Journal of Multidisciplinary Engineering
Published by DigitoPub | Digitodontics International Academy
DigitoPub
Open Access & Peer-reviewed
Indexed in: Crossref, ICI, Google Scholar

Artificial Neural Network Surrogate Model for Predicting the Overall Heat-Transfer Coefficient in a Glycol Heat Exchanger Using Aspen HYSYS Data

Najib M. Omar[1,2]* | Rabeeah H. Husayn[1] | Muhend Milad[3,4] | Areej M. Ali[1] | Hajer H. Ijdedhum[1]

[1]Chemical Engineering Department, Faculty of Engineering, Sabratha University, Libya.
[2]Advanced Membrane Technology Research Centre (AMTEC), Universiti Teknologi Malaysia, Malaysia.
[3]Department of Petroleum Engineering, Faculty of Chemical and Energy Engineering, Universiti Teknologi Malaysia, Malaysia.
[4]Department of Petroleum Engineering, Faculty of Engineering, University of Zawia, Libya.

ABSTRACT

Purpose: The overall heat-transfer coefficient (U) is an important descriptor of heat-exchanger performance in natural-gas dehydration. This study develops a compact artificial neural network (ANN) surrogate model for predicting U in a rich-monoethylene-glycol (MEG) heat exchanger represented by a steady-state Aspen HYSYS V14 model.

Materials and Methods: The dataset contains 1,500 structured computational operating points, comprising one converged HYSYS reference case and 1,499 points generated using a thermal calculation framework calibrated against the HYSYS reference effectiveness and effective heat capacities. Predictor auditing identified four varying operating variables: hot- and cold-stream inlet temperatures and the two stream mass flow rates. Exchanger geometry and fouling parameters were fixed. A 4-10-1 multilayer perceptron was developed using grouped operating-condition validation.

Results: On the fixed test partition, the ANN achieved $R^2 = 0.9999990$, RMSE = 1.39×10^{-3} W/(m².K), MAE = 1.15×10^{-3} W/(m².K), and MAPE = 0.0088%. Across ten repeated grouped partitions, the mean test RMSE was $(1.90 \pm 0.32) \times 10^{-3}$ W/(m².K). Flow-edge hold-outs produced RMSE values of 2.41×10^{-2} to 2.47×10^{-2} W/(m².K). The reference U was 13.320 W/(m².K). Because the generator holds effectiveness, effective heat capacities, geometry, and fouling resistance constant, the target exhibits a flow-dependent relationship with 16 distinct U values.

Conclusions: The ANN accurately reproduces the HYSYS-calibrated computational relationship within the investigated operating domain. The proposed ANN provides a compact reduced-order surrogate for rapid repeated evaluation of the specified exchanger configuration. *(J Multidiscip Eng 2026;1:125-149)*

Keywords: Artificial neural network, Glycol, Natural-gas dehydration, Aspen HYSYS, Overall heat-transfer coefficient.

ARTICLE INFO

*Corresponding author:
Najib M. Omar
Email: najib.omar@sabu.edu.ly

Received: 23-09-2026
Accepted: 17-11-2026
Published: 19-11-2026

DOI:
https://doi.org/10.26629/jme.2026.16

Copyright: © 2026 by the authors.
Articles published in the *J Multidiscip Eng* are licensed under a Creative Commons Attribution-Noncommercial 4.0 International License.
https://creativecommons.org/licenses/by-nc/4.0/

Audioabstract:

How to cite this article | Omar NM, Husayn RH, Milad M, Ali AM, Ijdedhum HH. Artificial neural network surrogate model for predicting the overall heat-transfer coefficient in a glycol heat exchanger using aspen HYSYS data. J Multidiscip Eng. 2026;1:125-149.

1. INTRODUCTION

Natural gas is rarely pipeline-quality at the wellhead and is often water-saturated, requiring dehydration to prevent hydrate formation, corrosion and two-phase flow in transmission lines.[1] Glycol-based absorption is widely used because of its relatively low capital cost, established thermodynamics and suitability for remote and offshore operation.[2,3] TEG is the most

common solvent, while MEG and DEG are also used offshore and at low temperatures, where dehydration and hydrate inhibition may be combined.[4] The rich/lean-glycol heat exchanger is an important thermal-integration element because it preheats the rich solvent leaving the contactor using hot lean solvent from the regenerator. Its performance, represented by the overall heat-transfer coefficient (U), is relevant to exchanger sizing, energy requirements, performance assessment and optimisation.[5,6] Mechanistic approaches such as the Kern and Bell–Delaware methods require detailed property and heat-transfer calculations and can be sensitive to fouling and off-design operation.[7] Consequently, data-driven surrogates, particularly artificial neural networks (ANNs) trained on simulated or experimental data, have increasingly been investigated for heat-transfer prediction.[8]

ANNs have also been applied to glycol-dehydration and heat-transfer problems. Goh et al.[9] developed an ANN surrogate for predicting several responses of a TEG dehydration process, including dehydrated-gas water content, whereas Akay and Das[10] modelled the overall heat-transfer coefficient of a nuclear-reactor primary loop. Bobič[11] demonstrated a digital-twin framework for a plate heat exchanger, and Omar et al.[12] applied digital-twin technology in the remediation industry. More recently, machine-learning surrogates have been developed from rigorous process simulations; Domingos et al.[13] trained and compared machine-learning models using an Aspen HYSYS crude-distillation model. Srinivasan et al.[14] emphasised the importance of reporting the data-generation mechanism, benchmark comparisons and domain of validity when interpreting surrogate accuracy, while Danesh et al.[15] demonstrated the value of interpretability methods for neural-network predictions. Simulator-calibrated data generation is likewise established in the surrogate-modelling literature,[16-18] where surrogate fidelity is inherently constrained by the fidelity of the underlying model.

Accordingly, a distinction is made between physics-calibrated and physics-informed surrogates. A physics-calibrated surrogate uses a physics-based model to generate or calibrate the target data, whereas a physics-informed model incorporates governing physical constraints directly into the training objective or architecture.[16,17] The present study adopts the former approach and therefore evaluates the ANN primarily in terms of its ability to reproduce the HYSYS-calibrated computational relationship within the declared domain.

For MEG-based dehydration, relatively few studies have combined a HYSYS-calibrated reduced-order surrogate with an explicit applicability domain, leakage-aware predictor selection, grouped validation and benchmark comparisons for a specified exchanger. Existing ANN studies of glycol dehydration primarily address process-level or TEG responses,[9] while studies predicting (U) commonly rely on experimental databases or mechanistic approaches.[8,10] The present study therefore investigates whether a compact ANN can reproduce the overall heat-transfer coefficient obtained from an Aspen HYSYS-based computational workflow for a rich-MEG exchanger, building on previous process-simulation and optimisation studies of gas-processing units.[19,20]

The study has five objectives: (i) document the steady-state HYSYS reference model; (ii) define and audit the HYSYS-based dataset; (iii) develop a compact four-input ANN surrogate for (U); (iv) evaluate its robustness using grouped operating-condition splits, repeated validation and edge-condition hold-outs; and (v) benchmark the ANN against simpler models while defining the scope of computational validation and the pathway to external validation.

2. PROCESS DESCRIPTION AND ASPEN HYSYS STEADY-STATE MODEL

The HYSYS case represents a steady-state MEG-based dehydration loop containing absorption, pressure reduction, heat exchange, pumping, regeneration and recycle. The case is "Dehydration with MEG.hsc", and the exchanger of interest is "Rich Ethylene Glycol Exch @Main" (Fig. 1). The Peng-Robinson fluid package is used for the sour service. The converged reference case contains absorber T-101, valve VLV-100, the

rich/lean glycol exchanger, pump P-100, regenerator, reboiler, condenser and recycle loop. The inlet natural-gas conditions are 29.44 °C, 6205 kPa and 9183 kg/h. In the exchanger, the hot-side inlet is Regen Bttms, the cold-side inlet is LP DEG, the hot-side outlet is Lean from L/R, and the cold-side outlet is Regen Feed; their reported temperatures and mass flow rates form the reference basis for the calibrated data generation.

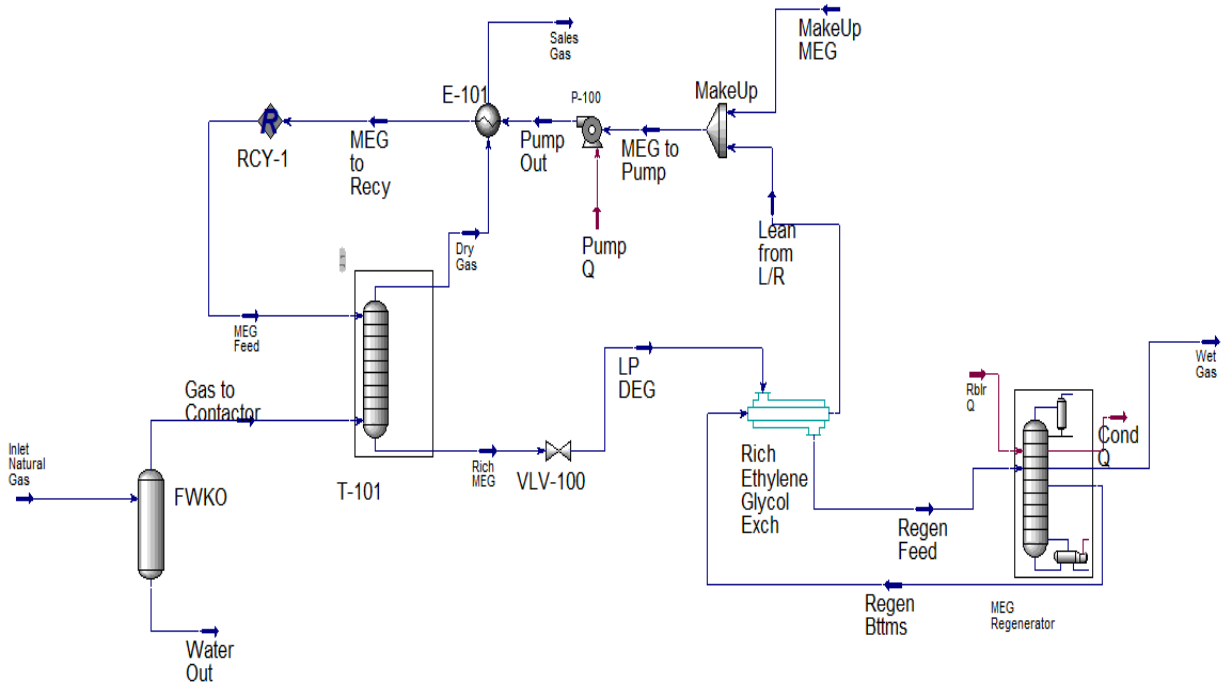


Fig. 1: Aspen HYSYS V14 process flowsheet.

2.1 Reference operating condition

The inlet conditions define the thermal and material boundary conditions of the process (Table 1). The gas is methane-rich, with 2.84 mol% CO_2 and 1.55 mol% H_2S. The exchanger duty obtained from the change in the reported HYSYS heat-flow values is approximately 27.5 kW, and the heat-transfer area of 21.545 m² is taken from the design basis recorded in the calculation workbook. Both are stated explicitly to support reproducibility and transparent interpretation of the resulting heat-transfer coefficient. Supplementary Table S1 provides the detailed reference-condition values.

Table 1: Reference conditions for the HYSYS model.

Quantity	HYSYS/reference value
Hot-stream inlet temperature	193 °C
Cold-stream inlet temperature	30.87 °C
Hot-stream mass flow	504.1 kg/h
Cold-stream mass flow	509.2 kg/h
Hot-stream outlet temperature	134.3 °C
Cold-stream outlet temperature	104.4 °C
Heat duty	27.5 kW
Assumed heat-transfer area	21.545 m²
Reference LMTD	95.824 K
Reference U	13.320 W/(m².K)

2.2 Lean-solvent and rich/lean stream inventory

Table 2 summarises the converged MEG-loop streams obtained from the steady-state simulation. The loop comprises absorption of water by the lean glycol followed by regeneration and return of the lean solvent to the contactor; the reported temperatures, pressures, flow rates and compositions provide the process context required to evaluate the rich-MEG exchanger and the corresponding simulation-derived ANN target values. The corresponding detailed stream inventory is provided in Supplementary Table S3.

Table 2: Main process-stream conditions.

