exactly what it was told, and the weights encode a humidity preference the designer did not consciously choose. This is precisely what a hard gate exists to catch, and it is the argument against expressing comfort as a penalty term alone. The PSO controllers behave differently because they optimise a schedule under a margin-inset constraint rather than an hourly trade-off: C3 clears both gates comfortably at 6.0% saving, while C4 achieves perfect compliance but saves only 0.8%, since holding the band in an empty building overnight consumes very nearly what it saves. The night-setback contrast is therefore asymmetric across families, the most operationally useful reading of Table 4. For PSO, setback is worth 5.2 points of energy (6.0% vs 0.8%) for 3.4 points of humidity compliance (96.6% vs 100%) — an excellent trade, both variants remaining eligible. For grid search the difference is small and slightly negative (19.0% vs 19.8%), because the setback branch relaxes the unoccupied comfort weight and permits overnight drift that must then be reconditioned. The intuition that night setback is universally beneficial is not supported; its value depends on what the controller does with the freed degrees of freedom.

5.3 IEQ and equipment health

Table 5. IEQ and equipment metrics, occupied hours, 2025.

Controller	Satisfaction PPD<20 (%)	Mean PPD (%)	Compressor cycles	Run-hours
C1 Grid + setback	98.9	6.2	≈ 1	≈ 8 720
C2 Grid – setback	100.0	5.6	≈ 10	≈ 8 700
C3 PSO + setback	100.0	7.8	≈ 1 120	≈ 6 800
C4 PSO – setback	100.0	5.2	≈ 380	≈ 8 100

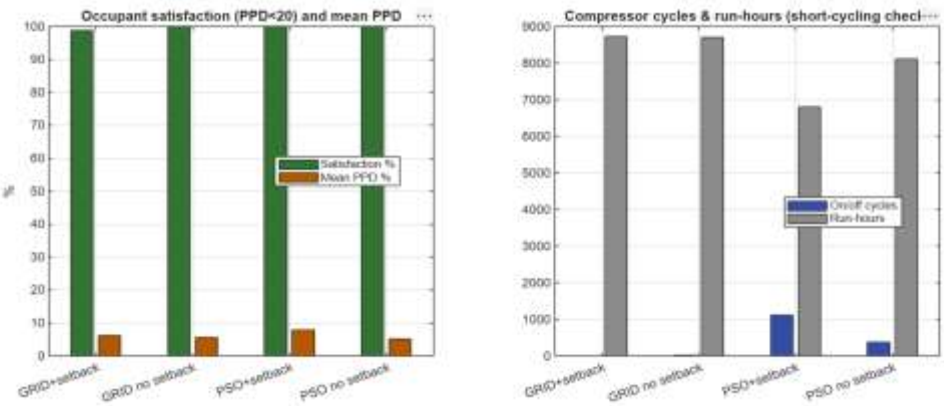

Figure 6. IEQ and hardware comparison. Left: occupant satisfaction and mean PPD. Right: compressor cycles and run-hours — the short-cycling check.

Satisfaction is essentially saturated for all four, mean PPD between 5 and 8%. This is expected: PPD is dominated by temperature, every controller holds temperature well, and satisfaction therefore fails to discriminate. The metric that does discriminate is humidity compliance — which is why it is gated separately rather than folded into a satisfaction score. A study reporting only PPD would conclude all four are equivalent on comfort, and would be wrong.

The equipment panel is more informative and less comfortable. The grid-search controllers run the compressor almost continuously, approaching 8 700 h with single-digit cycle counts; the PSO controllers run far fewer hours — 6 800 for C3 — but cycle far more, ~1 120 transitions against ~1 for C1. Every transition respects the two-hour hold, so none is short cycling, and 1 120 cycles across 8 760 h sits well inside relay duty. The trade is explicit: the locked controller earns its saving substantially by switching off when the zone is empty, recovering ~1 900 run-hours, and pays in switching events — concealed if expressed only in kWh. Condition-based triggers (air filter every 720 run-h, coil every 2 000, refrigerant every 4 000, relay every 15 000 cycles) give approximately 9, 3, 2 and fewer than 1 service events per year.

5.4 Locked controller: annual and seasonal performance

The locked controller is C3, PSO with night setback: 53 556 → 50 325 kWh, a saving of 3231 kWh or 6.0%.

