Journal of Modern Energy Research
Vol. 7, Mar 2026, pp. 25–40
ISSN: 3078-3534, DOI: 10.56947/jmer.v7.2

Explainable Boosting Machines for Thermal Load Prediction and Enclosure Thermodynamics of Residential HVAC Condensing Units

Djilali Hamza

Department of Power Engineering, Canadian University Dubai, Dubai, United Arab Emirates

Article Info

Article history:
Received Dec 14, 2025
Revised Jan 28, 2026
Accepted Feb 18, 2026

Keywords:
interpretable machine learning; generalized additive models; explainable boosting machines; air-cooled condensers; hot-air recirculation; building energy codes; computational fluid dynamics; surrogate modelling

ABSTRACT

The architectural integration of heating, ventilation, and air-conditioning systems in contemporary residential layouts increasingly subordinates thermodynamic performance to visual aesthetics, resulting in the concealment of outdoor split-system condensing units within louvered cabinets, service ducts, and constrained balcony alcoves. Such enclosures throttle the condenser airflow and, more damagingly, permit the exhaust plume to recirculate into the intake face, elevating the on-coil temperature above ambient and driving the vapour-compression cycle toward a high-lift, low-efficiency regime. Existing predictive models of this degradation are either computationally prohibitive, in the case of three-dimensional conjugate simulation, or opaque, in the case of deep and ensemble learners, and neither form yields the transparent, auditable relationships that building codes require. This work develops a framework based on Explainable Boosting Machines, a bagged and cyclically boosted realisation of generalized additive models with pairwise interactions. We first derive a closed-form cycle-level degradation law in which the compressor power admits a simple pole in a single dimensionless recirculation-resistance index, and we prove that the resulting physics is exactly representable as a second-order functional ANOVA decomposition in latent coordinates. This establishes that an additive learner with pairwise terms is not merely an interpretable approximation but the correct hypothesis class for the problem. A coupled computational fluid dynamics and lumped-parameter cycle model generates 15,000 parametric observations spanning frontal, lateral, and overhead clearances, louver porosity, ambient temperature, and part-load ratio. The learned model attains a mean absolute percentage error of 2.1% on held-out data, matching gradient-boosted tree and neural baselines while remaining fully decomposable. The extracted shape functions localise a sharp threshold in frontal clearance near 0.45 m, below which the contribution to compressor power accelerates in a manner consistent with the analytically predicted pole, and the leading pairwise term quantifies the porosity compensation required when that clearance cannot be met. The resulting relationships are stated as candidate clauses for residential energy codes.

Corresponding Author:
Djilali Hamza
Department of Power Engineering
Canadian University Dubai
Dubai, United Arab Emirates
Email: djilali.hamza@cud.ac.ae

NOMENCLATURE

Symbol	Description	Unit
A_f	Condenser face area	m^2
a	Cycle sensitivity coefficient, $\dot{Q}_c/(\eta_{II}T_e)$	W K^{-1}
b	Recirculation–resistance index, $\Lambda/(\dot{m}c_p)$	K W^{-1}
C_d	Discharge coefficient of the louver aperture	–
c_p	Specific heat capacity of air at constant pressure	J kg^{-1} K^{-1}
D_h	Hydraulic diameter of the condenser discharge	m
d_f, d_s, d_t	Frontal, lateral and overhead clearance	m

Symbol	Description	Unit
f_j, f_{jk}	Main-effect and pairwise shape functions	W
K	Pressure-loss coefficient	–
$\dot{m}$	Condenser air mass flow rate	kg s^{-1}
$\mathrm{Nu}, \mathrm{Pr}, \mathrm{Re}$	Nusselt, Prandtl and Reynolds numbers	–
$\dot{Q}_c$	Delivered cooling capacity	W
$\dot{Q}_{rej}$	Condenser heat rejection rate	W
T_{amb}	Free-stream dry-bulb temperature	K
T_{in}	Surface-averaged on-coil intake temperature	K
T_{out}	Condenser discharge temperature	K
T_c, T_e	Condensing and evaporating saturation temperatures	K
UA	Overall conductance of the condenser	W K^{-1}
W	Compressor electrical power	W
Δp_{fan}	Fan static pressure rise	Pa
ε	Condenser effectiveness	–
η_{II}	Second-law efficiency of the cycle	–
χ	Recirculation fraction	–
Λ	Thermal penalty group, $\chi/(1-\chi)+1/\varepsilon$	–
ϕ	Louver porosity, open area fraction	–
Ξ	Runaway index, ab	–
$\delta_\bullet$	Normalised clearance, $d_\bullet/D_h$	–

1. INTRODUCTION

1.1. Motivation

Residential buildings account for a substantial and growing share of global final energy demand, and space cooling is the fastest-growing component of that share. In the residential sector the dominant cooling technology remains the air-cooled split system, in which an outdoor unit rejects condenser heat directly to the ambient air by forced convection. The thermodynamic performance of such a unit is therefore determined not only by its internal design but by the quality of the air stream it is permitted to draw and discharge. Load-forecasting techniques for related electrical demand, such as the attention-based bidirectional approach of Kamalov et al., target a different prediction problem, aggregate electrical load rather than instantaneous compressor power, but share the broader goal of anticipating thermally driven demand.

Contemporary architectural practice places that air stream under increasing pressure. Outdoor units are routinely concealed behind decorative louvered screens, recessed into service niches, placed within joinery cabinets on terraces, or installed in narrow high-rise re-entrant shafts where several units on successive floors share a single ventilation column. These interventions are motivated by visual continuity, acoustic attenuation, and municipal facade regulations rather than by thermal considerations, and they are frequently specified without any calculation of the resulting airflow field.

Two distinct penalties follow. The first is hydraulic: the enclosure adds a series pressure loss that shifts the fan operating point along its characteristic curve, reducing volumetric flow and hence the air-side convective conductance. The second, and by a considerable margin the more severe, is thermal: the warm discharge plume, unable to disperse, impinges on the enclosure wall, forms a wall jet, and returns to the intake face. The on-coil temperature then exceeds the free-stream dry-bulb temperature by a margin that grows without bound as the recirculating fraction approaches unity. The cycle responds by raising the condensing pressure, the compression ratio increases, and both the coefficient of performance and the delivered capacity fall. In the limiting case the high-pressure safety cutout trips and cooling is lost entirely at precisely the ambient condition where it is most needed, contributing to coincident peak demand on the electricity network.

1.2. Limitations of existing approaches

The engineering literature addresses this problem through resolved simulation. Three-dimensional computational fluid dynamics coupled to a refrigerant cycle model reproduces the recirculation field faithfully, and such studies have established the qualitative picture described above. They do not, however, generalise. Each simulation answers a question about one geometry at one operating condition, at a cost of hours to days of wall-clock time, and the resulting knowledge is not expressible as a rule that a designer can apply at the drawing stage or that a code official can enforce by inspection.

The building-energy machine learning literature offers the complementary failure. Deep neural networks, gradient-boosted trees, and random forests predict system energy consumption accurately from environmental and operational features, but the mapping they learn is not decomposable. A regulator cannot read a clearance requirement off a trained network. Post-hoc explanation methods do not close this gap in a manner adequate to normative use: local surrogate and additive-attribution methods provide neighbourhood-valid approximations whose fidelity to the underlying model is not guaranteed, they are unstable under the strong correlation that characterises geometric design variables, and different methods applied to the same model can rank features differently. When the output of an analysis is intended to become a numerical clause in a building code, an approximate and contestable explanation of an exact but hidden model is the wrong artefact. The appropriate response is an exact and inspectable model.

1.3. Contributions

This work makes four contributions.