Stream	Vapor fraction	T (°C)	P (kPa)	Molar flow (kgmol/h)	Mass flow (kg/h)	Vol. flow (m^3/h)	Heat Flow (kJ/h)
Inlet Natural Gas	1	29.44	6205	498.1	9183	27.58	-4.257×10^7
MEG Feed (lean)	0	48.89	6205	8.130	504.5	0.4542	-3.676×10^6
Gas to Contactor	1	29.44	6205	498.1	9183	27.58	-4.257×10^7
Dry Gas	1	30.48	6205	497.9	9178	27.57	-4.253×10^7
Rich MEG	0	29.43	6205	8.289	509.2	0.4626	-3.716×10^6
LP DEG	0.0136	30.87	179.3	8.289	509.2	0.4626	-3.716×10^6
Regen Bttms	0	193.0	103.4	8.122	504.1	0.4538	-3.455×10^6
Lean from L/R	0	134.3	102.7	8.122	504.1	0.4538	-3.554×10^6
Regen Feed	0.0184	104.4	110.3	8.289	509.2	0.4626	-3.617×10^6
Wet Gas	1	101.0	101.4	0.1667	5.143	8.741×10^{-3}	-1.725×10^4
MakeUp MEG	0	15.56	102.7	7.454×10^{-3}	0.4594	4.137×10^{-4}	-3397
MEG to Pump	0	134.2	102.7	8.130	504.5	0.4543	-3.557×10^6
Pump Out	0	135.3	6274	8.130	504.5	0.4543	-3.553×10^6
MEG to Recy	0	48.89	6205	8.130	504.5	0.4543	-3.676×10^6
Sales Gas	1	35.64	6171	497.9	9178	27.57	-4.240×10^7

2.3 Energy streams

The energy streams define the principal heat and utility duties of the process, as shown in Table 3. The reported duties provide a consistency check on the converged simulation and establish the thermal loads associated with solvent circulation, regeneration and condensation, providing process-level context for the subsequent evaluation of the exchanger. The detailed energy-stream values are provided in Supplementary Table S4.

Table 3: Energy-stream conditions.

Energy stream	Heat flow (kJ/h)
Pump Q (P-100)	4072
Condenser Q (RCY-1)	177.6
Rboiler Q (E-101)	1.448×10^5

2.4 Unit operations

Table 4 summarises the major unit operations. The equipment configuration establishes the process context in which the rich-MEG exchanger is evaluated as a thermal-performance element while remaining integrated with the absorption, pressure-reduction, circulation and regeneration steps of the dehydration cycle. The complete unit-operation inventory is provided in Supplementary Table S5.

Table 4: Unit operations in the MEG dehydration model.

Operation	Type	Feeds	Products	Calc. level
Mixer	Lean from L/R, MEG to Pump	—	500	
FWKO	Separator	Inlet Natural Gas	Water Out, Gas to Contactor	500
T-101	Absorber	Gas to Contactor, MEG Feed	Dry Gas, Rich MEG	2500
VLV-100	Valve	Rich MEG	LP DEG	500
Rich Ethylene Glycol Exch	Heat Exchanger	LP DEG, Regen Feed, Pump Out, MEG to Recy	Regen Feed, Pump Out	500
E-101	Heat Exchanger	Dry Gas	Sales Gas	500
P-100	Pump	MEG to Pump	Pump Out	500
RCY-1	Recycle	MEG to Recy, Regen Feed, Regen Bttms	MEG Feed	3500
MEG Regenerator	Column Sub-Flowsheet	Rblr Q, Cond Q, Regen Feed	Wet Gas	2500

2.5 HYSYS-calibrated operating-domain generation

The dataset used in this work comprises HYSYS-based computational operating points over a predefined operating domain. The converged model provides the process and thermal basis for the exchanger calculation, while the accompanying workbook records the operating conditions and calculated coefficients and implements the generation sequence summarised in Fig. 2, so that the dataset can be reproduced without access to the original HYSYS case. The detailed operating-domain structure is provided in Supplementary Table S7.

Step 1: reference duty. The exchanger duty is taken from the reported hot-side heat-flow change:

$Q = |H_{out} - H_{in}|/3600 = 27.5$ kW; the cold-side change yields the same value, so the two sides are consistent to the precision reported by HYSYS.

Step 2: effective heat-capacity rates. Rather than assuming one heat capacity for both streams, effective values are derived from the reference duty and the reported reference temperatures:

$C_{ph,exx} = 3.3456$ and $C_{pc,exx} = 2.6441$ kJ/(kg.K), giving $C_h = m_h C_{ph,exx} = 0.46848$ kW/K and $C_c = m_c C_{pc,exx} = 0.373997$ kW/K, so that $C_{min} = C_c$ and $C_r = C_{min}/C_{max} = 0.7983$.

Step 3: reference effectiveness. The single calibrated dimensionless parameter of the generator is the effectiveness evaluated at the reference condition:

$\varepsilon = Q/[C_{min}(T_{h,in} - T_{c,in})] = 27.5/[0.373997(193.0 - 30.87)] = 0.4535$.

Because it is derived from the converged HYSYS duty and temperatures, it carries the thermal behaviour of the simulated exchanger into the generated operating domain.

Step 4: operating domain. The two inlet temperatures and the two stream mass flows are varied over the ranges given in Table 6, whereas the exchanger geometry and the tube- and shell-side fouling resistances are held at their design values (Fig. 3). The complete operating-domain and sampling structure is provided in Supplementary Table S7.

Step 5: duty and outlet temperatures. For each generated operating combination, the duty follows from the calibrated effectiveness, $Q = \varepsilon C_{min}(T_{h,in} - T_{c,in})$, and the outlet temperatures follow from the energy balance of each stream, $T_{h,out} = T_{h,in} - Q/C_h$ and $T_{c,out} = T_{c,in} + Q/C_c$.

Step 6: LMTD and target. The terminal differences and the counter-current logarithmic mean temperature difference are evaluated as $\Delta T_1 = T_{h,in} - T_{c,out}$, $\Delta T_2 = T_{h,out} - T_{c,in}$ and $LMTD = (\Delta T_1 - \Delta T_2)/\ln(\Delta T_1/\Delta T_2)$, after which the target follows from the overall heat-transfer equation, $U = Q/(A \times LMTD)$. The equivalent number of transfer units, $NTU = UA/C_{min}$, is computed only as a diagnostic. This treatment should be read as a physics-calibrated surrogate construction rather than a physics-informed one: the physical structure is embedded in the generator and in the target

construction, not in the network, whereas physics-informed models embed governing equations, boundary or initial conditions or residuals directly in the training objective.[16,17]

This construction removes the circularity of a procedure in which a fixed U was assumed to obtain Q, and U was then recalculated from that same Q. In the present generator, Q is fixed by the calibrated effectiveness and the operating temperatures alone, and U is evaluated once from Q and the resulting LMTD. Table 5 collects the reference-condition values obtained with this procedure and reconciles them with the original HYSYS report.

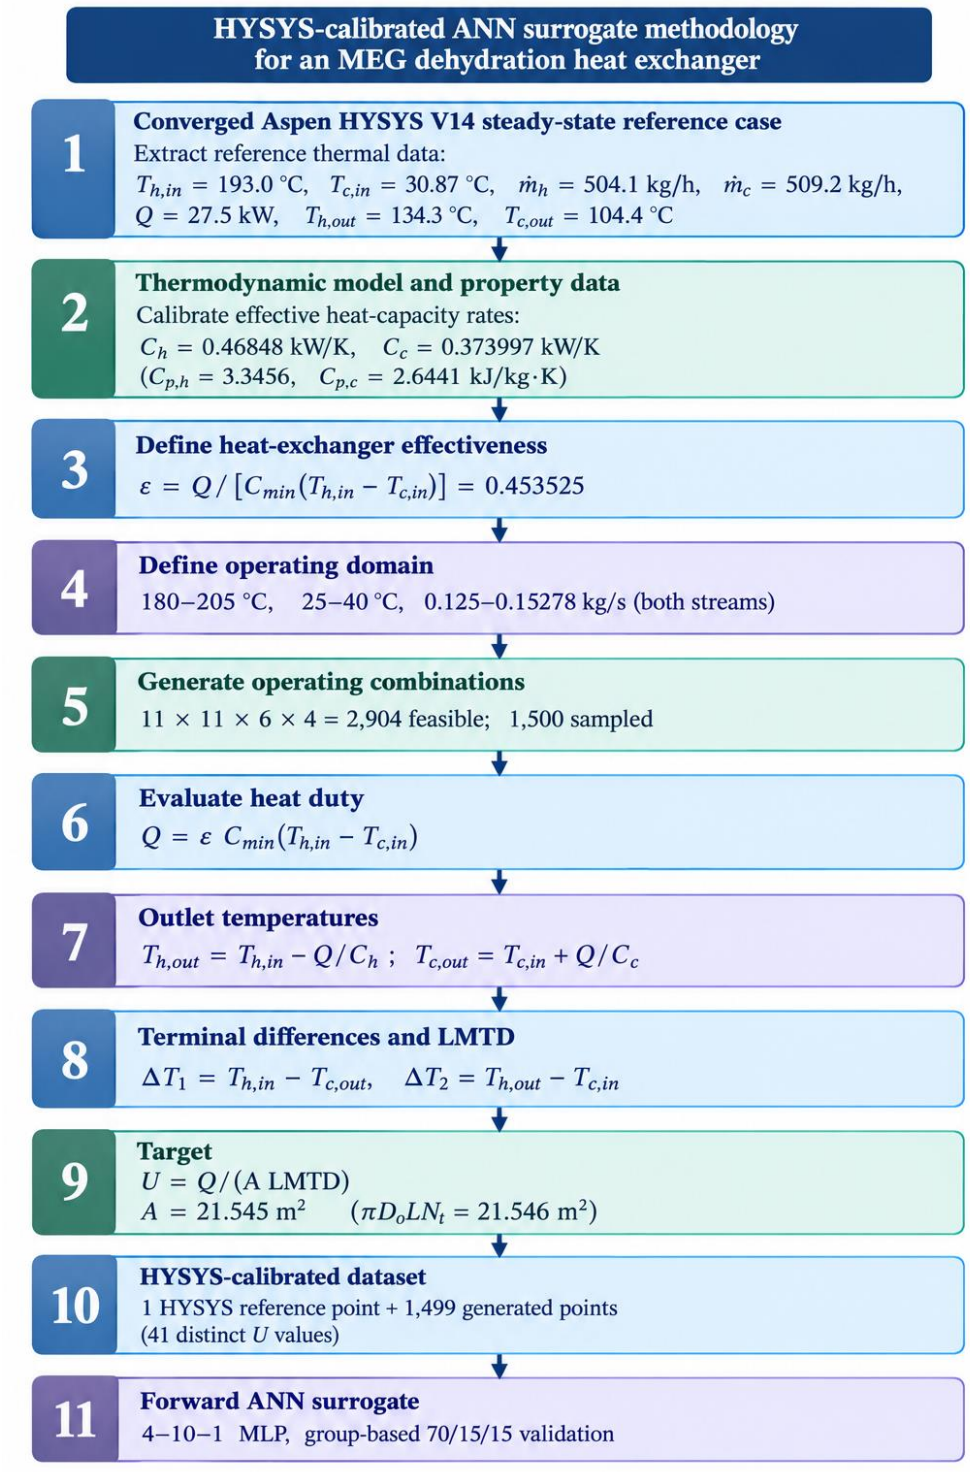


Fig. 2: HYSYS-based operating-domain workflow.

Table 5: Reference thermal basis and overall heat-transfer coefficient.

Parameter	Value	Basis
Heat-capacity rates (C_h/C_c)	0.46848 / 0.373997 kW/K / ($C_{min} = C_c$, $C_r = 0.7983$)	Computed
Reference effectiveness (e)	0.4535	$e = Q/[C_{min} (T_{h,in} - T_{c,in})]$, computed
Terminal differences (dT_1 / dT_2)	88.60 / 103.43 K	Computed from the reported temperatures
LMTD	95.8238 K	Computed from the reported terminal differences
Heat-transfer area A	21.5450 m²	Nominal outer tube area $A = \pi D_o L N_t = 21.5450$ m²
HYSYS-derived target (U)	13.3202 W/(m².K)	$U = Q/(A\ LMTD)$
Equivalent NTU at the reference condition	0.7673	Diagnostic only / $NTU = U A/C_{min}$

Table 6: Operating domain and sampling structure.