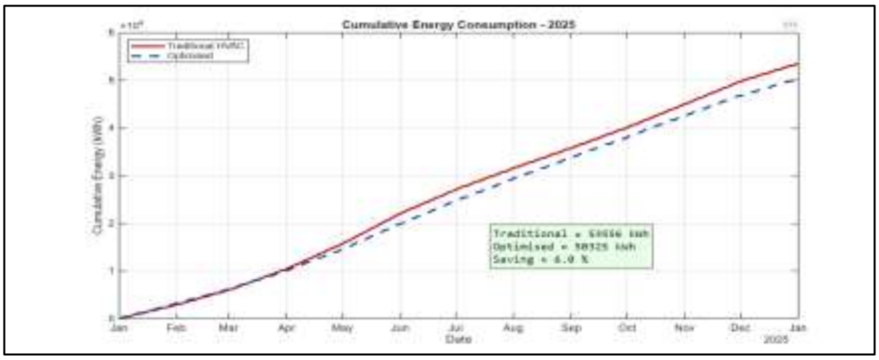

Figure 7. Cumulative 2025 energy, baseline against the locked controller.

5.4 Comfort attainment

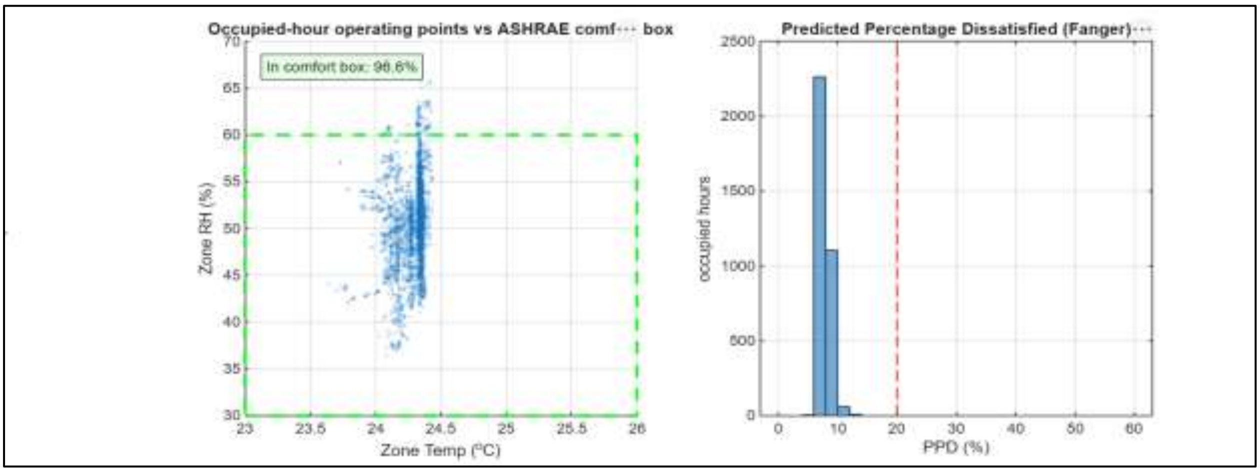

Figure 8. Locked controller, occupied hours. Left: operating points against the ASHRAE comfort box. Right: PPD distribution with the PPD = 20 threshold marked.

Figure 8 is the compliance evidence: occupied-hour operating points form a tight vertical cluster at 24.2–24.5 °C spanning roughly 40–63% RH, 96.6% inside the comfort box, with the entire PPD distribution below 13%. The narrowness in temperature against the spread in humidity is the visual statement of the whole paper — near-total authority over temperature, partial authority over humidity.

6. Discussion And Limitations

The energy-optimal controller was not the deployable one. Both grid-search controllers beat the winner by more than threefold on energy and were rejected on humidity; had comfort been a penalty inside the objective — common practice — this study would have reported 19.8% savings with temperature compliance above 95%, and the humidity shortfall would never have surfaced. The two search families failed differently: grid search is exact but myopic, optimal for the next hour, and hour-by-hour declines to pay for dehumidification whose benefit accrues over days; PSO optimises a schedule under a constraint, cannot make that hourly trade, and therefore sustains humidity compliance at the cost of energy. The distinction is not that one algorithm is better but that the one-step horizon and the humidity time constant are mismatched — precisely the argument for MPC with a proper receding horizon, or for DRL, over greedy exhaustive search. tested rather than assumed, Night setback is not universally beneficial and should be particularly where unoccupied hours are when moisture ingress is unopposed.