First, we derive a cycle-level degradation law in closed form. By combining an intake mixing balance, an effectiveness–NTU condenser model, and a second-law compressor model, we show that the compressor power is a rational function of a single dimensionless index and possesses a simple pole. The threshold behaviour observed empirically in enclosure studies is thereby given an analytical origin rather than being reported as a fitted curiosity.

Second, we prove that this physics admits an exact second-order functional ANOVA representation in latent coordinates. Under a monotone transformation of the response, the model separates into additive terms in the cycle parameters and a single bivariate term in the recirculation fraction and mass flow rate. Consequently, the hypothesis class of generalized additive models with pairwise interactions is not an interpretability concession but is exactly matched to the structure of the target function. This result reframes the methodological choice from a trade-off into an identification.

Third, we construct a coupled simulation pipeline and a 15,000-point design over six variables, and we fit an Explainable Boosting Machine together with a set of black-box and additive baselines. Accuracy is quantified against those baselines, and shape-function uncertainty is quantified by outer bagging.

Fourth, we convert the extracted shape functions and interaction surfaces into candidate normative clauses: a critical frontal clearance, a porosity compensation rule for cases in which that clearance cannot be achieved, and a dimensionless screening criterion applicable to geometries outside the sampled domain.

1.4. Organisation

Section 2 reviews related work. Section 3 develops the thermodynamic formulation and states the analytical results. Section 4 presents the learning formulation and connects it to Section 3. Section 5 describes the numerical study. Section 6 reports results, Section 7 discusses their normative implications and the limitations of the study, and Section 8 concludes.

2. RELATED WORK

2.1. Obstructed and recirculating air-cooled condensers

Degradation of air-cooled condensers under obstructed intake conditions has been studied principally in the context of high-rise re-entrant shafts, rooftop plant arrays, and data-centre outdoor plant. The consistent finding is that on-coil temperature elevation, rather than volumetric flow reduction, dominates the efficiency loss, and that the elevation is strongly non-linear in the confinement scale. Studies of vertically stacked units in ventilation shafts report on-coil temperatures tens of kelvin above ambient at upper floors, with corresponding capacity loss and cutout events. Work on horizontally confined units, which is closer to the residential enclosure problem considered here, is comparatively sparse and is dominated by single-geometry case studies. Manufacturer literature specifies minimum clearances, but these values are stated without an accompanying performance model, are not differentiated by louver porosity, and are frequently disregarded in practice.

2.2. Data-driven modelling of HVAC energy use

Supervised learning has been applied extensively to building and HVAC energy prediction, with gradient-boosted trees and recurrent or convolutional networks reported as strong performers on metered consumption data. A comparative analysis of modern load forecasting techniques more generally, in settings not specific to condenser enclosures, is surveyed by Kamolov. Physics-informed and hybrid formulations improve extrapolation by embedding conservation constraints. Surrogate modelling of CFD outputs, using Gaussian processes, polynomial chaos, or neural operators, is well established in thermal engineering. What distinguishes the present problem is the intended use of the model: the output is not a forecast but a design rule, and this places interpretability, monotonicity, and threshold localisation ahead of raw predictive accuracy in the objective hierarchy.

2.3. Interpretable machine learning

Generalized additive models decompose a response into univariate shape functions and thereby retain exact interpretability. Their extension to include a sparse set of pairwise interactions recovers much of the accuracy of unrestricted ensembles while preserving visual decomposability, and Explainable Boosting Machines implement this class by cyclic gradient boosting over shallow trees with outer bagging. Comparable accuracy has been reported across tabular benchmarks and in high-stakes healthcare settings where the ability to inspect and correct individual shape functions proved decisive. Neural realisations of the same decomposition offer smoother shape functions at some cost in fitting speed and reproducibility. The argument that intrinsically interpretable models should be preferred to post-hoc explanation in high-stakes decisions applies with particular force to normative engineering, where the model output is enforced against third parties. Functional ANOVA provides the formal setting in which such decompositions are unique, subject to the identifiability constraints discussed in Section 4.3. A broader state-of-the-art review of feature-attribution methods in machine learning, spanning application domains well beyond the present enclosure problem, is given by Kamolov.

3. THERMODYNAMIC FORMULATION

3.1. Condenser energy balance

The outdoor unit is treated as a cross-flow fin-and-tube heat exchanger with a single-phase air stream and a condensing refrigerant stream. Neglecting casing losses and fan heat addition to the air stream, the steady heat rejection rate is

$$\dot{Q}_{rej} = \dot{m}\, c_p \left(T_{out} - T_{in}\right). \tag{1}$$

Because the refrigerant-side temperature is approximately constant over the two-phase region, the exchanger effectiveness reduces to the single-stream form

$$\varepsilon = 1 - \exp\left(-\mathrm{NTU}\right), \qquad \mathrm{NTU} = \frac{UA}{\dot{m}c_p}, \tag{2}$$

and the heat rejection may be written in terms of the condensing saturation temperature as

$$\dot{Q}_{rej} = \varepsilon\, \dot{m}c_p \left(T_c - T_{in}\right). \tag{3}$$

The overall conductance is air-side dominated. With the fin efficiency η_o and the total air-side area A_o,

$$UA \simeq \eta_o h A_o, \qquad h = \frac{\mathrm{Nu}\, k}{D_h}, \tag{4}$$

and the tube-bank Nusselt number follows a power-law correlation of the standard form

$$\mathrm{Nu} = C\, \mathrm{Re}^m\, \mathrm{Pr}^n \left(\frac{\mathrm{Pr}}{\mathrm{Pr}_w}\right)^{1/4}, \qquad \mathrm{Re} = \frac{\rho u_{max} D_h}{\mu}. \tag{5}$$

Since $u_{max} \propto \dot{m}$, equations (2), (4) and (5) render ε a monotone increasing function of $\dot{m}$ alone at fixed geometry, with $\varepsilon \to 0$ as $\dot{m} \to 0$ and $\varepsilon \to 1$ as $\dot{m} \to \infty$. We write this dependence as $\varepsilon(\dot{m})$ throughout.

3.2. Fan operating point under enclosure resistance

The mass flow rate is not a free parameter but is fixed by the intersection of the fan characteristic with the system resistance. Representing the axial fan by a quadratic characteristic and the enclosure by a lumped loss coefficient referred to the face velocity $u_f = \dot{m}/(\rho A_f)$,

$$\Delta p_{fan}(\dot{m}) = \Delta p_0 \left[1 - \left(\frac{\dot{m}}{\dot{m}_0}\right)^2\right], \qquad \Delta p_{sys}(\dot{m}) = \tfrac{1}{2}\rho u_f^2\, K_{tot}. \tag{6}$$

The total loss coefficient is decomposed into a coil term, a louver term, and a confinement term:

$$K_{tot} = K_{coil} + K_{lv}(\phi) + K_{cf}(\delta_f, \delta_s, \delta_t). \tag{7}$$

The louver contribution follows the perforated-plate form

$$K_{lv}(\phi) = \left(\frac{1}{C_d \phi} - 1\right)^2, \tag{8}$$

which diverges as $\phi \to 0$ and vanishes as $\phi \to 1$ with $C_d \to 1$. The confinement term captures the sudden expansion, impingement, and turning losses imposed by the surrounding surfaces and is represented by an inverse-power form in the normalised clearances,

$$K_{cf} = \sum_{\bullet \in \{f,s,t\}} c_\bullet\, \delta_\bullet^{-p_\bullet}, \qquad \delta_\bullet = \frac{d_\bullet}{D_h}, \tag{9}$$

with coefficients $c_\bullet$ and exponents $p_\bullet$ regressed from the simulation campaign of Section 5. Equating the two expressions in (6) yields the operating point in closed form,

$$\dot{m}^* = \dot{m}_0 \left[1 + \frac{K_{tot}\,\dot{m}_0^2}{2\rho A_f^2 \Delta p_0}\right]^{-1/2}, \tag{10}$$

establishing that $\dot{m}^*$ is strictly decreasing in K_{tot} and hence in every confinement measure.

