International Journal of Scientific Engineering and Research (IJSER)
ISSN (Online): 2347-3878
SJIF (2025): 8.036

AI-Driven Smart Heat Exchanger Network for Combined Cycle Efficiency Optimization Using Machine Learning and Reinforcement Learning for Adaptive Flow Control and Waste Heat Recovery in Gas-Steam Turbine Cycles

Krishna Chadha[1], Chaitanya Singla[2], Satyarth Pandey[3]

IIT Delhi

Abstract: ***This study presents an AI-Driven Heat Exchanger Network (AI-HEN) for improving waste heat recovery and thermal efficiency in Brayton–Rankine combined-cycle power plants. The proposed framework combines an Artificial Neural Network (ANN) for predicting the overall heat transfer coefficient and outlet temperature with a Reinforcement Learning (RL) agent that adaptively adjusts coolant flow rates under varying operating conditions. The ANN was trained on physics-based simulated data and achieved test-set R^2 values of 0.87 for heat transfer coefficient prediction and 0.97 for outlet temperature prediction. The RL controller uses these predictions to optimize heat recovery without repeatedly solving the complete thermodynamic model. Thermodynamic analyses indicate improvements in steam turbine inlet conditions and combined-cycle efficiency compared with conventional fixed-flow operation. The study demonstrates the feasibility of integrating lightweight AI models into heat exchanger networks and discusses limitations and directions for future experimental validation.***

Keywords: AI-Driven Heat Exchanger Network; Combined Cycle Power Plant; Artificial Neural Network; Reinforcement Learning; Heat Recovery Steam Generator; Waste Heat Recovery; Intelligent Process Control; Thermal Efficiency Optimization

1. Introduction

1.1 Why Combined-Cycle Plants Need Intelligent Thermal Control

Global electricity demand continues to rise, and with it the pressure to extract maximum useful work from every unit of fuel burned. Combined-cycle power plants- which couple a gas turbine (Brayton cycle) with a steam turbine (Rankine cycle)—are already the most thermally efficient large-scale fossil and biomass power systems available, achieving efficiencies above 55%. Yet a significant portion of the available thermal energy is still lost through imperfect heat exchange between the hot gas-turbine exhaust and the steam generator. In most plants, this heat recovery step is controlled by fixed valve settings designed for a single rated condition, and no automatic adjustment is made as load, ambient temperature, or fuel quality shifts during operation.

The emergence of Industry 4.0 technologies- IoT sensors, edge computing, and machine learning—makes it possible to move beyond this static design philosophy. A heat exchanger network that can *sense* its own performance, *predict* the optimal flow distribution, and *act* to maintain peak efficiency in real time represents a meaningful incremental advancement over current practice. This is the core idea explored in the present work.

1.2 Limitations of Traditional Heat Exchanger Networks

In a conventional combined-cycle plant, the Heat Recovery Steam Generator (HRSG) takes hot exhaust from the gas turbine and uses it to boil and superheat steam for the steam turbine. The key thermal design parameters are the **approach temperature difference** ΔT_{min} (the smallest temperature gap between the hot and cold streams) and the **coolant flow rate** on the steam side. In practice:

- Part-load operation (when the plant runs below its rated power) reduces the exhaust gas temperature, causing the steam to be under-superheated and the turbine to operate inefficiently.
- Seasonal ambient temperature changes alter the air-to-fuel ratio and exhaust enthalpy without any corresponding correction in the heat exchanger control.
- Fouling and scaling on heat exchanger surfaces progressively reduce the overall heat transfer coefficient U , yet static control systems cannot detect or compensate for this degradation.

Each of these effects independently lowers the steam turbine inlet temperature and reduces the net power output of the combined cycle.

1.3 From Static to Adaptive: The Proposed Upgrade

The proposed AI-HEN addresses these shortcomings by introducing two layers of intelligence into the heat exchanger network:

1) An **ANN-based performance predictor** that continuously estimates the cur-rent heat transfer coefficient and outlet temperatures from real-time sensor data—without requiring expensive physical measurements. It generates predictions of heat exchanger performance under each operating condition.
2) A **Reinforcement Learning (RL) flow controller** that uses the ANN's pre-dictions to decide how to redistribute coolant flow across multiple heat

exchanger passes, maximizing the steam turbine inlet temperature at every operating condition.