Several threats to validity qualify these findings. All control results are obtained through the surrogate rather than the physical building: inter-controller comparisons are exact since all four are scored on identical plant, but absolute savings inherit surrogate error, and although an energy MAE of 0.098 kWh against ~53 000 kWh annually is small in expectation, the heteroscedastic residuals of Figure 4 mean a systematic component would not cancel — the 6.0% is a well-founded estimate, not a metered result. The surrogate learns dehumidifier and off-nominal setpoint behaviour from a calibrated physical model rather than measurement, so its fidelity in exactly the region the controller most exploits is inherited from that model; this is unavoidable without an instrumented campaign, and is why κ is free. The swarm budget is modest, and the Config B log showed best cost flat across all ten iterations, indicating premature convergence rather than a converged optimum — the PSO results should be read as a lower bound. PSO tuning uses representative weeks (final scoring uses all 8760 h, so no result is contaminated, but tuned parameters are optimal for the sampled weeks only). The Fanger calculation assumes fixed metabolic rate, clothing and air speed, whereas adaptive models would likely widen the acceptable band in an Indian office. Finally, one zone, one building and two years support no claim of generalisation: the framework transfers, the numbers do not.

For an operator, three findings are directly actionable: supervisory retuning of an existing schedule delivered 6% annual energy reduction while improving comfort, with no capital equipment assumed; the residual comfort deficiency is confined to July and August, converting a diffuse control problem into a specific capacity question; and savings here are purchased by switching plant off when the zone is empty, recovering ~1900 compressor run-hours, so maintenance should be re-based on run-hours rather than calendar time.

7. Conclusion

Three gradient-boosted surrogates trained on measured data augmented by a calibrated psychrometric plant reproduced zone temperature, humidity and energy on an unseen year at R^2 = 0.998, 0.994 and 0.999, and served as both optimisation model and evaluation plant for four supervisory controllers spanning grid search and PSO, each with and without night setback. The central result is that the two highest-saving controllers, at 19.0% and 19.8%, were both rejected by an ASHRAE humidity gate, while the locked controller — a PSO-tuned schedule with night setback — delivered 6.0% annual saving (53556 → 50325 kWh) at 100% temperature compliance, 96.6% humidity compliance, 96.6% of occupied hours inside the comfort box and no occupied hour above PPD 20, with cumulative occupied discomfort down ~58%. A sensitivity analysis established that the assumed dehumidifier energy coefficient moves the same controller between −12.3% and +6.0% — an asymmetry structural to any humid-climate study benchmarked against a less-capable measured baseline.