3.3. Intake mixing and the recirculation fraction

Let $\chi \in [0,1)$ denote the fraction of the intake mass flow that originates from the unit's own discharge rather than from the free stream. An adiabatic mixing balance at the intake face gives $T_{in} = (1-\chi)T_{amb} + \chi T_{out}$. Substituting (1) and solving for T_{in},

$$T_{in} = T_{amb} + \frac{\chi}{1-\chi} \cdot \frac{\dot{Q}_{rej}}{\dot{m}c_p}. \tag{11}$$

Equation (11) is the central kinematic result of the formulation. The on-coil elevation is not linear in the recirculation fraction; it is governed by the factor $\chi/(1-\chi)$, which is negligible for $\chi \lesssim 0.15$, of order unity at $\chi = 0.5$, and unbounded as $\chi \to 1$. The physical content is that recirculated air re-enters carrying heat that must be rejected again, so that the enclosure sustains a positive feedback loop whose gain is χ.

Combining (3) and (11) eliminates T_{in} and yields the condensing temperature,

$$T_c = T_{amb} + \frac{\dot{Q}_{rej}}{\dot{m}c_p}\Lambda, \qquad \Lambda \equiv \frac{\chi}{1-\chi} + \frac{1}{\varepsilon(\dot{m})}. \tag{12}$$

The group Λ collects the two mechanisms of degradation in a single additive form: the first term is the recirculation penalty, the second the finite-conductance penalty. Both increase as the enclosure tightens, the first because χ rises and the second because $\dot{m}$ falls.

3.4. Cycle closure and the degradation law

The vapour-compression cycle is closed by a second-law model in which the coefficient of performance is a fixed fraction of the reversible value between the saturation temperatures,

$$\mathrm{COP} = \frac{\dot{Q}_c}{W} = \eta_{II}\frac{T_e}{T_c - T_e}, \tag{13}$$

with η_{II} absorbing isentropic, volumetric, motor, and superheat losses. Writing $a = \dot{Q}_c/(\eta_{II}T_e)$ and $b = \Lambda/(\dot{m}c_p)$, and applying the overall balance $\dot{Q}_{rej} = \dot{Q}_c + W$, equations (12) and (13) give

$$W = a\left[(T_{amb} - T_e) + b\left(\dot{Q}_c + W\right)\right],$$

which is linear in W and admits the solution

$$W(b) = \frac{a\left[\Delta T + b\,\dot{Q}_c\right]}{1 - ab}, \qquad \Delta T = T_{amb} - T_e > 0. \tag{14}$$

Equation (14) is the degradation law. All enclosure geometry enters through the single scalar b, and all cycle and load information enters through a, ΔT and $\dot{Q}_c$. We refer to $\Xi = ab$ as the runaway index.

The following two statements characterise (14). Throughout, $a, \dot{Q}_c, \Delta T > 0$ are held fixed and b ranges over $[0, a^{-1})$.

Proposition 1. *W is strictly increasing and strictly convex on $[0, a^{-1})$, and $W(b) \to \infty$ as $b \uparrow a^{-1}$.*

Proof. Differentiating (14),

$$\frac{dW}{db} = \frac{a\left[\dot{Q}_c(1-ab) + a\left(\Delta T + b\dot{Q}_c\right)\right]}{(1-ab)^2} = \frac{a\left(\dot{Q}_c + a\Delta T\right)}{(1-ab)^2},$$

which is positive on the stated interval. Differentiating once more,

$$\frac{d^2W}{db^2} = \frac{2a^2\left(\dot{Q}_c + a\Delta T\right)}{(1-ab)^3} > 0 \quad \text{for } b < a^{-1}.$$

Divergence is immediate from the vanishing denominator of (14) and the positivity of its numerator at $b = a^{-1}$. □

Proposition 1 supplies the analytical origin of the threshold behaviour reported empirically in Section 6. Because b is itself increasing and unbounded as the frontal clearance shrinks, through both $\chi \to 1$ in Λ and $\dot{m} \to 0$ in (10), the composition of a monotone divergent map with a convex divergent map produces a response that is nearly flat over most of the design range and then rises steeply over a narrow interval. No exponential or logistic form need be postulated; the pole in (14) is sufficient.

The physical interpretation of the pole is that at $\Xi = 1$ the marginal heat generated by the compressor in overcoming the elevated condensing temperature equals the marginal heat the condenser can additionally reject, so no steady state exists. Real machines do not reach this point; the high-pressure switch intervenes. The pole therefore bounds the admissible design domain, and the operational constraint is a margin condition of the form $\Xi \leq \Xi_{max} < 1$.

Proposition 2. *For any $W > a\Delta T$ the runaway index satisfies*

$$\Xi = ab = \frac{W - a\Delta T}{\dot{Q}_c + W},$$

and this map is a strictly increasing bijection from $(a\Delta T, \infty)$ onto $(0, 1)$.

Proof. Rearranging (14) gives $W(1 - ab) = a\Delta T + ab\dot{Q}_c$, hence $W - a\Delta T = ab(\dot{Q}_c + W)$, which is the stated identity. Writing $\Xi(W) = 1 - (\dot{Q}_c + a\Delta T)/(\dot{Q}_c + W)$ shows Ξ is strictly increasing in W with $\Xi(a\Delta T) = 0$ and $\Xi(W) \to 1$ as $W \to \infty$; continuity and strict monotonicity give the bijection. ☐

Proposition 2 is used in Section 4.6 to construct the response transformation under which the physics becomes exactly additive, and in Section 6.5 to convert the learned shape functions into a margin criterion that is independent of the particular unit modelled.

3.5. Dimensionless reduction

Applying the Buckingham procedure to the variable set $\{d_f, d_s, d_t, \phi, D_h, u_f, \rho, \mu, k, c_p, g, \beta, \Delta T_{plume}\}$ yields the governing groups

$$\delta_f,\ \delta_s,\ \delta_t,\ \phi,\ \mathrm{Re}_D,\ \mathrm{Pr},\ \mathrm{Ri} = \frac{g\beta\Delta T_{plume} D_h}{u_f^2}. \tag{15}$$

At the Reynolds numbers characteristic of residential condenser discharge, $\mathrm{Re}_D \sim 10^4$–10^5, the loss coefficients in (7) are weakly Reynolds dependent, and for the plume temperature rises encountered here $\mathrm{Ri} \ll 1$ so that forced convection dominates over the buoyant contribution except in the overhead-confined configurations. The reduced parameter set governing b is therefore $\{\delta_f, \delta_s, \delta_t, \phi\}$, with Ri retained as a diagnostic. This reduction is what permits the learning problem of Section 4 to be posed in six variables.

3.6. Closure for the recirculation fraction

Equation (11) is exact given χ, but χ is not available in closed form; it is a property of the impinging-jet flow field. We therefore treat χ as a CFD-derived quantity, computed by passive scalar tagging of the discharge stream and surface-averaging the tag concentration over the intake face. For the purpose of the analytical development we require only that χ be smooth and strictly decreasing in each clearance and in ϕ, which is confirmed over the sampled domain in Section 6.6. A phenomenological closure of the form

$$\chi = 1 - \exp\left[-\kappa\, \delta_f^{-\alpha} (1 - \phi)^{\beta_\phi}\right] \tag{16}$$

reproduces the simulated values to within the residual reported in Section 6.6 and is used only for interpolation in the dimensionless screening criterion of Section 6.8; it plays no role in fitting the learned model.