Design element	Domain structure	Levels
Hot-inlet temperature (Th,in)	180 - 205 °C	11
Cold-inlet temperature (Tc,in)	25 - 40 °C	11
Hot-stream mass flow mh	0.12500 - 0.152778 kg/s	6
Cold-stream mass flow mc	0.12500 - 0.152778 kg/s	4
Exchanger geometry and fouling	$D_o = 0.01905$ m, $D_i = 0.01575$ m, $L = 3$ m, $N_t = 120$, $D_s = 0.35$ m, baffle spacing = 0.12 m, / $R_{f,tube} = R_{f,shell} = 1 \times 10^{-4}$ m².K/W	constant
Full-factorial size implied by the four levels	11 x 11 x 6 x 4 = 2,904 combinations	-
Realised sampling	100 of the 121 possible ($T_{h,in}$, $T_{c,in}$) combinations, each evaluated at the 15 structured flow settings (one combination at 14 settings and the reference pair at one setting), giving 1,499 grid points	1,499
Reference point	Direct HYSYS solution at $T_{h,in} = 193$ °C, $T_{c,in} = 30.87$ °C, $m_h = 0.14003$ kg/s, $m_c = 0.14144$ kg/s	1
Total dataset	1,500 computational operating points	1,500
Temperature combinations realised	101 distinct ($T_{h,in}$, $T_{c,in}$) pairs	101
Distinct mass-flow combinations	15 structured grid pairs plus the reference pair	16
Distinct target values	16 distinct U values resolved to four or more significant figures; the 49 distinct floating-point entries of the target column are round-off of these 16 values	16
Target range	11.1644 - 15.3447 W/(m².K)	-
Duty range	20.986 - 32.977 kW	-
LMTD range	77.583 - 112.172 K	-
Diagnostic NTU range	0.7278 - 0.8184	-

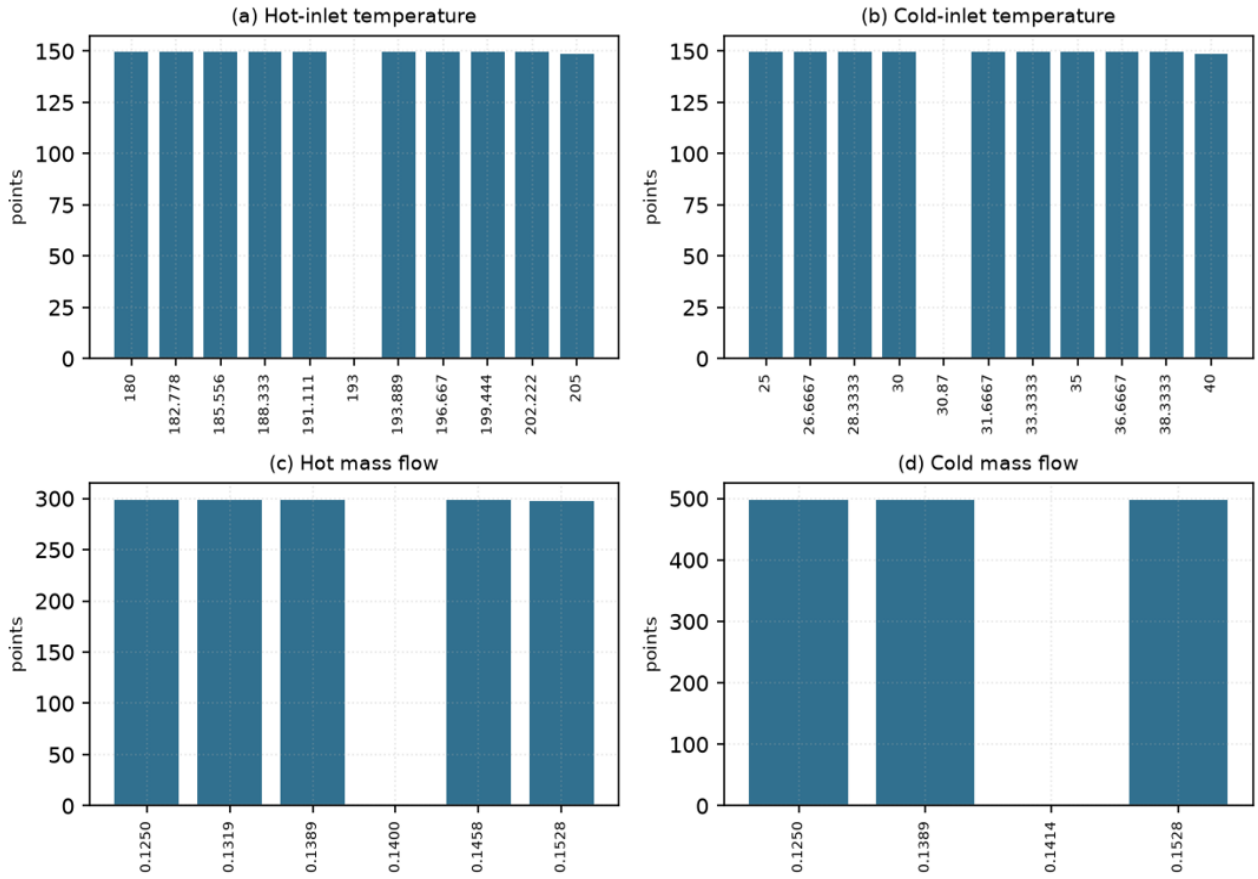


Fig. 3: Operating-domain design.

3. REDUCED-ORDER SURROGATE METHODOLOGY

3.1 Target definition and predictor audit

The target variable is the overall heat-transfer coefficient, U, of the rich-MEG exchanger, defined by the overall heat-transfer equation for the calibrated reference geometry,

$$U = Q / (A \times LMTD) \quad (1)$$

Where: Q is the exchanger duty, A is the nominal outer-tube heat-transfer area, and LMTD is the counter-current logarithmic mean temperature difference. The reference calculation uses the converged HYSYS duty of 27.5 kW and the area A = 21.5450 m², which corresponds to the nominal outer-tube surface $A = \pi D_o L N_t$ of the design basis (D_o = 0.01905 m, L = 3 m, N_t = 120). With the corresponding LMTD = 95.8238 K, this gives the reference value U = 13.3202 W/(m².K), which is retained throughout this work. The detailed thermal-basis values are provided in Supplementary Table S6.

The predictor matrix supplied to the surrogate contains twelve columns, listed in Table 7. Eight are constant across all 1,500 operating points: tube outer diameter, tube inner diameter, tube length, number of tubes, shell diameter, baffle spacing, and the tube- and shell-side fouling resistances. They are retained in the exported workbook only because they define the geometry of the exchanger for which the surrogate is valid; they are not predictors, and MATLAB removes them through the removeconstantrows preprocessing step. Only four columns vary: the hot-stream inlet temperature, the cold-stream inlet temperature and the two stream mass flow rates, so the problem solved by the network is a four-dimensional regression. Two further variables exported in earlier versions of the workflow, the hot- and cold-stream outlet temperatures, are not predictors here: both are evaluated inside the LMTD that produces the target, and including them would constitute target-information leakage. They are used only in the diagnostic model of Section 4.7. Table 7 states the role of every column explicitly. The full predictor

audit is provided in Supplementary Table S8.

Table 7: Predictor variables used in the surrogate model.

No.	Exported column	Unique values	Range or fixed value	Role in the model
1	Hot-inlet temperature, $T_{h,in}$ (°C)	11	180 - 205	Predictor (independent)
2	Cold-inlet temperature, $T_{c,in}$ (°C)	11	25- 40	Predictor (independent)
3	Hot-stream mass flow, m_h (kg/s)	6	0.12500 - 0.152778	Predictor (independent)
4	Cold-stream mass flow, m_c (kg/s)	4	0.12500 - 0.152778	Predictor (independent)
5	Tube outer diameter, D_o (m)	1	0.01905	Constant; defines the domain
6	Tube inner diameter, D_i (m)	1	0.01575	Constant; defines the domain
7	Tube length, L (m)	1	3.000	Constant; defines the domain
8	Number of tubes, N_t	1	120	Constant; defines the domain
9	Shell diameter, D_s (m)	1	0.35	Constant; defines the domain
10	Baffle spacing (m)	1	0.12	Constant; defines the domain
11	Tube-side fouling resistance (m^2.K/W)	1	1×10^{-4}	Constant; defines the domain
12	Shell-side fouling resistance (m^2.K/W)	1	1×10^{-4}	Constant; defines the domain
13	Hot-stream outlet temperature $T_{h,out}$ (°C)	varies	124.4 - 144.5	Excluded: coupled to the LMTD that defines the target (diagnostic model only)
14	Cold-stream outlet temperature $T_{c,out}$ (°C)	varies	95.3 - 110.5	Excluded: coupled to the LMTD that defines the target (diagnostic model only)

3.2 Network architecture and training

The adopted network is a two-layer feed-forward multilayer perceptron with a single hidden layer of ten neurons (Table 8, Fig. 4).[21,22] The hidden-layer transfer function is the hyperbolic tangent sigmoid, and the output layer is linear, so the network realises a smooth, single-output nonlinear mapping of the four effective inputs, which is appropriate for a continuous and smooth target such as U. Inputs and outputs are pre- and post-processed by removing constant rows and applying a min-max normalisation to the interval $[-1, 1]$.[23] The compact architecture was selected to provide sufficient nonlinear representation while keeping the model small enough for rapid surrogate evaluation: the 4-10-1 network contains 61 adjustable parameters (40 input weights, 10 hidden biases, 10 output weights and one output bias) for 1,050 training observations, a parameter-to-observation ratio of approximately 1:17. The effective input dimension is four, so network complexity remains closely aligned with the information-bearing variables of the computational mapping. The complete training configuration is provided in Supplementary Table S9.

Table 8: ANN configuration and training settings.

Setting	Value	Description
Training function	trainlm	Levenberg–Marquardt back-propagation[23]
Division function	divideind	User-defined index division (operating-condition groups)
Division basis	Operating-condition groups	Groups formed from ($T_{h,in}$, $T_{c,in}$) pairs; no group split across subsets
Division mode	sample	Indices are sample indices within the grouped partition
Number of groups	101	Distinct ($T_{h,in}$, $T_{c,in}$) combinations in the 1,500-point dataset
Training/validation/test groups	71 / 15 / 15	Nominal 70 / 15 / 15 of the groups
Training/validation/test points	~1,050 / 225 / 225	Realised partition of the 1,500 operating points
Random seed	rng(42,"twister")	Fixed, for reproducible group assignment
Performance function	mse	Mean-squared error
Inputs evaluated by the network	12 supplied columns	8 removed at runtime by removeconstantrows
Effective input dimension	4	Th,in, Tc,in, mh and mc (see Table 7)
Hidden-layer size	10	Single hidden layer, hyperbolic-tangent-sigmoid neurons
Outputs	1	Overall heat-transfer coefficient U
Initialisation	Nguyen-Widrow	Default MATLAB hidden-layer weight initialisation, under the fixed seed
Maximum epochs	1,000	Prescribed training limit
Early stopping	Not triggered	Training reached 1,000 epochs with zero validation checks
Preprocessing	removeconstantrows, mapminmax	Constant columns removed; remaining inputs and the output scaled to [-1, 1]
Plot functions (enabled)	plotperform, plottrainstate, ploterrhist, plotregression, plotfit	Training diagnostics generated by the trainer

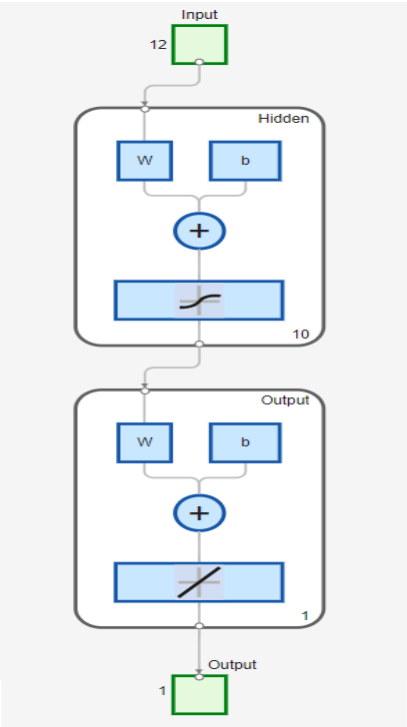


Fig. 4: ANN architecture.

3.3 Data partitioning and target-information control

The data were partitioned by operating condition rather than by individual observation. Observations were grouped according to their ($T_{h,in}$, $T_{c,in}$) combination, and complete groups were assigned to the training, validation and test subsets, so that no temperature combination appears in more than one subset. This group-based scheme is deliberately more conservative than a random sample-by-sample division, because the dataset is a structured grid in which all 15 flow settings of a given temperature pair share the same thermal operating condition. After grouping, the 1,500 operating points form 101 groups: 15 % of the groups were assigned to validation and 15 % to testing, and the remainder to training, giving the nominal 70/15/15 proportions. The resulting partition contains approximately 1,050 training, 225 validation and 225 test operating points, and

the reported metrics are quoted per subset. The random seed is fixed (rng(42, 'twister')) and the group indices are generated once and then held constant, so the partition is reproducible.

Target-information control. Two variables that drive the target calculation, the hot- and cold-stream outlet temperatures, are themselves outputs of the heat balance that produces U, and they enter the LMTD directly. Excluding these target-coupled variables ensures that the primary model represents the forward mapping $U = f(T_{h,in}, T_{c,in}, m_h, m_c)$ for a fixed exchanger configuration and can be evaluated directly from operating conditions. To quantify the effect of this modelling choice, a five-input diagnostic model that additionally receives the two outlet temperatures is trained and reported in Section 4.7. The comparison provides an interpretability check, while the four-input model remains the primary predictive surrogate.