References

1. L. Pérez-Lombard, J. Ortiz, and C. Pout, “A review on buildings energy consumption information,” Energy and Buildings, vol. 40, no. 3, pp. 394–398, 2008.
2. International Energy Agency, The Future of Cooling. Paris: IEA, 2018.
3. REN21, Renewables 2025 Global Status Report — Buildings. Paris: REN21 Secretariat, 2025.
4. Bureau of Energy Efficiency, Energy Conservation Building Code 2017. New Delhi: Ministry of Power, Government of India, 2017.
5. ASHRAE, ANSI/ASHRAE Standard 62.1-2022: Ventilation and Acceptable Indoor Air Quality. Atlanta, GA: ASHRAE, 2022.
6. “A temperature and humidity setback demand response strategy for HVAC systems,” Sustainable Cities and Society, 2021. [Author list and volume to be verified.]
7. K. Amasyali and N. M. El-Gohary, “A review of data-driven building energy consumption prediction studies,” Renewable and Sustainable Energy Reviews, vol. 81, pp. 1192–1205, 2018.
8. T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in Proc. 22nd ACM SIGKDD (KDD ’16), San Francisco, CA, 2016, pp. 785–794.
9. J. H. Friedman, “Greedy function approximation: A gradient boosting machine,” The Annals of Statistics, vol. 29, no. 5, pp. 1189–1232, 2001.
10. L. Breiman, “Random forests,” Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
11. H. Wang, D. Mai, Q. Li, and Z. Ding, “Evaluating machine learning models for HVAC demand response: The impact of prediction accuracy on model predictive control performance,” Buildings, vol. 14, no. 7, 2212, 2024.
12. A. Afram and F. Janabi-Sharifi, “Theory and applications of HVAC control systems — A review of model predictive control (MPC),” Building and Environment, vol. 72, pp. 343–355, 2014.
13. Y. Yao and D. K. Shekhar, “State of the art review on model predictive control (MPC) in heating ventilation and air-conditioning (HVAC) field,” Building and Environment, vol. 200, 107952, 2021.
14. Y. Wang, K. Velswamy, and B. Huang, “A long-short term memory recurrent neural network based reinforcement learning controller for office heating ventilation and air conditioning systems,” Processes, vol. 5, no. 4, 46, 2017.
15. A. Manjavacas, A. Campoy-Nieves, J. Jiménez-Raboso, M. Molina-Solana, and J. Gómez-Romero, “An experimental evaluation of deep reinforcement learning algorithms for HVAC control,” Artificial Intelligence Review, vol. 57, no. 7, 173, 2024.
16. Z. Wang and T. Hong, “Reinforcement learning for building controls: The opportunities and challenges,” Applied Energy, vol. 269, 115036, 2020.
17. M. Han, R. May, X. Zhang, X. Wang, S. Pan, D. Yan, Y. Jin, and L. Xu, “A review of reinforcement learning methodologies for controlling occupant comfort in buildings,” Sustainable Cities and Society, vol. 51, 101748, 2019.
18. J. Kennedy and R. Eberhart, “Particle swarm optimization,” in Proc. IEEE Int. Conf. Neural Networks, Perth, Australia, 1995, pp. 1942–1948.
19. Y. Shi and R. Eberhart, “A modified particle swarm optimizer,” in Proc. IEEE Int. Conf. Evolutionary Computation, Anchorage, AK, 1998, pp. 69–73.
20. A. Kusiak, M. Li, and F. Tang, “Modeling and optimization of HVAC energy consumption,” Applied Energy, vol. 87, no. 10, pp. 3092–3102, 2010.
21. A. Kusiak and G. Xu, “Modeling and optimization of HVAC systems using a dynamic neural network,” Energy, vol. 42, no. 1, pp. 241–250, 2012.
22. R. Poli, J. Kennedy, and T. Blackwell, “Particle swarm optimization: An overview,” Swarm Intelligence, vol. 1, no. 1, pp. 33–57, 2007.
23. J. Mao, R. Grammenos, and K. Karagiannis, “Data analysis and interpretable machine learning for HVAC predictive control: A case-study based implementation,” Science and Technology for the Built Environment, vol. 29, no. 7, pp. 1–21, 2023.
24. J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, “All you need to know about model predictive control for buildings,” Annual Reviews in Control, vol. 50, pp. 190–232, 2020.

25. A.-T. Nguyen, S. Reiter, and P. Rigo, "A review on simulation-based optimization methods applied to building performance analysis," Applied Energy, vol. 113, pp. 1043–1058, 2014.
26. J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," Journal of Machine Learning Research, vol. 13, pp. 281–305, 2012.
27. T. Wei, Y. Wang, and Q. Zhu, "Deep reinforcement learning for building HVAC control," in Proc. 54th Design Automation Conference (DAC '17), Austin, TX, 2017, pp. 1–6.
28. L. Yu, Y. Sun, Z. Xu, C. Shen, D. Yue, T. Jiang, and X. Guan, "Multi-agent deep reinforcement learning for HVAC control in commercial buildings," IEEE Transactions on Smart Grid, vol. 12, no. 1, pp. 407–419, 2021.
29. G. Gao, J. Li, and Y. Wen, "DeepComfort: Energy-efficient thermal comfort control in buildings via reinforcement learning," IEEE Internet of Things Journal, vol. 7, no. 9, pp. 8472–8484, 2020.
30. ASHRAE, ANSI/ASHRAE Standard 55-2020: Thermal Environmental Conditions for Human Occupancy. Atlanta, GA: ASHRAE, 2020.
31. ISO, ISO 7730:2005 — Ergonomics of the Thermal Environment: Analytical Determination and Interpretation of Thermal Comfort Using PMV and PPD Indices. Geneva: ISO, 2005.
32. P. O. Fanger, Thermal Comfort: Analysis and Applications in Environmental Engineering. Copenhagen: Danish Technical Press, 1970.
33. ASHRAE, ASHRAE Handbook — Fundamentals, ch. 1: Psychrometrics. Atlanta, GA: ASHRAE, 2021.