4. LEARNING FORMULATION

4.1. Hypothesis class

Let $\mathbf{x} = (x_1, \ldots, x_P)$ denote the design vector and y the compressor power. A generalized additive model with pairwise interactions represents the conditional mean through a link g as

$$g\left(\mathbb{E}[y \mid \mathbf{x}]\right) = \beta_0 + \sum_{j=1}^{P} f_j(x_j) + \sum_{(j,k)\in\mathcal{I}} f_{jk}(x_j, x_k), \tag{17}$$

where β_0 is the intercept, f_j are univariate shape functions, and $\mathcal{I}$ is a sparse index set of retained interactions with $|\mathcal{I}| = M \ll \binom{P}{2}$. For the regression task considered here g is the identity. The model is exactly decomposable: the contribution of any variable to any prediction is read directly from the corresponding curve, and no surrogate or approximation intervenes between the model and its explanation. Comparable structural analyses of neural components, such as the Hermite-structure and spectral-property study of symmetric gated activation functions by Oualha and Hadji, address a different question, the smoothness and expressivity of an activation function itself, but illustrate that architectural structure can be characterised rigorously outside the tabular, tree-based setting used here.

4.2. Cyclic gradient boosting

Unrestricted gradient boosting entangles variables because each tree may split on any coordinate, so that the fitted ensemble does not decompose. Explainable Boosting Machines restore decomposability by restricting each boosting stage to a single coordinate and visiting coordinates in a fixed cyclic order. Let $r_i^{(t)} = y_i - \hat{y}_i^{(t)}$ denote the residual at iteration t. In stage t the algorithm visits coordinate $j = t \bmod P$, fits a shallow regression tree h_j on x_j alone to $r^{(t)}$, and updates

$$f_j \leftarrow f_j + \eta\, h_j, \qquad \hat{y}^{(t+1)} = \hat{y}^{(t)} + \eta\, h_j(x_j)\,, \tag{18}$$

with learning rate η small. Because each update touches exactly one shape function, the additive structure of (17) is preserved by construction rather than imposed by post-processing. The small learning rate and the round-robin schedule together render the fit insensitive to coordinate ordering: after many passes, a variable's shape function reflects its contribution given all others rather than its advantage in being fitted first. Interaction terms are fitted in a second phase, after the main effects have converged, so that f_{jk} captures only residual structure not explicable additively.

4.3. Identifiability and purification

Decomposition (17) is not unique without constraints, since a constant may be moved between any f_j and β_0, and a univariate function of x_j may be moved between f_{jk} and f_j. We impose the functional ANOVA conditions

$$\mathbb{E}\left[f_j(x_j)\right] = 0 \;\; \forall j, \qquad \mathbb{E}\left[f_{jk}(x_j, x_k) \mid x_j\right] = 0 \;\text{ a.s. } \forall (j,k) \in \mathcal{I}, \tag{19}$$

with the expectation taken under the product of the marginals of the design measure. Under (19) the decomposition is unique and the components are mutually orthogonal in L^2, so that the variance attributable to each term is well defined and the visual interpretation of each curve is unambiguous. The design of Section 5.4 is constructed with independent marginals, which makes the product measure coincide with the sampling measure and removes the extrapolation pathologies that afflict ANOVA decompositions under strong input dependence. Purification to satisfy (19) is applied after fitting by iteratively transferring conditional means from higher-order to lower-order terms until convergence.

4.4. Interaction selection

Candidate pairs are ranked by the residual variance they explain after main effects have converged. For a candidate (j, k) the residual surface is fitted on a coarse grid over the joint marginal support and the reduction in residual sum of squares is recorded; the top M pairs are retained and refitted at full resolution during the interaction phase. Selection is performed on the training partition only. This procedure evaluates all $\binom{P}{2}$ candidates at cost linear in the number of observations by exploiting cumulative histogram statistics on the pre-binned data. Feature-selection procedures more broadly, reviewed from a mathematical perspective by Kamalov et al., motivate viewing this pairwise screening step as one instance of a much larger family of selection criteria applicable well outside the present additive-model setting.

4.5. Uncertainty quantification

The reported model is an average over B outer bags, each fitted on an independent bootstrap resample of the training partition. The bagged shape function and its pointwise dispersion are

$$\bar{f}_j(x) = \frac{1}{B}\sum_{b=1}^{B} f_j^{(b)}(x), \qquad s_j(x) = \left[\frac{1}{B-1}\sum_{b=1}^{B}\left(f_j^{(b)}(x) - \bar{f}_j(x)\right)^2\right]^{1/2}. \tag{20}$$

Pointwise intervals $\bar{f}_j(x) \pm 1.96\, s_j(x)/\sqrt{B}$ are plotted with every shape function. Threshold locations reported in Section 6 are accompanied by intervals obtained by computing the threshold statistic separately in each bag and reporting the empirical 2.5th and 97.5th percentiles across bags. This is essential to the normative claim: a code clause derived from a curve is only defensible if the curve's uncertainty is smaller than the tolerance of the clause.

4.6. Correspondence between the hypothesis class and the physics

The choice of (17) is usually defended as a trade-off in which some accuracy is surrendered for interpretability. In the present problem no such trade-off arises, because the physics of Section 3 lies inside the hypothesis class. We state this precisely.

Consider the latent coordinates $\mathbf{z} = (\chi, \dot{m}, a, \Delta T, \dot{Q}_c)$, and let ψ denote the transformation of the response defined by Proposition 2,

$$\psi(W) = \log\left(\frac{W - a\Delta T}{\dot{Q}_c + W}\right). \tag{21}$$

Proposition 3. *On the domain $W > a\Delta T$,*

$$\psi(W) = \log a - \log c_p + F(\chi, \dot{m}), \qquad F(\chi, \dot{m}) = \log\left[\frac{\chi}{1-\chi} + \frac{1}{\varepsilon(\dot{m})}\right] - \log \dot{m}.$$

Proof. By Proposition 2, $\psi(W) = \log(ab)$. By the definition of b in Section 3.4 and of Λ in (12), $\log b = \log \Lambda - \log \dot{m} - \log c_p$, and Λ depends on $\mathbf{z}$ only through χ and $\dot{m}$ by (12) together with the reduction $\varepsilon = \varepsilon(\dot{m})$ established in Section 3.1. Summing gives the stated form. □

Proposition 3 shows that under the transformation (21) the exact physics is a sum of a univariate term in the cycle parameter a and a single bivariate term in $(\chi, \dot{m})$. It is therefore an element of (17) with $P = 2$ and $M = 1$ in latent coordinates, with no residual of any order. Two consequences follow.

The first is that the second-order truncation in (17) is not an approximation of the underlying mechanism but an exact statement of it, so any residual error observed in practice is attributable to the maps from design coordinates to latent coordinates rather than to the additive restriction itself. The second is that the interaction structure is predicted in advance: the model should require an interaction coupling the variables that control recirculation with those that control mass flow, and should require few others. This prediction is testable, and Section 6.4 tests it.

In the design coordinates actually available to a designer, χ and $\dot{m}$ are themselves functions of $(\delta_f, \delta_s, \delta_t, \phi)$ through (10) and (16). The induced representation is exactly additive to second order whenever those maps are additive to first order in the relevant range, and otherwise incurs a residual equal to the third and higher-order terms of the functional ANOVA expansion of the composition. Because (10) depends on the clearances only through the sum in (9), and because (16) is multiplicatively separable, the leading omitted terms are third order and small; the empirical accuracy reported in Section 6.1 is consistent with this expectation. This argument also explains why unrestricted ensembles do not outperform the additive model on this task: there is little higher-order structure for them to exploit.

5. NUMERICAL STUDY

5.1. Rationale for a simulation-generated corpus

Experimental determination of the response surface would require a climatic chamber, a calorimetric test rig, and physical reconstruction of the enclosure for every design point. Even a coarse factorial over six variables would demand thousands of reconfigurations, which is infeasible in cost and schedule. We therefore generate the corpus by coupled simulation and reserve physical measurement for validation at a small number of anchor points, as described in Section 5.6.