3.4 Implementation and supplementary code

The reported models were implemented in the MATLAB Neural Network Toolbox and independently reproduced in Python with the scikit-learn MLPRegressor,[24] using an identical group partition so that the two implementations can be compared directly. In MATLAB, the input and target data were arranged as matrices with observations in columns; the network was constructed with fitnet (10), and training used the Levenberg-Marquardt algorithm (trainlm), the standard quasi-Newton method for small and medium-sized fitting networks.[25] The preprocessing functions were removeconstantrows and mapminmax, and the data division function was divideind, with the training, validation and test indices supplied explicitly from the operating-condition groups described in Section 3.3. Mean squared error was used as the performance function, and the maximum number of epochs was set to 1,000. Training terminated at the maximum epoch limit with zero validation checks; validation-based early stopping was therefore not the termination mechanism, and the reported validation error is the error at the final epoch rather than a restored best-validation checkpoint. This distinction is stated explicitly because it affects how the validation curve should be interpreted. The implementation computes network predictions and prediction errors and generates the performance, training-state, regression and error-distribution diagnostics used in this study; deployment-related functions remain disabled. An independent Python implementation computes the benchmark models, the repeated-partition statistics and the operating-region hold-out results reported in Section 4, so that the robustness findings do not depend on one software stack. The complete MATLAB training and evaluation script is provided as supplementary material (Listing S2) to support reproducibility.

3.5 Model selection, benchmarking and validation design

An accurate fit to a deterministic computational dataset is not, by itself, evidence that a surrogate is useful, because several model classes can interpolate a smooth low-dimensional grid equally well. The study therefore reports (i) an explicit benchmark family, (ii) repeated group-based resampling and (iii) validation on complete operating regions withheld from training. The benchmark family comprises ordinary least-squares linear regression, second- and third-order polynomial regression with interaction terms, radial-basis-function support-vector regression, and two ensemble tree methods (random forest[26] and gradient boosting). All models use the same four effective predictors and the same grouped partition, so the comparison isolates the model class. Repeated resampling is performed by regenerating the group partition ten times and reporting the mean and standard deviation of the test-set metrics, which separates genuine model performance from the particular realisation of one favourable split. Operating-region hold-out removes six distinct regions of the domain in turn: the two temperature edges of the hot stream, the two edges of the cold stream, and the high-flow and low-flow edges, and trains on the remainder, which tests whether the surrogate interpolates reliably inside a domain whose boundary was never seen. Model accuracy is reported through R^2, RMSE, MAE, MAPE and maximum absolute error rather than a correlation coefficient alone; the combined use of RMSE and MAE follows the discussion of their complementary behaviour as average-error

measures.[27]

4. RESULTS AND DISCUSSION

4.1 Calibrated dataset and reference condition

The calibrated generator reproduces the reference case exactly: with $\varepsilon = 0.4535$, $C_{min} = 0.373997$ kW/K and $(T_{h,in} - T_{c,in}) = 162.13$ K, Step 5 returns $Q = 27.5$ kW, and the resulting LMTD of 95.8238 K together with $A = 21.5450$ m² returns $U = 13.3202$ W/(m².K). The two sides of the exchanger therefore balance to the precision of the HYSYS report, and the reference point is retained unchanged as the first row of the dataset. Over the declared domain, the generated points span a duty of 20.986-32.977 kW, an LMTD of 77.583-112.172 K and an equivalent NTU of 0.7278-0.8184, and the target U ranges from 11.1644 to 15.3447 W/(m²·K), a variation of about 37 % about the reference value. The dataset therefore covers a materially wider range of operating duty than the reference case alone would suggest, and the surrogate is required to interpolate within that range rather than to reproduce an operating point.

One structural property of the generator governs the interpretation of all subsequent results. Because the effectiveness and the two effective heat-capacity rates are constant, the duty is proportional to the driving temperature difference $(T_{h,in} - T_{c,in})$, and both terminal differences that form the LMTD are themselves proportional to the same quantity. The ratio Q/LMTD is consequently independent of the two inlet temperatures and depends only on the mass-flow pair, so the generator produces exactly one target value for each flow combination, that is, 16 distinct values in total. The dataset therefore contains 16 resolvable target values, and the 49 distinct floating-point entries of the target column are round-off of these 16 values. This is a property of the calibrated reduced model and is reported as a limitation in Section 4.2, not as a physical finding.

The reported network was evaluated on a partition formed from operating-condition groups, as described in Section 3.3, so that no $(T_{h,in}, T_{c,in})$ combination is shared between the training, validation and test subsets, and the test set is not a near-duplicate of a training case. Because the target values repeat across a temperature group, this grouped partition provides a conservative assessment of interpolation across operating conditions. The region hold-out analysis in Section 4.8 provides an additional assessment by withholding complete operating regions.

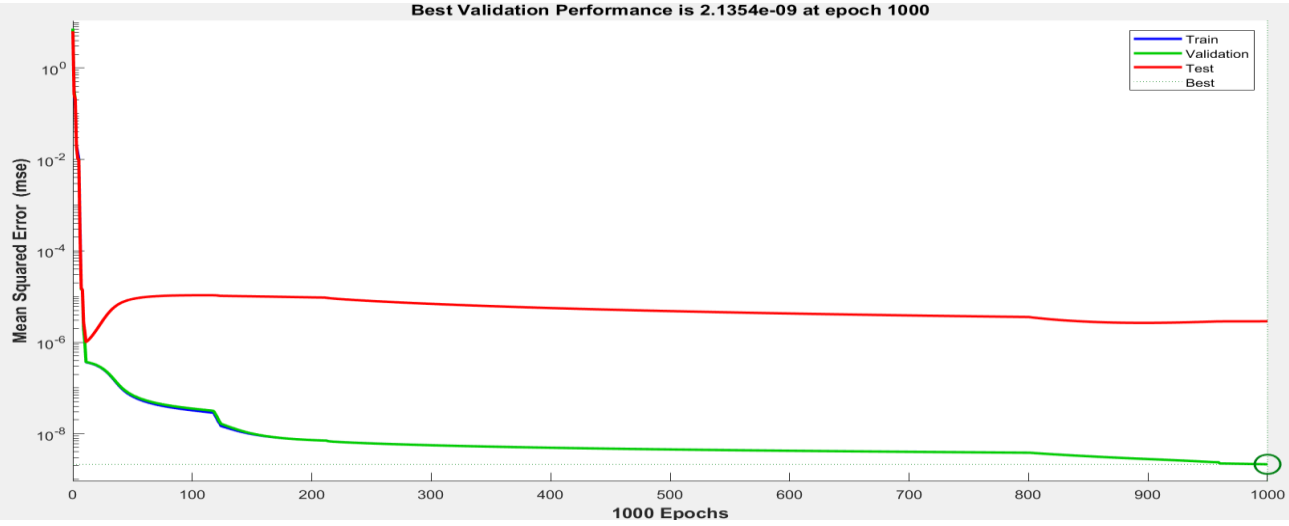


Fig. 5: Training, validation, and test performance.

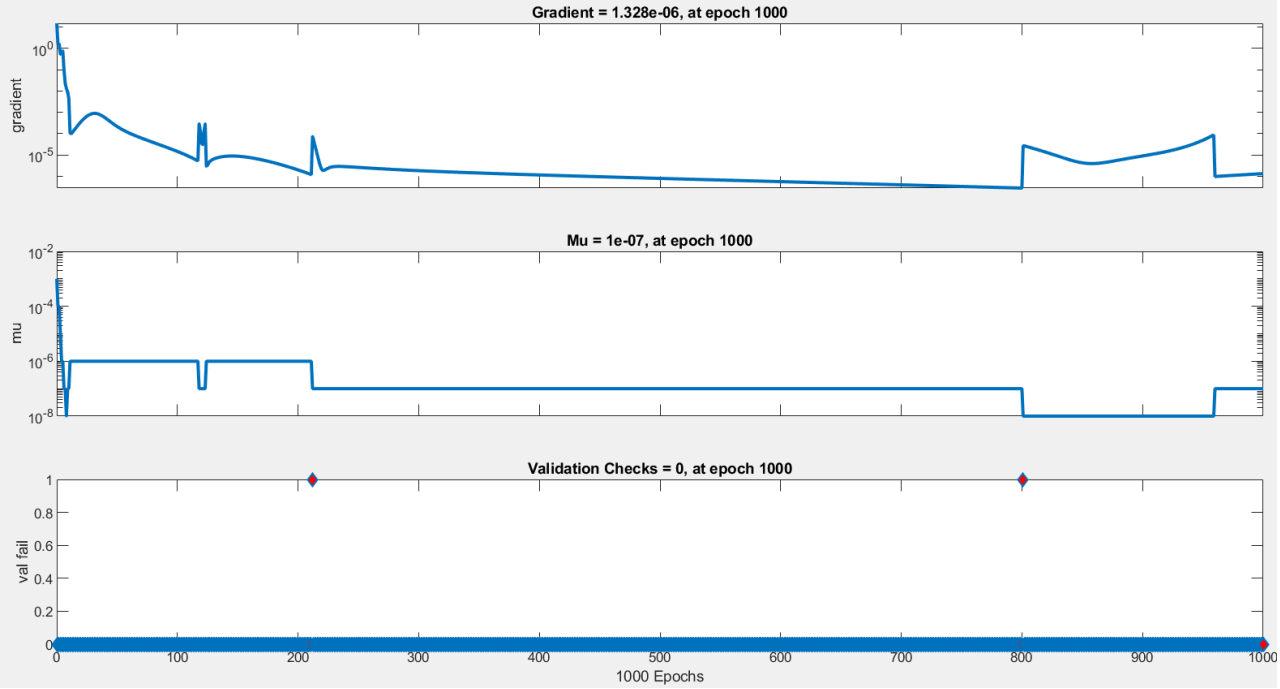


Fig. 6: Training-state variables during optimisation.

The training and validation mean-squared error decreased through optimisation and reached their lowest reported values at the end of the prescribed schedule of 1,000 epochs; the training and validation curves remain close in the later stages, while the test error is slightly higher. The training-state record (Fig. 6) shows that optimisation continued to the prescribed maximum: at the final epoch the performance function had fallen to approximately 2.1×10^{-9}, the gradient magnitude was 1.33×10^{-6}, the step size was of the order of 1×10^{-7}, and no validation checks were triggered. Training therefore terminated at the epoch limit rather than through validation-based early stopping. The quantitative performance of the reported models is summarised in Table 9.

Table 9: Primary model performance.

Model	R^2 (-)	RMSE (W/(m²·K))	MAE (W/(m²·K))	Max. abs. error (W/(m²·K))	MAPE (%)
ANN MLP-4-10-1 (primary forward model)	0.9999990	1.393×10^{-3}	1.153×10^{-3}	4.316×10^{-3}	0.0088

Independently of the MATLAB run, the same four-input mapping was retrained in the Python implementation on the identical grouped partition to provide a reproducible reference result and a basis for the benchmark comparison. That model, MLP-4-10-1, is referred to as the primary forward model. It reproduces the calibrated targets with a test-set R^2 of 0.9999990, RMSE of 1.39×10^{-3} W/(m².K), MAE of 1.15×10^{-3} W/(m².K), MAPE of 0.0088 % and a maximum absolute error of 4.32×10^{-3} W/(m²·K), with a residual bias of -4.4×10^{-5} W/(m².K) (Table 9). The residual bias is small relative to the residual scale. Training, validation and test metrics are closely aligned, consistent with the smooth deterministic target used for surrogate development. Because the target is computational and deterministic, the reported errors quantify reproduction of the calibrated computational relationship rather than independent experimental or plant-level validation. Detailed primary-model performance metrics are provided in Supplementary Table S10.

4.2 Scope of the calibrated target and model limitations

The present model is a physics-calibrated surrogate of a HYSYS-generated target, and its performance should therefore be interpreted within that computational framework. The physical

structure enters the workflow upstream through the calibrated effectiveness ε = 0.4535 and the two effective heat-capacity rates derived from the converged reference case; the network itself is trained to reproduce the resulting calibrated mapping. Accordingly, the reported accuracy characterises agreement with the deterministic computational target. The following scope statements define the applicability of the model and provide context for interpreting the numerical results.

Computational validation basis. The reported accuracy is measured against a HYSYS-calibrated target consisting of one converged reference observation (Q = 27.5 kW, A = 21.5450 m², LMTD = 95.8238 K, U = 13.3202 W/(m².K)) and 1,499 points generated with the calibrated thermal calculation. No plant measurement or experimental dataset enters the calibration. The results therefore characterise the simulated process relationship within the declared computational domain and provide a reproducible basis for subsequent external validation.