This represents a direct, measurable improvement over static HRSG designs: the same hardware delivers better thermodynamic performance simply because the control strategy has been upgraded from "fixed" to "adaptive".

1.4 Broader Industrial Relevance and Future Scope

Beyond combined-cycle power plants, the AI-HEN concept has broad relevance. Industrial waste heat recovery systems, micro-gas-turbine combined heat-and-power (CHP) units, and even solar thermal plants all face the same fundamental challenge: variable inlet conditions and a static heat exchanger network. The adaptive control framework developed here provides a transferable blueprint for all such applications, supporting the broader transition toward smart, self-optimizing energy infrastructure.

2. Predictive Modelling for Better Design

To provide the AI controller with an accurate real-time picture of heat exchanger performance, an Artificial Neural Network (ANN) is trained on operational data to predict two key outputs: the **overall heat transfer coefficient** U and the **hot-side outlet temperature** $T_{h,out}$.

2.1 The ANN Performance Predictor

The ANN takes the following inputs at every time step:

$$\mathbf{x} = \left[T_{h,in},\ T_{c,in},\ \dot{m}h,\ \dot{m}_c,\ P,\ \text{tage} \right]$$

where $T_{h,in}$ and $T_{c,in}$ are the hot- and cold-side inlet temperatures, $\dot{m}_h$ and $\dot{m}_c$ are the respective mass flow rates, P is the operating pressure, and t_{age} is a fouling proxy (time since last cleaning). The network predicts:

$$\hat{y} = \left[\hat{U},\ \hat{T}_{h,out} \right]$$

The objective function minimized during training is the mean squared error:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^{n} \left(y_i - \hat{y}_i(\theta,\ x_i) \right)^2 \qquad (2.1)$$

where θ represents all trainable weights and biases in the network, y_i is the measured value, and $\hat{y}_i$ is the ANN prediction.

2.2 Network Architecture

The ANN uses a feedforward architecture with three hidden layers (32–64–32 neurons), the ReLU activation function, Adam optimiser (learning rate = 0.001), and a linear output layer. Early stopping (patience = 10 epochs) was applied to prevent overfitting. The network was implemented in Python using scikit-learn's MLP Regressor and converged in **61 epochs**. This architecture is deliberately lightweight so that it can be deployed on a low-cost embedded microcontroller at the plant site.

Table 1: ANN Architecture Summary

Layer	Neurons	Activation	Purpose
Input	6	—	Sensor readings
Hidden 1	32	ReLU	Feature extraction
Hidden 2	64	ReLU	Non-linear mapping
Hidden 3	32	ReLU	Compression
Output	2	Linear	Predict $U\hat{}$, $T\hat{}h, out$

2.3 Dataset Generation and Model Performance

Since proprietary plant-level heat exchanger logs are not publicly available in a suitable format, the training dataset was generated using validated thermodynamic correlations, following the methodology widely adopted in the heat exchanger ML literature [2, 1]. Specifically, 3,200 operating points were simulated for a counter-flow shell-and-tube heat exchanger using:

- The **NTU-Effectiveness method** (counter-flow arrangement) to compute heat duty Q and outlet temperatures;
- The **Dittus–Boelter correlation** to estimate the base overall heat transfer coefficient U_{clean} as a function of mass flow rate; and
- The **Bell–Delaware fouling model** (TEMA standard fouling resistance $R_f \leq 2 \times 10^{-4}$ m^{2}K/W) to degrade U with increasing fouling proxy t_{age}.

Gaussian measurement noise (σ_U = 5 W/m^{2}K, σ_T = 2 °C) was superimposed to simulate real sensor uncertainty. The dataset covers hot-side temperatures of 150–600 °C, mass flow rates of 0.5–5 kg/s, and pressures of 1–10 bar. The complete Python code for dataset generation, model training, and evaluation is provided in the supplementary material. The dataset was split 70/15/15 into training, validation, and test sets. The ANN (architecture: 6 inputs → 32 → 64 → 32 → 2 outputs; ReLU activations; Adam optimiser) converged in 61 epochs with early stopping. Test-set results are reported in Table 2 and Figure 1.