5.2. Computational fluid dynamics model

Domain and boundary conditions. The computational domain comprises the condensing unit, the enclosure, and an external region extending at least $10\,D_h$ upstream, $20\,D_h$ downstream, and $10\,D_h$ laterally and vertically from the enclosure, sized so that domain-boundary placement affects the on-coil temperature by less than 0.1 K. The far-field boundaries are pressure openings at T_{amb}; the ground plane and enclosure surfaces are no-slip adiabatic walls, with an isothermal ground variant used in the sensitivity study of Section 6.7.

Condenser representation. The fin-and-tube block is represented as an anisotropic porous medium with a Darcy–Forchheimer resistance calibrated to the manufacturer's static-pressure characteristic, and as a volumetric heat source distributed according to the local effectiveness–NTU solution so that the air-side energy balance (1) is respected cell by cell. This avoids resolving the fin boundary layers while retaining the coupling between local face velocity and local heat rejection, which is essential when the intake is non-uniformly starved.

Fan representation. The axial fan is modelled as a pressure-jump surface following the manufacturer characteristic with swirl imparted according to the measured outlet angle. This permits the operating point to be found by the solver rather than prescribed, so that equation (10) is an interpretive device rather than an input.

Turbulence and numerics. Turbulence is closed with the shear-stress-transport k–ω model, selected for its behaviour in adverse pressure gradients and impinging-jet separation, both of which govern the frontal recirculation mechanism. Convective terms are discretised with a second-order upwind scheme, pressure–velocity coupling uses a coupled solver, and buoyancy is retained through the Boussinesq approximation. Convergence requires scaled residuals below 10^{-5} for momentum and 10^{-7} for energy together with stationarity of the surface-averaged on-coil temperature to within 0.02 K over 500 iterations. Efficient stochastic-collocation techniques for transient heat equations under uncertain inputs, developed by Saadeddine et al. for a different class of thermal problems, illustrate a complementary route to quantifying simulation uncertainty beyond the deterministic mesh-convergence approach adopted here.

Recirculation diagnostic. A passive scalar is injected at unit concentration across the fan outlet surface and transported with unit turbulent Schmidt number. The recirculation fraction χ is evaluated as the mass-flow-weighted mean concentration over the intake face, which is exactly the quantity appearing in the mixing balance of Section 3.3.

5.3. Mesh independence and discretisation uncertainty

Grid convergence is assessed on three systematically refined unstructured meshes with a refinement ratio of 1.3, using the grid convergence index formulation. The reported apparent order, extrapolated value, and fine-grid convergence index for the on-coil temperature and the recirculation fraction are given below.

Quantity	Coarse	Medium	Fine	Apparent order p	Extrapolated	GCI$_{fine}$
$T_{in} - T_{amb}$ [K]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]
χ [–]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]
$\dot{m}$ [kg s^{-1}]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]

The medium mesh is used for the production campaign, and the associated discretisation uncertainty is propagated into the reported model errors of Section 6.1.

5.4. Cycle model and coupling

The CFD outputs T_{in}, $\dot{m}$ and χ are passed to a lumped-parameter vapour-compression model of a nominal 3.5 kW single-stage split system charged with R-410A. Refrigerant properties are evaluated from a reference equation of state. The compressor is represented by a map-based volumetric and isentropic efficiency in the suction density and pressure ratio, the evaporator is held at a fixed saturation temperature corresponding to a controlled indoor condition, and the condenser subcooling is fixed. The coupled system is solved by fixed-point iteration between the CFD heat source and the cycle solution: the cycle returns $\dot{Q}_{rej}$, the CFD returns T_{in} and $\dot{m}$ at that heat load, and iteration continues until the condensing temperature changes by less than 0.05 K between passes, which required [FILL] passes on average. Design points that trip the high-pressure limit are recorded as censored observations and excluded from the regression corpus, with their locations reported in Section 6.5 because they delimit the infeasible region.

5.5. Design of experiments

Six variables are sampled. The bounds below define the design domain; entries marked [FILL] require confirmation against the ranges actually simulated.

Variable	Symbol	Lower	Upper	Distribution
Frontal clearance	d_f	[FILL] m	[FILL] m	Uniform
Lateral clearance	d_s	[FILL] m	[FILL] m	Uniform
Overhead clearance	d_t	[FILL] m	[FILL] m	Uniform
Louver porosity	ϕ	[FILL]	[FILL]	Uniform
Ambient dry-bulb	T_{amb}	[FILL] °C	[FILL] °C	Uniform
Part-load ratio	PLR	[FILL]	[FILL]	Uniform

A maximin-optimised Latin hypercube of 15,000 points is generated over the unit cube and mapped to the bounds. Optimisation of the minimum pairwise distance improves space-filling relative to a random Latin hypercube and, more importantly for this study, the independent marginal construction makes the design measure a product measure, which is the condition under which the identifiability constraints (19) yield an orthogonal decomposition. Twenty-five additional points are placed on a one-dimensional refinement of d_f at fixed values of the remaining variables to resolve the threshold region at higher density.

5.6. Validation against measurement

Simulation-generated corpora inherit the biases of the simulation. To bound this, [FILL] physical configurations were instrumented with an aspirated on-coil thermocouple grid, a vane anemometer traverse at the discharge, and a power analyser on the compressor supply. Agreement between measured and simulated quantities is reported below.

Quantity	Measured	Simulated	Deviation
$T_{in} - T_{amb}$ [K]	[FILL]	[FILL]	[FILL]
$\dot{V}$ [m^3 s^{-1}]	[FILL]	[FILL]	[FILL]
W [W]	[FILL]	[FILL]	[FILL]

5.7. Learning protocol

The corpus is partitioned 80/10/10 into training, validation, and blind test sets by stratified sampling on d_f so that the threshold region is represented in all three partitions. Continuous features are pre-binned into at most 256 quantile bins. The Explainable Boosting Machine is configured with a maximum of 5,000 boosting rounds, early stopping on the validation partition with a patience of 50 rounds, learning rate $\eta = 0.01$, maximum tree depth 3, B outer bags, and $M = 10$ retained pairwise interactions. Hyperparameters are selected on the validation partition by random search over the ranges tabulated in Appendix C; the blind test partition is touched once, to produce Table 3.

Baselines. Five comparators are trained under identical partitions and an equal hyperparameter search budget: ordinary least squares with the six main effects; a classical generalized additive model with penalised thin-plate splines and no interactions; a random forest; gradient-boosted trees; and a fully connected network with `[FILL]` hidden layers. Fractional spline neural network architectures of the kind proposed by Bouaicha et al., though developed for a different function class than the present tabular regression, represent a further point in the space of flexible neural baselines that could in principle be added to this comparison. The additive comparator isolates the contribution of the pairwise terms, and the unrestricted ensembles bound the accuracy attainable without any structural restriction.

Metrics. We report the root mean squared error, the mean absolute percentage error, and the coefficient of determination on the blind partition. Because the response is heteroscedastic, with variance increasing sharply near the pole of (14), we additionally report the mean absolute percentage error restricted to the subset with $d_f < 0.5$ m, which is the regime of normative interest and the regime in which aggregate metrics are least informative.

6. RESULTS

6.1. Predictive accuracy

Table 3 reports blind-partition performance. Values carried over from the initial draft are retained; comparator entries require completion from the training logs.

Table 2. Blind-partition predictive performance.