Structured target relationship. Because the generator holds the effectiveness, the two effective heat-capacity rates, the exchanger geometry (A = 21.5450 m²) and the tube- and shell-side fouling resistances constant, the duty is proportional to ($T_{h,in} - T_{c,in}$) and both terminal differences that form the LMTD are proportional to the same quantity, so the driving difference cancels in U = Q/(A×LMTD). The target therefore reduces to a flow-only function with exactly 16 distinct values across the 1,500 operating points; the 49 distinct floating-point entries of the target column arise from numerical round-off of these levels. This structure explains the high coefficients of determination and defines the scope of the surrogate as a calibrated reduced-order relationship rather than a general heat-transfer model.

Interpolation within the investigated domain. Across ten repeated group-based partitions, the mean test RMSE is $(1.90 \pm 0.32) \times 10^{-3}$ W/(m².K), and the four temperature-edge hold-outs retain RMSE between 1.95×10^{-3} and 2.20×10^{-3} W/(m²·K). When a complete flow level is withheld, RMSE increases to 2.41×10^{-2} and 2.47×10^{-2} W/(m².K), with residual biases of $+2.14 \times 10^{-2}$ and -2.45×10^{-2} W/(m².K). The target range of 11.1644-15.3447 W/(m².K) is also the range examined in this study. The results support interpolation within the investigated operating region, while predictions beyond the sampled flow range should be treated as extrapolation and checked against a new HYSYS evaluation.

Benchmark interpretation. The benchmark comparison of Section 4.5 indicates that model performance depends on the validation partition. The linear model attains $R^2 = 0.99751$ with RMSE = 0.07026 W/(m².K), while the ensemble tree methods reproduce the training targets closely and achieve comparable mean errors on some partitions. Their RMSE standard deviations (0.0081 and 0.0080 W/(m².K)) are larger than their respective means (0.0038 and 0.0040 W/(m².K)). Under the tested partitions, the ANN provides a combination of low error and relatively consistent performance, supporting its use as the selected surrogate for the investigated computational mapping.

Physics-calibrated rather than physics-informed. A physics-calibrated surrogate reproduces the output of a calibrated simulator-generated target, whereas a physics-informed model embeds governing physics in the training objective or architecture, typically through differential-equation residuals, boundary or initial conditions, or other physical constraints in the loss function.[16,17] In the present workflow, the physical structure resides in the generator and target construction through the effectiveness calibration, stream energy balances and LMTD relation, while the network is trained to match the calibrated targets. The model is therefore described as a physics-calibrated surrogate of a HYSYS-generated target. Future incorporation of governing equations or conservation constraints into the training objective could extend the framework toward a physics-informed formulation.

Future development. The framework can be extended by allowing effective heat capacities and effectiveness to vary with temperature and flow through additional HYSYS runs or a validated property model, by introducing geometry and

fouling resistance as predictors to broaden the applicability domain, and by incorporating measured exchanger duties and temperatures from plant or experimental campaigns to support external validation.

4.3 Error and residual analysis

The prediction-error distribution was examined to assess the magnitude and shape of the residuals. The residual is defined as the difference between the target and the network output; the resulting distributions for the training, validation and test subsets are presented in Fig. 7. The residual distributions are unimodal and centred close to zero, with narrow tails and no pronounced evidence of outlying predictions. The largest absolute residual from the forward model is 4.32×10^{-3} W/(m².K), about 0.03 % of the reference coefficient. A slight left-skew is visible in the test-set distribution, consistent with the small bias reported in Table 9 and indicating a tendency toward positive prediction error for a small number of cases.

The test-set residuals are examined separately in Fig. 8. Almost all residuals fall within ±0.005 W/(m².K) of zero, the largest absolute deviation being about 0.0043 W/(m².K), equivalent to about 0.03 % of the corresponding target value. Because the target values were generated deterministically, these residuals do not represent experimental measurement scatter. Rather, they quantify the estimate error of the network relative to the deterministic forward-model outputs, and they should be interpreted together with the independent hold-out results reported in Section 4.8 rather than as a standalone measure of accuracy. The absence of structure is the relevant finding: a physically meaningful systematic error would appear as a trend of the residual with one of the operating variables, and no such trend is present for the forward model. The small number of larger residuals identify the cases most sensitive to the network architecture and define the numerical tolerance within which the surrogate can be expected to reproduce the calibrated target.

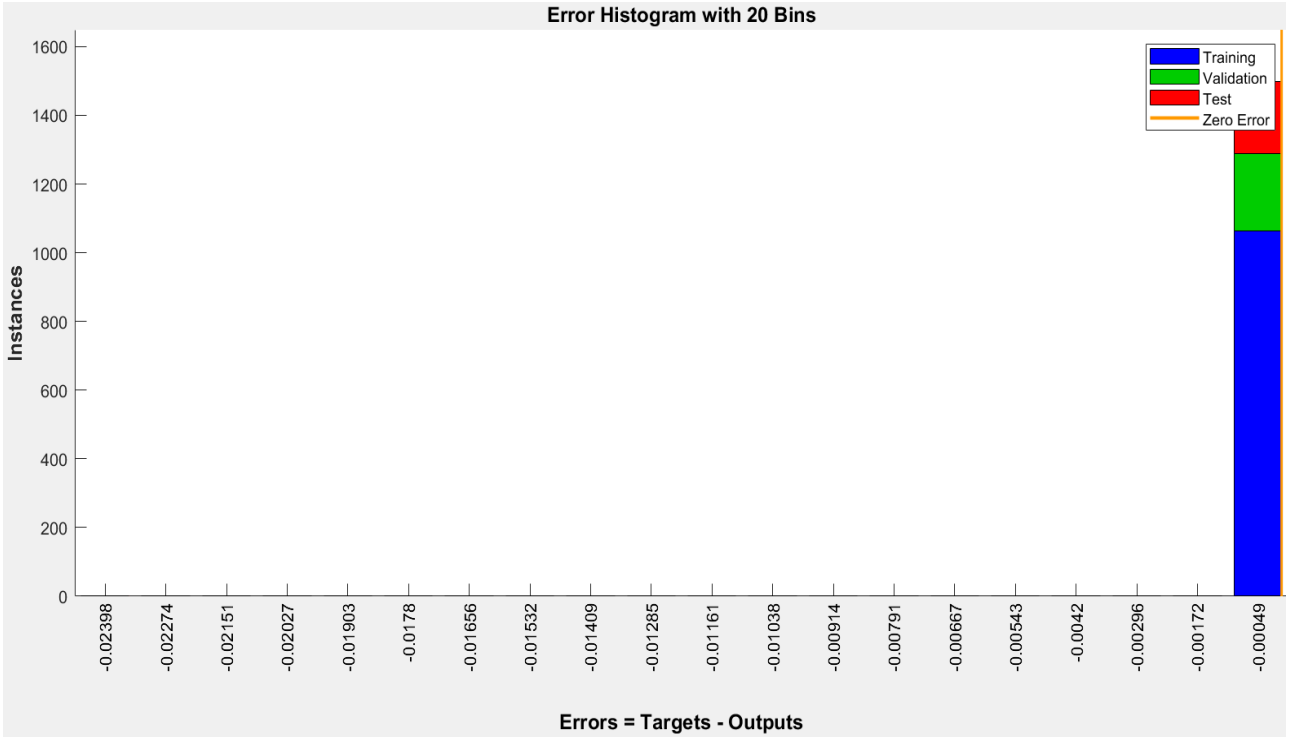


Fig. 7: Prediction-error histogram.

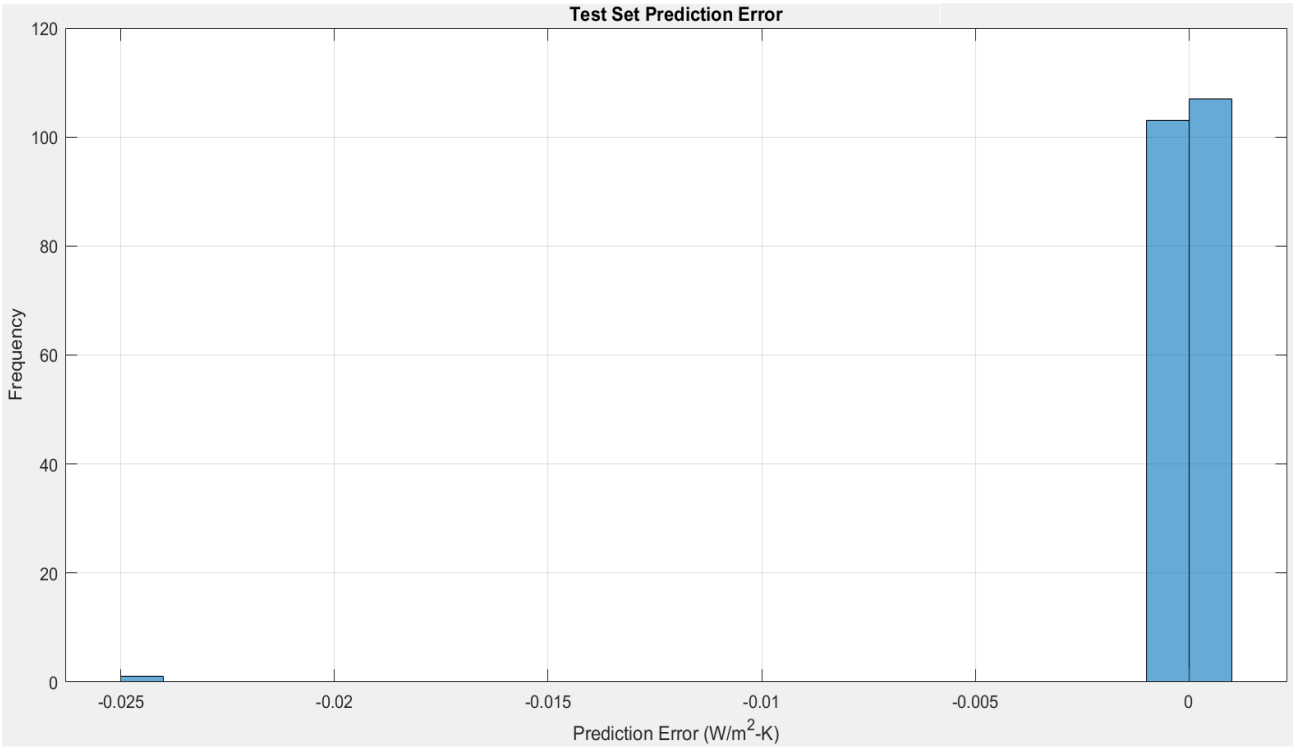


Fig. 8: Test-set prediction-error distribution.

4.4 Regression diagnostics and prediction plots

The regression diagnostics in Fig. 9 show the relationship between the predicted and target values for the training, validation, test and combined subsets. The fitted regression lines closely follow the ideal Y = T line in all four panels, and the corresponding correlation coefficients are very close to unity. However, a correlation coefficient close to unity alone is not sufficient to establish predictive accuracy: a high correlation can coexist with a small systematic offset and does not quantify the magnitude or distribution of the prediction errors. The regression plots are therefore used primarily as a visual assessment of the agreement between predicted and target values and of any apparent nonlinearity or systematic deviation from the ideal relationship, while quantitative model performance is assessed using the metrics reported in Table 9.

The direct comparison of predicted and target U-values for the test subset is shown in Fig. 10. The points follow the 45° ideal line over the whole investigated range of about 11.16-15.34 W/(m².K), and the largest visible departure is of the order of the maximum absolute error quoted in Table 9. Figure 11 presents the same information in a form that makes the residual scale explicit by plotting the residuals against the target together with the residual distribution. The residuals remain confined to a narrow band of a few thousandths of W/(m².K), with no discernible systematic trend across the target range, which indicates that the prediction error does not increase appreciably toward either end of the U range. These results demonstrate numerical agreement with the calibrated dataset but do not constitute evidence of experimental accuracy or plant-level generalisation; this distinction is maintained deliberately to define the scope and conditions under which the surrogate can be applied.