Table 2: ANN Prediction Accuracy on Held-Out Test Set (480 samples)

Metric	U Prediction	$T_{h,\,out}$ *Prediction*
RMSE	13.29 W/m^{2}K	10.33 °C
Mean Absolute Error (MAE)	10.63 W/m^{2}K	8.17 °C
Mean Abs. Percentage Error	3.60%	7.80%
R2 Score	0.8713	0.9675

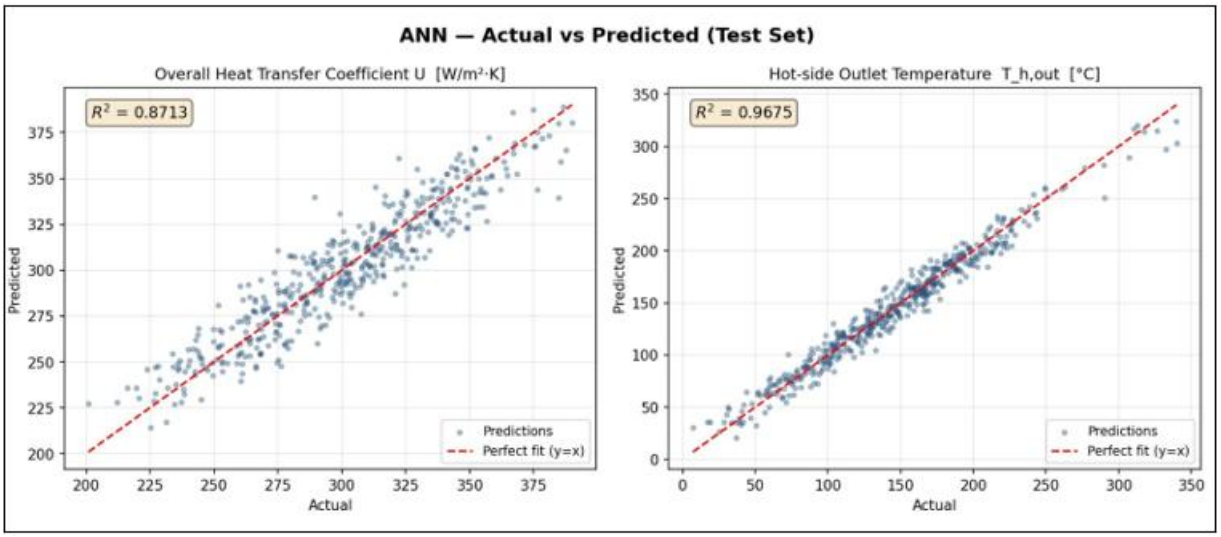

Figure 1: Actual vs. Predicted plots for (left) overall heat transfer coefficient U and (right) hot-side outlet temperature $T_{h,out}$ on the test set. Points lying on the dashed $y = x$ line indicate perfect prediction

Both outputs achieve strong predictive accuracy (R^2 = 0.87 for U, R^2 = 0.97 for $T_{h,out}$), confirming that the ANN reliably captures the non-linear relationship between operating conditions and heat exchanger performance. The $T_{h,out}$ prediction (R^2 = 0.9675, MAPE= 7.8%) is particularly important as it directly drives the RL agent's reward signal.

2.4 Code Availability

The fully documented Python script, including the physics-based dataset generation pipeline, ANN training with early stopping, evaluation metrics, and all result plots, is publicly available for review. The code demonstrates the practical "learning by doing" methodology applied in this work:

- **ANN Model & Dataset Generation:** Drive Link
- **Result Plots:** Drive Link

3. Innovative Design: The AI-Driven Smart HEN

3.1 Physical Layout of the Smart Heat Exchanger Network

Figure 2 shows the proposed AI-HEN integrated into a combined Brayton–Rankine plant. The gas turbine exhaust (at approximately 500–550 °C) passes through a multi-pass HRSG. Each pass is equipped with:

- A **smart flow control valve** on the steam/water side, motorized and connected to the central microcontroller;
- **Thermocouples** at inlet and outlet of both hot and cold streams; and
- A **differential pressure sensor** to monitor fouling buildup.