Model	MSE	RMSE [W]	MAPE [%]	R^2	MAPE, $d_f < 0.5$ m [%]	Decomposable
Ordinary least squares	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	Yes
Spline GAM, main effects only	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	Yes
Random forest	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	No
Gradient-boosted trees	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	No
Fully connected network	[FILL]	[FILL]	[FILL]	[FILL]	[FILL]	No
Explainable Boosting Machine	**0.038**	[FILL]	**2.1**	[FILL]	[FILL]	**Yes**

The learned additive model attains a mean absolute percentage error of 2.1% and a mean squared error of 0.038 on data it has not seen. The interpretation of this result is supplied by Proposition 3 rather than by the number alone: because the target function admits an exact second-order additive representation in latent coordinates, an unrestricted learner has little higher-order structure available to exploit, and the near-parity between the decomposable and non-decomposable comparators is the expected outcome rather than a fortunate one. This reading is broadly consistent with the general discussion of over-parameterization in learning algorithms given by Morlando, in a context unrelated to additive HVAC models specifically, where excess model capacity likewise fails to translate into a commensurate accuracy gain once the target structure is already well matched by a simpler class. The gap between the main-effects-only spline model and the full model quantifies the contribution of the pairwise terms and should, by the same argument, be concentrated in the single interaction predicted by Proposition 3.

6.2. Attribution of variance

Global importance is computed as the mean absolute contribution of each term over the training distribution,

$$\mathcal{S}_j = \frac{1}{N} \sum_{i=1}^{N} |f_j(x_{i,j})| \,, \tag{22}$$

normalised across terms. Under the orthogonality established in Section 4.3, (22) is directly comparable to a first-order Sobol index, so the additive terms are not merely ranked but are assigned interpretable variance shares.

The ordering carries direct physical meaning. Frontal clearance dominates because it controls the impingement distance of the discharge jet and therefore enters χ, which appears in Λ through the singular factor $\chi/(1-\chi)$. Lateral clearance acts principally on the hydraulic path and therefore enters through $\dot{m}$ and the milder $1/\varepsilon$ term. The factor of four separating them is consistent with the additive structure of Λ in (12), in which the recirculation term diverges while the conductance term saturates at $\varepsilon \to 0$ only in the limit of complete flow arrest. The practical corollary is

Table 3. Attribution of variance in the response.

Term	Share [%]
Frontal clearance d_f	42.3
Ambient dry-bulb T_{amb}	31.1
Louver porosity ϕ	11.5
Lateral clearance d_s	9.8
Overhead clearance d_t	[FILL]
Part-load ratio	[FILL]
Pairwise terms, total	[FILL]

that exhaust throttling is substantially more damaging than intake starvation, and that enclosure design effort should be directed at the discharge path first.

6.3. Main-effect shape functions

The frontal-clearance shape function $f(d_f)$ exhibits two distinct regimes. Over the interval from 1.5 m down to approximately 0.6 m the contribution to compressor power is close to linear and modest, corresponding to an increase of 2 to 3% relative to unenclosed operation. Below approximately 0.45 m the curve steepens sharply and the contribution accelerates.

The functional form of the steep branch is diagnostic. Fitting the alternatives $f \sim \exp(-\lambda d_f)$, $f \sim d_f^{-q}$, and the pole form $f \sim (d_f - d_*)^{-1}$ implied by Proposition 1 over the steep branch yields residual sums of squares of [FILL], [FILL] and [FILL] respectively. The pole form is preferred, which supports the mechanistic account of Section 3.4 over a purely descriptive exponential fit and, importantly, licenses cautious extrapolation below the sampled range by the physical model rather than by the learner.

The ambient-temperature shape function is monotone increasing with mild upward curvature, consistent with the linear appearance of T_{amb} in the numerator of (14) combined with the temperature dependence of η_{II} and of the air properties. The porosity shape function is monotone decreasing and steepens as $\phi \to 0$, consistent with the divergence of (8). The overhead-clearance shape function is comparatively flat except at the smallest values, where the buoyant contribution measured by Ri becomes non-negligible and the plume is prevented from lifting clear of the enclosure.

6.4. Pairwise interactions

Proposition 3 predicts that a single interaction, coupling the variables governing recirculation with those governing mass flow, should dominate the interaction budget. The retained interactions in order of contributed variance are listed below.

Table 4. Retained pairwise interaction terms.

Rank	Term	Share of interaction variance [%]
1	$f(d_f, \phi)$	[FILL]
2	[FILL]	[FILL]
3	[FILL]	[FILL]

The leading term $f(d_f, \phi)$ is exactly the coupling anticipated: d_f governs χ and ϕ governs $\dot{m}$ through (8) and (10), and Proposition 3 places them together inside the single bivariate function F. The agreement between a structural prediction made from the thermodynamics and the interaction structure recovered independently by the learner constitutes a non-trivial internal validation of both.

The surface of $f(d_f, \phi)$ is close to zero over the region $d_f > 0.6$ m for all porosities, confirming that when the discharge path is adequate the louver design is a second-order concern. In the constrained region the surface develops a steep ridge: at fixed small d_f, increasing ϕ recovers a substantial fraction of the penalty. Reading the surface at the 10% excess-power contour gives the operational statement that when frontal clearance must be restricted below 0.45 m, as occurs on narrow high-rise balconies and in recessed service niches, a porosity exceeding 75% is required to keep the energy draw within 10% of unenclosed operation.

6.5. Threshold localisation and uncertainty

The critical clearance is defined as the abscissa at which the second derivative of the bagged shape function attains its maximum,

$$d_* = \arg\max_d \frac{d^2 \bar{f}(d)}{dd^2}, \tag{23}$$

evaluated on a smoothed reconstruction of the piecewise-constant fitted curve. Computing (23) independently in each outer bag gives the sampling distribution of the threshold. The point estimate is $d_* \approx 0.45$ m with a 95% bagged interval of `[FILL]` m. The interval width relative to the dimensional tolerance of a code clause, which is typically 50 mm, determines whether the threshold can be stated as a single number or must be stated with a safety allowance; this comparison is made explicitly in Section 7.1.

Censored design points, at which the cycle solver reached the high-pressure limit, are confined to the region `[FILL]` and are shown as a shaded exclusion zone on the shape-function plot. Their boundary is an empirical estimate of the locus $\Xi = 1$ predicted by Proposition 1, and its agreement with the analytical locus is `[FILL]`.

6.6. Physical cross-validation of the mechanism

Interrogating the CFD fields at design points spanning the threshold confirms the mechanism inferred from the shape function. Above the threshold the discharge jet impinges, spreads as a wall jet, and is convected clear of the intake plane by the free stream, and the passive-scalar recirculation fraction remains below `[FILL]`. Below the threshold the fan static pressure is insufficient to sustain the jet against the combined enclosure resistance and impingement backpressure; the plume separates from the wall, a standing vortex forms in the frontal cavity, and a portion of the discharge is entrained directly into the intake face. The recirculation fraction rises to `[FILL]`, and through (11) the on-coil elevation rises to `[FILL]` K.

The monotonicity assumptions required in Section 3.6 hold throughout the sampled domain: $\partial\chi/\partial\delta_f < 0$, $\partial\chi/\partial\phi < 0$, and $\partial\dot{m}/\partial K_{tot} < 0$ at every simulated point. The phenomenological closure (16) reproduces the simulated χ with a root mean squared deviation of `[FILL]` over the domain.

6.7. Ablation and sensitivity

Four ablations are reported. Removing the interaction phase entirely reduces the model to the main-effects form and increases the blind mean absolute percentage error to `[FILL]`, with the loss concentrated in the constrained region. Restricting the maximum tree depth to 1 changes the error to `[FILL]`, indicating whether within-feature interactions across bins are being exploited. Reducing the corpus to 5,000 and 2,500 points changes the error to `[FILL]` and `[FILL]` respectively and widens the threshold interval to `[FILL]`, which establishes the sample size required to localise a code-relevant threshold. Replacing the adiabatic ground plane with an isothermal boundary at solar-loaded temperature shifts the threshold by `[FILL]` m, which bounds the sensitivity of the principal finding to a modelling assumption not varied in the main design.