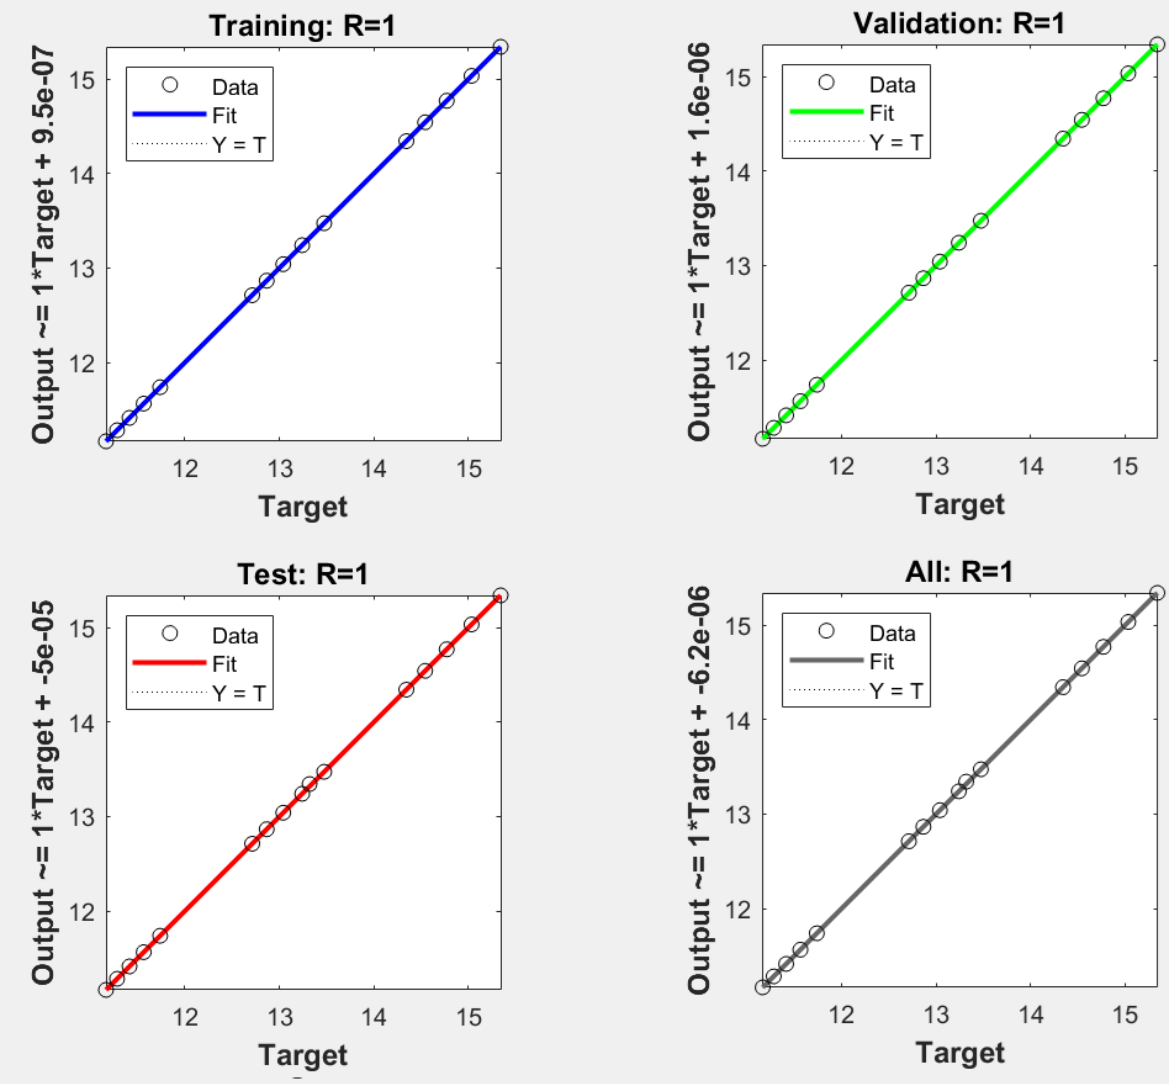


Fig. 9: Predicted versus target U-values across data subsets.

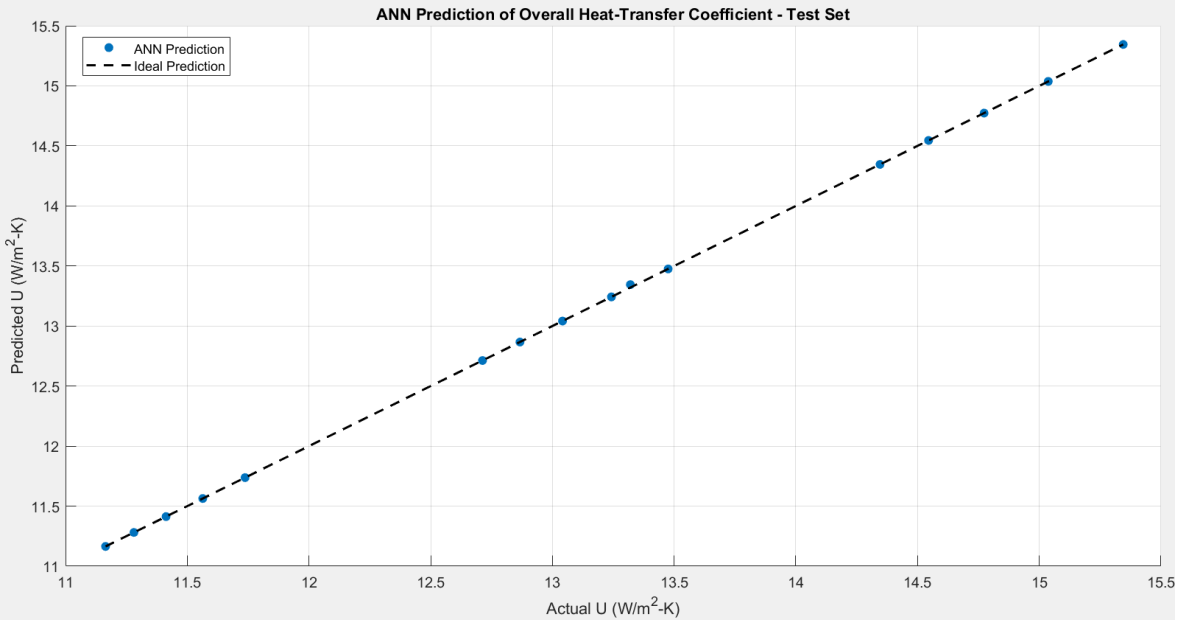


Fig. 10: Direct comparison of predicted and target U-values.

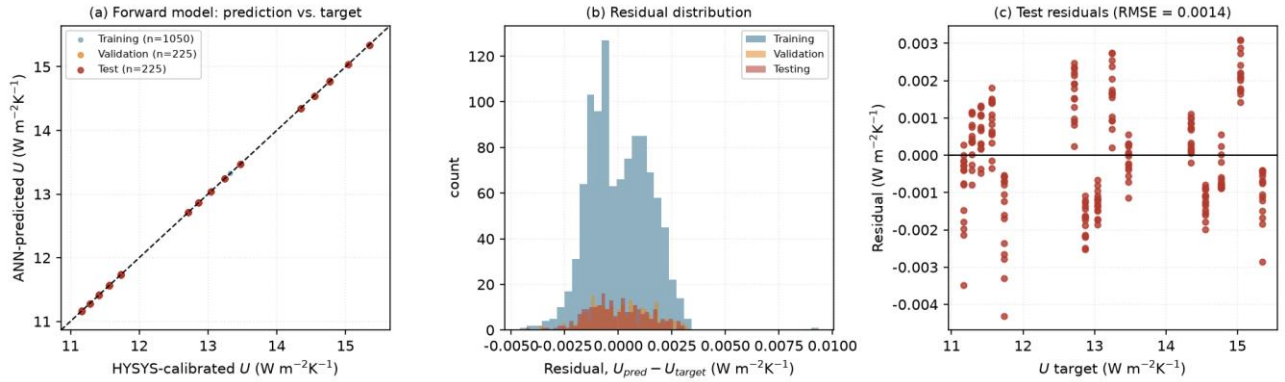


Fig. 11: Predicted versus target U-values and residuals.

4.5 Benchmarking against alternative model classes

Table 10 and Fig. 12 compare the forward model with six benchmark models trained on the identical four effective predictors and evaluated on the identical test groups: ordinary least-squares linear regression, second- and third-order polynomial regression with interaction terms, radial-basis-function support-vector regression, random forest and gradient boosting. This comparison provides a common basis for assessing how different model classes represent the same calibrated computational relationship. The linear model achieves R^2 = 0.99751 with an RMSE of 0.07026 W/(m².K), while the second-order polynomial model gives an RMSE of 0.09429 W/(m².K). The neural network gives lower test-set RMSE on the fixed partition, whereas the ensemble tree methods reproduce the training targets closely and show partition-dependent behaviour. These results indicate that the ANN provides a suitable balance of accuracy and consistency for the investigated surrogate task.

The comparison is more demanding when the partition itself is varied. Table 11 reports the mean and standard deviation of the test-set metrics over ten independent group-based partitions. The network is both the most accurate and the most stable model: $R^2 = 0.9999981 \pm 6 \times 10^{-7}$ and RMSE = $(1.90 \pm 0.32) \times 10^{-3}$ W/(m².K) across the ten partitions. The linear and polynomial models are stable but consistently less accurate by one to two orders of magnitude in RMSE, whereas the random forest and gradient boosting show comparable mean errors but standard deviations that exceed their means (RMSE 0.0038 ± 0.0081 and 0.0040 ± 0.0080 W/(m².K), respectively), because in some partitions the trees place an entire flow level in the test set and cannot interpolate it. The network therefore offers the accuracy of the ensemble methods without their partition sensitivity, which is the practically relevant advantage for a surrogate that is intended to be applied at unseen flow settings. The detailed repeated-partition results are provided in Supplementary Table S11.

Table 10: Benchmark comparison on the grouped test partition.

Model	R^2 (-)	RMSE (W/(m².K))	MAE (W/(m².K))	Max. abs. error (W/(m².K))	MAPE (%)
ANN MLP-4-10-1 (primary forward model)	0.9999990	1.393×10^{-3}	1.153×10^{-3}	4.316×10^{-3}	0.0088
Linear regression (4 terms)	0.9975084	7.026×10^{-2}	5.533×10^{-2}	1.655×10^{-1}	0.4246
Polynomial regression, 2nd order (15 terms)	0.9955127	9.429×10^{-2}	7.142×10^{-2}	3.097×10^{-1}	0.5510
Polynomial regression, 3rd order (35 terms)	0.9896069	1.435×10^{-1}	1.103×10^{-1}	4.559×10^{-1}	0.8528
Support-vector regression, RBF kernel	0.9983470	5.723×10^{-2}	4.887×10^{-2}	1.201×10^{-1}	0.3727
Random forest (400 trees)	1.0000000	1.05×10^{-13}	9.34×10^{-14}	2.20×10^{-13}	7.2×10^{-13}
Gradient boosting	0.99999999	1.489×10^{-4}	1.179×10^{-4}	3.067×10^{-4}	0.00089

Table 11: Repeated group-based partitioning results.

Model	R^2 mean	R^2 s.d.	RMSE mean	RMSE s.d.	MAE mean	Max. abs. error mean
ANN MLP-4-10-1	0.9999981	6.3×10^{-7}	1.90×10^{-3}	3.2×10^{-4}	1.54×10^{-3}	7.9×10^{-3}
Linear regression	0.9975073	3.2×10^{-6}	7.03×10^{-2}	1.6×10^{-5}	5.53×10^{-2}	1.65×10^{-1}
Polynomial regression, 2nd order	0.9957948	5×10^{-4}	9.11×10^{-2}	5.5×10^{-3}	6.90×10^{-2}	3.05×10^{-1}
Support-vector regression, RBF	0.9980685	5.8×10^{-4}	6.13×10^{-2}	8.5×10^{-3}	5.07×10^{-2}	1.52×10^{-1}
Random forest	0.9999627	7.9×10^{-5}	3.84×10^{-3}	8.1×10-^{3}	2.64×10^{-4}	5.6×10^{-2}
Gradient boosting	0.9999627	7.9×10^{-5}	3.96×10^{-3}	8×10^{-3}	3.84×10^{-3}	5.6×10^{-2}

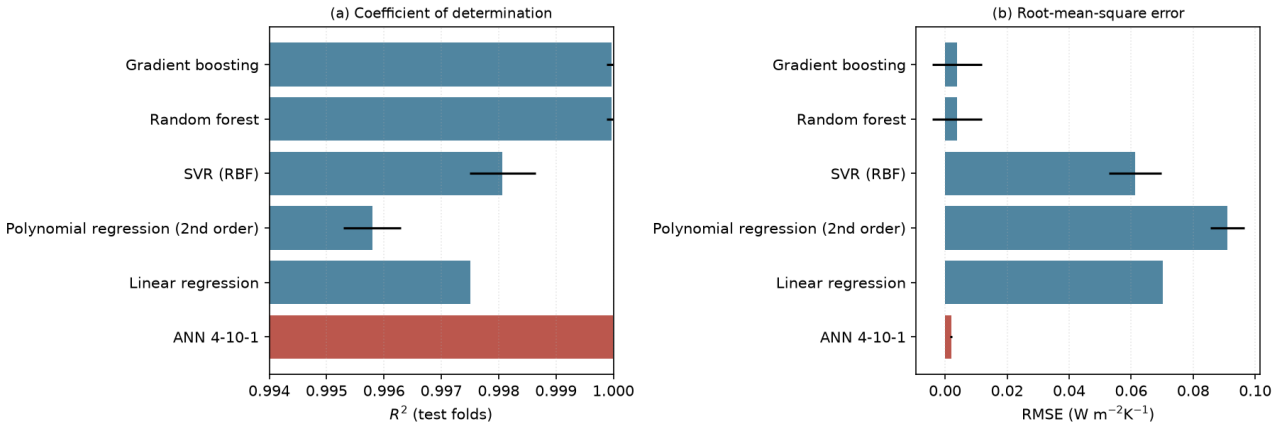


Fig. 12: Benchmark performance over repeated group-based partitions.

4.6 Converged results and the reconciled reference basis

The converged steady-state results provide the process context for the exchanger analysis. Table 12 reports the stream conditions and utility duties transcribed from the converged HYSYS case, and Table 13 reports the resulting reference-condition heat-transfer basis after the reconciliation described in Section 2.5. The two tables define the single reference point from which the entire calibrated dataset is derived, so any inconsistency between them would propagate into all 1,500 operating points. The exchanger that is the subject of this study appears as the Rich Ethylene Glycol Exch, with Regen Bttms as the hot inlet, Lean from L/R as the hot outlet, LP DEG as the cold inlet and Regen Feed as the cold outlet; the corresponding temperature and flow values are those used for the calibration. The complete converged process results are provided in Supplementary Table S12.

The heat-transfer basis is stated once and used consistently throughout. The duty of 27.5 kW is obtained independently from the hot-side and the cold-side heat-flow change of the converged case, both of which return the same value. The area is the nominal outer-tube surface of the design basis, $A = \pi D_o L N_t = 21.5450$ m², and the LMTD computed from the reported terminal differences is 95.8238 K. These three quantities give the reference value $U = 13.3202$ W/(m².K). That value is also the first data point of the surrogate dataset, which makes the reference case and the generated data mutually consistent rather than two separate results, and the reconciled basis links the process simulation, the exchanger duty, the LMTD and the resulting U-value through a single chain of quantities, each of which can be checked against the HYSYS report. At the reference condition, the primary forward model returns $U = 13.3296$ W/(m².K), a deviation of 9.384×10^{-3} W/(m².K), or about 0.0705 %. The reference point lies in the training/validation part of the group-based partition, so the agreement is reported as an internal-consistency check of the

surrogate at the calibrated condition and not as an independent validation result. The detailed reference-condition heat-transfer basis is provided in Supplementary Table S13.

A quantitative comparison with the Kern or Bell–Delaware correlations is deliberately not attempted. Those methods require the tube- and shell-side film coefficients, the wall and fouling resistances and the thermophysical properties of both streams at the operating condition, none of which is reported in the HYSYS case report for this exchanger; a comparison built on assumed values would be a comparison between assumptions rather than between methods. The comparison reported here is therefore made against data-driven benchmarks trained on the identical predictors and the identical partition (Section 4.5), where the input data are common by construction. Computational efficiency is assessed structurally rather than through a wall-clock comparison, because the available benchmark records do not provide identical hardware, software, preprocessing and post-processing conditions across methods. What can be stated is a structural property: once trained, the surrogate evaluates U from four operating variables through 61 multiplications and additions, so a single evaluation is a closed-form arithmetic operation with no iteration and no thermodynamic property call. For workflows that require repeated evaluations of the exchanger within the declared domain, such as sensitivity screening, optimisation loops or plant-wide calculations, this removes the need to re-converge the flowsheet at every evaluation; the practical benefit is therefore expressed as a reduction in the evaluation cost of the calibrated model, not as demonstrated superiority over other modelling approaches. The detailed converged process results are provided in Supplementary Table S12.

Table 12: Converged process results at the reference condition.

Stream/utility	Molar flow (kgmol/h)	Mass flow (kg/h)	P (kPa), T (°C)	Notes/composition
Inlet natural-gas feed	498.1	9183	6205, 29.44	89.89 % CH_4, 2.84 % CO_2, 1.55 % H_2S, 5.91 % C_{2+}
Lean MEG feed (MEG Feed)	8.130	504.5	6205, 48.89	99.96 mol% MEG / Makeup solvent before contactor
Dry Gas (absorber overhead)	497.9	9178	6205, 30.48	0 mol% H_2O / Specification met
Sales Gas (E-101 outlet)	497.9	9178	6171, 35.64	0 mol% H_2O / Pipeline-quality dry gas
Rich MEG (absorber bottoms)	8.289	509.2	6205, 29.43	98.00 mol% / MEG after water pickup
LP DEG (VLV-100 outlet)	8.289	509.2	179.3, 30.87	VF = 0.0136 / Dissolved gases flashed off
Regen Feed (lean/rich exchanger outlet)	8.289	509.2	110.3, 104.4	VF = 0.0184 / Inlet to the MEG Regenerator
Regen Bttms (regenerator lean MEG)	8.122	504.1	103.4, 193.0	99.96 mol% MEG reboiled Lean solvent
Lean from L/R (exchanger hot outlet)	8.122	504.1	102.7, 134.3	99.96 mol% / MEG-cooled lean solvent
MEG to Recy (recycle to MEG Feed)	8.130	504.5	6205, 48.89	99.96 mol% MEG / Loop closure
Reboiler Q / (E-101 kettle reboiler)	—	1.448×10^5 kJ/h		
Condenser Q / (regenerator overhead)	—	177.6 kJ/h		
Pump Q / (P-100 glycol pump)	—	4072 kJ/h		

Table 13: Reference-condition heat-transfer basis.

Parameter	Value	Basis / Interpretation
Heat duty (kW)	27.5	Reference condition
Heat-transfer area (m²)	21.545	Assumed exchanger area
LMTD (K)	95.8238	Calculated from the reported terminal differences
HYSYS-derived U-value (W/(m².K))	13.320	Reference target
ANN-predicted U-value (W/(m².K))	13.330	Agreement with reference target

4.7 Global sensitivity, interpretability and target-information leakage

The sensitivity analysis quantifies how the trained surrogate responds to each of its four effective inputs within the declared domain. The one-factor-at-a-time screening used previously, whose perturbation magnitudes, normalisation and aggregation rules were not fully documented, has been replaced by a global, model-agnostic analysis: permutation importance measured as the increase in test-set RMSE when a single predictor is randomly permuted, and partial-dependence profiles that trace the mean predicted U as one predictor is swept across its sampled range while the others retain their observed values.[28,29] Both measures are computed on the test groups only, so they do not use information from the training partition. Sensitivity of a surrogate is a statement about the model, not about the physics, and the two are compared explicitly at the end of this subsection because the underlying function is known analytically.

The results are collected in Supplementary Table S14 and plotted in Fig. 13. Judged by both measures, only the two mass flows carry information. Permuting the cold-stream mass flow raises the test RMSE by 1.936 W/(m².K), and permuting the hot-stream mass flow by 0.393 W/(m².K), both far exceeding the baseline test RMSE of 1.39×10^{-3} W/(m².K), whereas permuting either inlet temperature changes the RMSE by less than 4×10^{-5} W/(m².K) in absolute value, that is, by less than 3 % of the baseline error. The partial-dependence profiles agree: sweeping the cold-stream flow across its range moves the mean predicted U by 3.377 W/(m².K) and sweeping the hot-stream flow moves it by 0.777 W/(m²·K), while sweeping either inlet temperature moves it by at most 1.6×10^{-3} W/(m².K). The cold stream is the stronger influence because it has the smaller heat-capacity rate at every operating point and therefore sets C_{min}, which controls both the duty and the LMTD in the generator.

The ranking is not merely a property of the fitted network: it is a direct consequence of the generator, and the correspondence can be shown analytically. Because the calibrated effectiveness and the two effective heat-capacity rates are constant, the duty is $Q = \varepsilon C_{min}(T_{h,in} - T_{c,in})$ and both terminal differences are proportional to the same driving difference $(T_{h,in} - T_{c,in})$, with proportionality factors that depend only on the flow rates through C_h/C_{min} and C_c/C_{min}. The driving difference therefore cancels between the numerator and the denominator of $U = Q/(A \times LMTD)$, so the target is exactly independent of the two inlet temperatures and depends only on the mass-flow pair. This is confirmed numerically: all 1,500 operating points collapse onto 16 distinct targets, with a maximum spread of 6×10^{-14} W/(m².K) within a pair, which is floating-point round-off. Sensitivity analysis and analytical inspection of the generator therefore agree, and both identify the mass-flow rates as the only influential inputs (Fig. 13). Two insights follow. First, the temperature invariance of U is a property of the calibrated reduced-order model and provides a clear diagnostic of the assumptions embedded in the generator; a physically richer generator can extend this framework by allowing $C_{p,exx}$ and ε to vary with temperature and flow through additional HYSYS runs or a validated property model. Second, the high R^2 values reflect the structured information content of the target, which comprises

16 distinct levels associated with the flow pairs; the benchmark comparison of Section 4.5 and the operating-region validation of Section 4.8 therefore provide complementary evidence on surrogate performance and robustness.

The influence of target-coupled variables was quantified directly. A second network was trained with five inputs, the four operating predictors plus the hot- and cold-stream outlet temperatures, on the identical group partition. Both outlet temperatures are computed inside the LMTD that produces the target, so the five-input model represents a reconstruction of the target calculation rather than the forward operating-condition mapping and cannot be applied to a new operating condition without first solving the heat balance. The five-input model attains a test R^2 of 0.9999949 and a test RMSE of 3.18×10^{-3} W/(m².K), whereas the four-input forward model attains $R^2 = 0.9999990$ and an RMSE of 1.39×10^{-3} W/(m².K) on the same partition. The four-input model therefore provides the stronger and more direct predictive representation, while the diagnostic comparison confirms that the target-coupled outlet temperatures are unnecessary for the primary surrogate. The corresponding model comparison is reported in Supplementary Table S15.

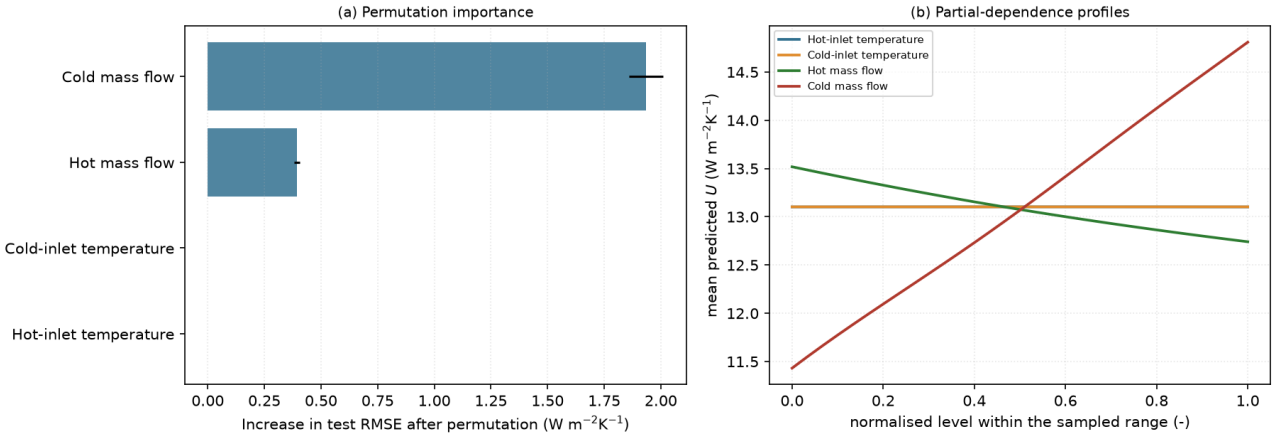


Fig. 13: Global sensitivity of the forward surrogate.

4.8 Validation on independent operating regions

Independent operating-region validation was used to assess interpolation and boundary behaviour within the declared domain. Six complete regions were withheld in turn; the model was retrained on the remaining operating points, and the withheld region was then predicted. The regions are the two temperature edges of the hot stream ($T_{h,in} > 200$ °C and $T_{h,in} < 185$ °C), the two edges of the cold stream ($T_{c,in} < 27$ °C and $T_{c,in} > 38$ °C), and the high-flow and low-flow edges ($m_h \geq 0.1458$ kg/s and $m_h \leq 0.1319$ kg/s). Because the target depends on the flow pair only, the temperature-edge tests primarily assess interpolation in the flow dimension, whereas the flow-edge tests examine behaviour beyond sampled flow levels. The results in Table 16 and Fig. 14 distinguish these two cases. The complete hold-out results are reported in Supplementary Table S16.