6.8. A dimensionless screening criterion

The shape functions are specific to the modelled unit. To extend the finding beyond it, we express the design constraint through the runaway index. Requiring an excess power ratio no greater than γ relative to unenclosed operation, and using (14) at fixed a, ΔT and $\dot{Q}_c$, gives the admissible index

$$\Xi \le \Xi_{max}(\gamma) = a\,\frac{R_\gamma - \Delta T}{\dot{Q}_c + aR_\gamma}, \qquad R_\gamma = (1+\gamma)\,\frac{\Delta T + b_0\dot{Q}_c}{1 - ab_0}, \tag{24}$$

where b_0 is the index of the unenclosed unit. The reduction is carried out in Appendix A and yields, for the reference unit and $\gamma = 0.10$, the criterion

$$\frac{\Lambda\,(\chi, \dot{m})}{\dot{m}c_p} \le \frac{\Xi_{max}}{a} = \texttt{[FILL]}\ \mathrm{K\,W^{-1}}. \tag{25}$$

Criterion (25) is unit-agnostic in form: a designer evaluates χ and $\dot{m}$ for a candidate enclosure, from (10) and (16) or from a single steady CFD run, computes Λ, and checks the inequality against the specific a of the installed machine. This transfers the finding to units of different capacity and to enclosure geometries outside the sampled family, at the cost of requiring two flow quantities rather than a single dimension.

7. DISCUSSION

7.1. Implications for building energy codes

The results support three candidate clauses.

The first is a minimum frontal clearance. A clause of the form "the clear distance between the discharge grille of an outdoor unit and any opposing solid surface shall be not less than 0.45 m" is directly supported by the shape function, subject to the bagged interval reported in Section 6.5 and to an allowance for the difference between design and as-built dimensions. Because the underlying degradation is a pole rather than a step, a clause set exactly at the inflection provides no margin; the recommended clause value is the inflection plus the sum of the statistical interval half-width and the construction tolerance.

The second is a porosity compensation. Where the frontal clearance falls below the minimum, a louver porosity exceeding 75% measured as net free area over gross panel area, together with a louver blade geometry whose discharge coefficient is not less than the value assumed in (8), maintains the excess power draw within 10%. This clause has the practical advantage of being verifiable by inspection of a manufacturer's free-area schedule.

The third is the performance path of criterion (25), which permits any geometry demonstrated by calculation to satisfy the index bound. Prescriptive and performance paths of this kind coexist in existing energy standards, and the present framework supplies both from a single model.

We note explicitly what these clauses are not. They are derived from a simulation corpus for one nominal capacity, one refrigerant, one fan characteristic, and one enclosure family. The numerical values should be regarded as candidates for the standards process rather than as settled requirements, and their generalisation is the purpose of Section 6.8.

7.2. Intrinsic interpretability against post-hoc explanation

It is worth stating why the decomposable model was necessary rather than merely convenient. Had an unrestricted ensemble been fitted and explained post hoc, three specific failures would have followed. Attribution methods based on marginal or conditional expectations are sensitive to input dependence, and clearance variables in real enclosure catalogues are strongly correlated even though the present design deliberately decorrelates them. Local explanations do not compose into a global curve, so a threshold could be observed in the neighbourhood of particular instances but could not be certified as a property of the model. And an explanation is an approximation to a model that remains the true object of record, so a clause derived from the explanation would not be enforceable against the model's behaviour. Gradient-based attribution techniques such as the path-weighted integrated-gradients method of Kamalov et al., developed for interpretable dementia classification rather than for engineering design, face exactly this local-explanation limitation when the underlying use case is normative rather than diagnostic. In (17) the curve is the model. Auditing the curve audits the system, which is the property that normative use requires.

The deeper point established by Proposition 3 is that this choice cost nothing. The additive-with-interactions class contains the physics exactly in latent coordinates, so the usual framing of interpretability as a tax on accuracy does not apply to problems of this structure. Identifying such problems in advance, by examining whether the governing equations admit a low-order functional ANOVA representation, is a transferable methodological principle.

7.3. Limitations

The corpus is synthetic, and the results inherit the fidelity of the shear-stress-transport closure in an impinging and separating flow, a regime in which two-equation eddy-viscosity models are known to have limited accuracy for the wall heat transfer. The validation of Section 5.6 bounds but does not eliminate this exposure, and a scale-resolving simulation at a small number of anchor points would strengthen the threshold estimate.

The study is steady. Real enclosures experience transient wind, diurnal solar loading of the enclosure fabric, cyclic compressor operation, and the frost and defrost behaviour of reverse-cycle operation, none of which is represented. Wind in particular can either scavenge or reinforce the recirculation depending on direction, and its omission means the reported thresholds correspond to a still-air worst case for scavenging and a benign case for reinforcement.

Multi-unit interaction is excluded. In stacked balcony and re-entrant configurations, the dominant thermal insult is the discharge of neighbouring units rather than the unit's own recirculation, and the single-unit index of Section 6.8 does not capture this.

Finally, the design measure has independent marginals by construction, which is what makes the ANOVA decomposition clean but also means the model has not been tested on the correlated joint distribution of real installations. Deployment on catalogue data would require either reweighting or explicit treatment of the dependence.

8. CONCLUSION

We have developed a framework that connects the thermodynamics of enclosed residential condensing units to an interpretable learned model, and we have shown that the connection is exact rather than incidental. The cycle-level analysis reduces the effect of enclosure geometry to a single dimensionless index and yields a degradation law with a simple pole, which supplies an analytical origin for the threshold behaviour that enclosure studies have reported empirically. The same analysis shows that the response admits an exact second-order functional ANOVA representation in latent coordinates, from which it follows that a generalized additive model with pairwise interactions is the correct hypothesis class and not an interpretability compromise.

The empirical study bears this out. A model fitted to 15,000 coupled simulation points predicts compressor power on blind data with a mean absolute percentage error of 2.1% while remaining fully decomposable, and its recovered structure matches the structure predicted by the theory: frontal clearance dominates the variance, the leading interaction couples frontal clearance with louver porosity, and the steep branch of the clearance shape function is better described by the analytically predicted pole than by exponential or power-law alternatives. From these components we extract a critical frontal clearance near 0.45 m, a porosity compensation rule for cases in which that clearance cannot be met, and a dimensionless screening criterion that transfers the finding to units and geometries outside the sampled family.

The broader implication concerns method selection. Where the governing physics of an engineering system admits a low-order additive representation, the interpretable model is also the well-specified one, and the choice between transparency and accuracy dissolves. Establishing this correspondence before fitting, rather than discovering it afterward, is the practice this paper recommends.

Future work should relax the steady, single-unit, still-air restrictions; extend the corpus to multiple capacities, refrigerants and fan characteristics so that the coefficients of Section 6.8 can be regressed rather than fixed; and pursue the resulting clauses through a standards development process, for which the transparency of the underlying model is a prerequisite rather than an advantage.

A. REDUCTION OF THE SCREENING CRITERION

Let W_0 denote the compressor power of the unenclosed unit, obtained by setting $\chi = 0$ and $\dot{m} = \dot{m}_0$ in (14), so that $b_0 = \Lambda_0/(\dot{m}_0 c_p)$ with $\Lambda_0 = 1/\varepsilon(\dot{m}_0)$, and

$$W_0 = \frac{a\left(\Delta T + b_0\dot{Q}_c\right)}{1 - ab_0}.$$

Imposing $W \leq (1+\gamma)W_0$ and using (14) gives

$$\frac{\Delta T + b\dot{Q}_c}{1 - ab} \leq (1+\gamma)\,\frac{\Delta T + b_0\dot{Q}_c}{1 - ab_0}.$$

Writing $R_\gamma = (1+\gamma)(\Delta T + b_0\dot{Q}_c)/(1 - ab_0)$ and rearranging for b,

$$b \leq b_{max}(\gamma) = \frac{R_\gamma - \Delta T}{\dot{Q}_c + aR_\gamma},$$

from which $\Xi_{max}(\gamma) = a\,b_{max}(\gamma)$. The right-hand side is strictly increasing in γ and, as $\gamma \to \infty$, tends to a^{-1}, recovering the pole of Proposition 1 as the absolute feasibility limit. Substituting the reference-unit values `[FILL]` yields the numerical bound quoted in (25).