For the four temperature-edge hold-outs, the surrogate retains R^2 between 0.9999976 and 0.9999981 with RMSE between 1.95×10^{-3} and 2.20×10^{-3} W/(m².K), comparable to the random group-partition results. For the flow-edge hold-outs, RMSE increases to 2.41×10^{-2} and 2.47×10^{-2} W/(m².K), while R^2 remains above 0.9996, mean biases are $+2.14 \times 10^{-2}$ W/(m².K) for the highest withheld flow levels and -2.45×10^{-2} W/(m².K) for the lowest withheld levels. These results indicate that the surrogate retains strong agreement for withheld temperature regions and provides useful, but less precise, predictions when an influential flow level is outside the training data. For application beyond the sampled flow range, a fresh HYSYS evaluation is recommended to verify the surrogate response. These results are also tabulated in Supplementary Table S16.

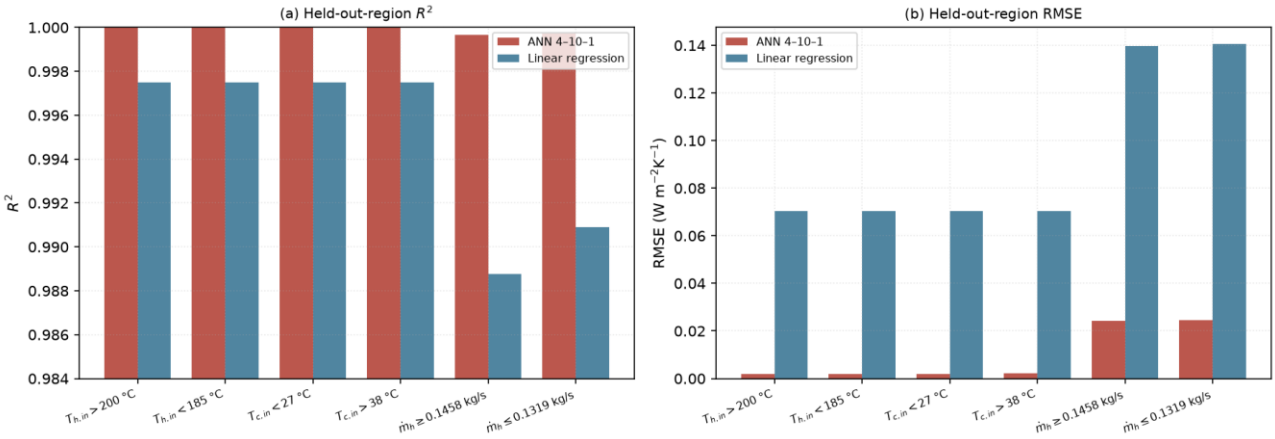


Fig. 14: Independent operating-region validation.

4.9 Implications and deployment scope

The practical value of the surrogate lies in the cost of a single evaluation. Once the reference case has been calibrated, U can be evaluated from four operating variables by a closed-form arithmetic expression, without re-converging the flowsheet and without any thermodynamic property calculation. Within the declared domain and over the sampled flow levels, this evaluation reproduces the calibrated target to about 0.01 %, and it retains R^2 above 0.9996 even when a complete flow edge of the domain is withheld from training. The surrogate can therefore support repeated evaluation of the same exchanger, such as screening studies, sensitivity or uncertainty sweeps, surrogate-assisted optimisation, and as a fast sub-model inside a plant-wide calculation. Because the error grows when a complete flow level is withheld, the surrogate should be used as an interpolator between the sampled flow levels; evaluation outside the sampled flow range is an extrapolation and should be accompanied by a fresh HYSYS evaluation.

The study establishes a defined computational validation framework in which a converged HYSYS reference case is used to calibrate the data-generation procedure, a declared operating domain is sampled, and the resulting surrogate is evaluated using grouped partitions, benchmark models and withheld operating regions. Because the dataset is computational and the exchanger geometry and fouling state are fixed, the resulting applicability domain is specific and reproducible. The framework provides a basis for subsequent external validation using independent operating cases, measured plant or experimental data, and temperature- and flow-dependent calibration. The same workflow can provide a useful template for developing reduced-order surrogates for other simulated unit operations, subject to validation for each new application.

5. CONCLUSIONS

A physics-calibrated reduced-order surrogate of a HYSYS-generated target was developed for the overall heat-transfer coefficient of a rich-MEG exchanger in a natural-gas dehydration unit. A converged Aspen HYSYS V14 model provided the reference basis, and 1,500 computational operating points were generated within the defined domain. Twelve predictor variables were initially supplied, of which eight were fixed geometry and fouling parameters; the information-bearing inputs were therefore the two inlet temperatures and two stream mass flow rates. Under the fixed generator assumptions, the target reduces to a flow-dependent relationship with 16 distinct values. The 4-10-1 neural network reproduced the calibrated targets with $R^2 = 0.9999990$, RMSE = 1.39×10^{-3} W/(m².K), and a maximum absolute error of 4.32×10^{-3} W/(m².K) on the held-out test set. Repeated grouped validation and operating-region hold-outs showed consistent numerical agreement within the investigated domain. Benchmark comparisons also showed lower test errors for the ANN than for the linear and polynomial models on the reported partitions, while the ensemble methods exhibited partition-dependent performance. The workflow is

therefore appropriately characterised as a physics-calibrated surrogate of a HYSYS-generated target rather than a physics-informed model.

Sensitivity analysis identified the two stream mass flows as the dominant predictors under the calibrated generator assumptions. Excluding target-coupled outlet temperatures kept the primary model focused on operating-condition inputs, while the diagnostic model using these variables did not improve the reported test performance. Overall, the results support the use of the ANN as a compact interpolative surrogate within the declared HYSYS-calibrated domain. The principal contribution is a transparent and reproducible surrogate-modelling workflow that combines a converged process model, controlled data generation, leakage-aware input selection, grouped validation, benchmark comparison and operating-region testing. The reported accuracy represents agreement with the computational HYSYS dataset. Further validation with independent plant or experimental measurements, together with broader temperature- and flow-dependent calibration, would support extension of the surrogate to wider applications.

AUTHOR CONTRIBUTIONS

All authors contributed equally for: Conceptualization, methodology, writing – original draft, investigation, data curation, validation, manuscript preparation, critical review, writing – review & editing, and final approval of the manuscript.

CONFLICT OF INTEREST

The authors declare no conflicts of interest concerning the research, authorship, and publication of this article.

FUNDING

This research received no external funding.

DATA AVAILABILITY

The dataset used for surrogate development is provided in the accompanying Excel workbook, which contains the ANN predictor and target matrix, the full heat-balance and LMTD calculation chain for all 1,500 operating points, the HYSYS reference conditions and modelling assumptions, and the documented step-by-step data-generation procedure. The MATLAB training and evaluation script is provided as supplementary Listing S1, and the benchmark, repeated-partition and operating-region hold-out analyses are reproduced with the Python implementation described in Section 3.4. The manuscript documents the effective input dimension (four), the partition function (dividing with operating-condition groups), the fixed random seed and all performance metrics, so that the reported results can be reproduced from the supplied workbook without access to the original HYSYS case. Supplementary Tables S1-S13 provide the detailed tabulations corresponding to Tables 1-9, 11-13. The Supplementary Information file contains Supplementary Tables S1-S13.

REFERENCES

1. S. Mokhatab, W. A. Poe, and J. Y. Mak, Handbook of Natural Gas Transmission and Processing, 4th ed. Cambridge, MA, USA: Gulf Professional Publishing, 2018.
2. A. L. Kohl and R. B. Nielsen, Gas Purification, 5th ed. Houston, TX, USA: Gulf Publishing, 1997.
3. C. O. Wosu, J. G. Akpa, A. A. Wordu, E. Ehirim, and E. M. Ezeh, "Design modification and comparative analysis of glycol-based natural gas dehydration plant," Applied Research, vol. 3, no. 5, 2024.
4. D. J. S. Chong, D. C. Y. Foo, and Z. A. Putra, "A reduced order model for triethylene glycol natural gas dehydration system," South African Journal of Chemical Engineering, vol. 44, pp. 51–67, 2023.
5. D. K. Mohanty and P. M. Singru, "Use of C-factor for monitoring of fouling in a shell and tube heat exchanger," Energy, vol. 36, no. 5, pp. 2899–2904, 2011.
6. J. Zou, T. Hirokawa, J. An, L. Huang, and J. Camm, "Recent advances in the applications of machine learning methods for heat exchanger modeling-A review," Frontiers in Energy Research, vol. 11, no. 1294531, 2023.
7. H. Ezuber and S. M. Zakir Hossain, "A review of corrosion failures in shell and tube heat exchangers: Roots and advanced counteractive," Heat and Mass Transfer, vol. 59, no. 6, pp. 971–987, 2023.
8. A. M. Alsayah, M. J. Alshukri, S. Ali, J. Faraj, and M. Khaled, "Prediction of overall heat

transfer coefficient in concentric tube heat exchangers using artificial neural networks: A comparative study with empirical correlations," International Communications in Heat and Mass Transfer, vol. 163, no. 108723, 2025.

9. B. Z. H. Goh, R. Elyas, A. Hamid, and D. C. Y. Foo, "Artificial neural network modelling of natural gas dehydration process," Discover Chemical Engineering, vol. 4, no. 10, 2024.
10. O. E. Akay and M. Das, "Modeling the total heat transfer coefficient of a nuclear research reactor cooling system by different methods," Case Studies in Thermal Engineering, vol. 25, no. 100914, Jun. 2021.
11. M. Bobič, "Digital twin of heat exchanger," in Heat Transfer—Fundamentals, Enhancement and Applications, S. N. Kazi, Ed. London, U.K.: IntechOpen, 2022.
12. N. M. A. Omar, A. O. A. Amhamed, N. S. E. Zhmoul, L. M. Alhallak, N. A. M. Krir, and M. Milad, "Digital twin-assisted next-generation ceramic membranes for sustainable heavy metal remediation," Chemical Engineering and Processing: Process Intensification, vol. 224, no. 110805, 2026.
13. A. P. Domingos, H. Bispo, T. V. da Costa et al., "Data-driven simulation of crude distillation using Aspen HYSYS and comparative machine learning models," Canadian Journal of Chemical Engineering, vol. 104, no. 8, pp. 4079–4100, 2026.
14. K. Srinivasan, A. Puliyanda, D. Thosar et al., "Artificial intelligence and machine learning at various stages and scales of process systems engineering," Canadian Journal of Chemical Engineering, vol. 103, no. 3, pp. 1004–1035, 2025.
15. T. Danesh, R. Ouaret, P. Floquet, and S. Negny, "Hybridization of model-specific and model-agnostic methods for interpretability of neural network predictions: Application to a power plant," Computers & Chemical Engineering, vol. 176, no. 108306, 2023.
16. M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics, vol. 378, pp. 686–707, 2019.
17. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," Nature Reviews Physics, vol. 3, no. 6, pp. 422–440, 2021.
18. N. V. Queipo, R. T. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, and P. K. Tucker, "Surrogate-based analysis and optimization," Progress in Aerospace Sciences, vol. 41, no. 1, pp. 1–28, 2005.
19. N. M. A. Omar, A. Aboljam, M. Milad, and R. A. Sheliqa, "Optimization of energy efficiency in amine regeneration using a novel side draw flow modification method," African Journal of Advanced Pure and Applied Sciences, vol. 4, no. 1, pp. 41–55, 2025.
20. N. M. Omar, "Simulation and optimization of gas sweetening process at Mellitah gas plant using different blends of amines," University Bulletin, no. 19, pp. 1–20, 2017.
21. M. T. Hagan, H. B. Demuth, M. H. Beale, and O. De Jesús, Neural Network Design, 2nd ed. Stillwater, OK, USA: Martin Hagan, 2014.
22. M. H. Beale, M. T. Hagan, and H. B. Demuth, Neural Network Toolbox User's Guide. Natick, MA, USA: The MathWorks, Inc., 2023.
23. The MathWorks, Inc., MATLAB and Neural Network Toolbox Release R2024b Documentation. Natick, MA, USA: The MathWorks, Inc., 2024.
24. F. Pedregosa, G. Varoquaux, A. Gramfort et al., "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
25. M. T. Hagan and M. B. Menhaj, "Training feedforward networks with the Marquardt algorithm," IEEE Transactions on Neural Networks, vol. 5, no. 6, pp. 989–993, 1994.
26. L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
27. C. J. Willmott and K. Matsuura, "Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance," Climate Research, vol. 30, pp. 79–82, 2005.
28. C. Molnar, Interpretable Machine Learning: A Guide for Making Black Box Models Explainable, 2nd ed., 2022.
29. D. W. Apley and J. Zhu, "Visualizing the effects of predictor variables in black box supervised learning models," Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 82, no. 4, pp. 1059–1086, 2020.