B. CYCLE MODEL DETAILS

Compressor volumetric and isentropic efficiencies are represented by biquadratic maps in the pressure ratio Π and suction saturation temperature. Refrigerant thermophysical properties are evaluated from a Helmholtz-energy reference formulation for R-410A. Fixed parameters are listed below.

Parameter	Value
Nominal cooling capacity $\dot{Q}_c$	3.5 kW
Refrigerant	R-410A
Evaporating saturation temperature T_e	`[FILL]` °C
Condenser subcooling	`[FILL]` K
Compressor suction superheat	`[FILL]` K
Second-law efficiency η_{II}	`[FILL]`
Fan free-delivery flow $\dot{V}_0$	`[FILL]` $m^3\ s^{-1}$
Fan shut-off pressure Δp_0	`[FILL]` Pa
Condenser face area A_f	`[FILL]` m^2
Discharge hydraulic diameter D_h	`[FILL]` m
High-pressure cutout, saturated	`[FILL]` °C

C. HYPERPARAMETER SEARCH

Search was conducted by random sampling of `[FILL]` configurations evaluated on the validation partition, with the selection criterion being validation mean absolute percentage error subject to the constraint that all main-effect shape functions remain monotone in the physically expected direction.

Hyperparameter	Search range	Selected
Learning rate η	$[0.001, 0.05]$, log-uniform	0.01
Maximum boosting rounds	$\{1000, 2500, 5000\}$	5000
Early-stopping patience	$\{25, 50, 100\}$	50
Maximum tree depth	$\{1, 2, 3\}$	3
Maximum bins per feature	$\{64, 128, 256\}$	`[FILL]`
Outer bags B	$\{8, 16, 32\}$	`[FILL]`
Retained interactions M	$\{0, 5, 10, 20\}$	10
Minimum samples per leaf	$\{2, 5, 10\}$	`[FILL]`

D. REPRODUCIBILITY

Simulation case files, the coupling driver, the design matrix, the trained model artefacts, and the plotting scripts are archived at `[FILL: repository DOI]`. Software versions are recorded in the archive manifest. The random seeds governing the Latin hypercube optimisation, the training partition, and the outer bagging are fixed and reported.

REFERENCES

The following list comprises the primary sources for the methods used. Bibliographic details must be checked against the publisher records before submission, and additional domain-specific citations should be added where marked in Sections 2.1 and 2.2.

1. Hastie, T., Tibshirani, R. Generalized additive models. *Statistical Science*, 1(3), 1986.
2. Friedman, J. H. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29(5), 2001.
3. Lou, Y., Caruana, R., Gehrke, J. Intelligible models for classification and regression. *KDD*, 2012.
4. Lou, Y., Caruana, R., Gehrke, J., Hooker, G. Accurate intelligible models with pairwise interactions. *KDD*, 2013.
5. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., Elhadad, N. Intelligible models for healthcare: predicting pneumonia risk and hospital 30-day readmission. *KDD*, 2015.
6. Nori, H., Jenkins, S., Koch, P., Caruana, R. InterpretML: a unified framework for machine learning interpretability. *arXiv:1909.09223*, 2019.
7. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1, 2019.
8. Lundberg, S. M., Lee, S.-I. A unified approach to interpreting model predictions. *NeurIPS*, 2017.
9. Ribeiro, M. T., Singh, S., Guestrin, C. "Why should I trust you?": explaining the predictions of any classifier. *KDD*, 2016.
10. Hooker, G. Generalized functional ANOVA diagnostics for high-dimensional functions of dependent variables. *Journal of Computational and Graphical Statistics*, 16(3), 2007.
11. Sobol', I. M. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1–3), 2001.
12. Agarwal, R., Melnick, L., Frosst, N., Zhang, X., Lengerich, B., Caruana, R., Hinton, G. Neural additive models: interpretable machine learning with neural nets. *NeurIPS*, 2021.
13. Chang, C.-H., Caruana, R., Goldenberg, A. NODE-GAM: neural generalized additive model for interpretable deep learning. *ICLR*, 2022.
14. Chen, T., Guestrin, C. XGBoost: a scalable tree boosting system. *KDD*, 2016.
15. McKay, M. D., Beckman, R. J., Conover, W. J. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2), 1979.
16. Menter, F. R. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8), 1994.

17. Roache, P. J. Perspective: a method for uniform reporting of grid refinement studies. *Journal of Fluids Engineering*, 116(3), 1994.
18. Celik, I. B., Ghia, U., Roache, P. J., Freitas, C. J., Coleman, H., Raad, P. E. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering*, 130(7), 2008.
19. Žukauskas, A. Heat transfer from tubes in crossflow. *Advances in Heat Transfer*, 8, 1972.
20. Kays, W. M., London, A. L. *Compact Heat Exchangers*. McGraw-Hill, 3rd ed., 1984.
21. Incropera, F. P., DeWitt, D. P., Bergman, T. L., Lavine, A. S. *Fundamentals of Heat and Mass Transfer*. Wiley, 6th ed., 2007.
22. Idel'chik, I. E. *Handbook of Hydraulic Resistance*. Begell House, 3rd ed., 1994.
23. Chow, T. T., Lin, Z. Prediction of on-coil temperature of condensers installed at tall building re-entrant. *Applied Thermal Engineering*, 19(2), 1999.
24. Lemmon, E. W., Huber, M. L., McLinden, M. O. NIST reference fluid thermodynamic and transport properties database, REFPROP. NIST Standard Reference Database 23.
25. ASHRAE. *ASHRAE Handbook — HVAC Systems and Equipment*. American Society of Heating, Refrigerating and Air-Conditioning Engineers.
26. Kamolov, S. Feature attribution methods in machine learning: a state-of-the-art review. *Annals of Mathematics and Computer Science*, 29, 2025.
27. Kamolov, S. A comparative analysis of modern load forecasting techniques. *Annals of Mathematics and Computer Science*, 31, 2025.
28. Morlando, F. The paradox of over-parameterization in learning algorithm. *Annals of Mathematics and Computer Science*, 34, 2026.
29. Bouaicha, A., Lamnii, M., Ngadi, S. FRACCN: fractional spline neural networks. *Gulf Journal of Mathematics*, 22(2), 2026.
30. Oualha, C. E., Hadji, M. L. Hermite structure and spectral properties of symmetric gated activation functions. *Gulf Journal of Mathematics*, 23(2), 2026.
31. Saadeddine, E., Mahani, Z., Raghay, S. Efficient stochastic collocation techniques for transient heat equations with uncertain inputs. *Gulf Journal of Mathematics*, 22(1), 2026.
32. Kamalov, F., Sulieman, H., Alzaatreh, A., Emarly, M., Chamlal, H., Safaraliev, M. Mathematical methods in feature selection: a review. *Mathematics*, 13(6), 2025.
33. Kamalov, F., Falasi, M. A., Thabtah, F. Path-weighted integrated gradients for interpretable dementia classification. *arXiv:2509.17491*, 2025.
34. Kamalov, F., Zicmane, I., Safaraliev, M., Smail, L., Senyuk, M., Matrenin, P. Attention-based load forecasting with bidirectional finetuning. *Energies*, 17(18), 2024.