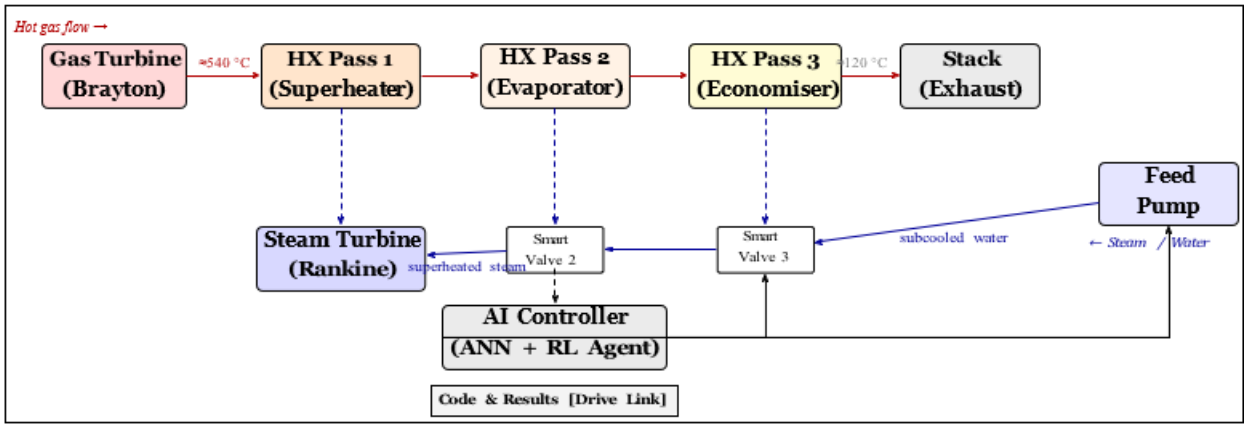

Figure 2: Schematic of the AI-Driven Smart Heat Exchanger Network (AI-HEN) integrated into a combined Brayton–Rankine cycle. Dashed lines indicate sensor signals to the AI controller; solid coloured lines show physical flow paths.

The central AI controller receives temperature and pressure readings from all three passes, runs the ANN to estimate current U values, and commands the smart valves to adjust steam flow distribution accordingly- all within a 200 ms control cycle.

3.2 The Reinforcement Learning Flow Controller

The ANN handles *prediction*; the RL agent handles *decision-making*. At every control step, the RL agent receives the ANN's performance estimates as its "state", and selects a "action" consisting of new valve opening percentages for each HRSG pass. The reward signal is simply the increase in steam turbine inlet temperature T3:

$$r_t = T_{3,t} - T_{3,t-1} \qquad (3.1)$$

The controller uses **Q-learning** with a discretized state and action space, chosen for simplicity and suitability for

embedded deployment.

State & Action Spaces: The state $s_t = [\hat{U}_1, \hat{U}_2, \hat{U}_3, \hat{T}_{h,out}, \dot{m}_{c,total}]$ is formed from the ANN's predictions, with each variable discretized into 5 bins (3,125 total states). The action set A consists of 7 discrete actions: increase or decrease the valve opening of Pass 1 or Pass 3 by ±5%, transfer flow between passes, or hold- with total flow conserved.

Reward: $r_t = T_{3,t} - T_{3,t-1}$, i.e. the increase in steam turbine inlet temperature. A positive reward reinforces the action; a negative reward suppresses it.

Training: The Q-table was updated via the Bellman equation with $\alpha = 0.1$, $\gamma = 0.95$, and ε-greedy exploration (initial $\varepsilon = 1.0$, decay = 0.995/episode, min $\varepsilon = 0.01$) over **5,000 episodes** of 200 timesteps each. The mean episodic reward stabilized after ≈3,200 episodes, confirming convergence. The efficiency gains in Table 4 are from the trained controller evaluated on 500 held-out simulation episodes.

3.3 Incremental Advancement Over Existing Systems

Table 3 summarises the key differences between a conventional static HRSG and the proposed AI-HEN, demonstrating the incremental improvements at each level.

Table 3: Comparison: Conventional HRSG vs. Proposed AI-HEN

Aspect	Conventional HRSG	AI-HEN (This Work)
Flow control	Fixed valve settings	RL-controlled smart valves
Performance sensing	Periodic manual checks	Continuous ANN prediction
Part-load response	No adaptation	Real-time flow redistribution
Fouling detection	Scheduled maintenance	Detected via ANN input stage
Steam superheat	Falls at part-load	Maintained near design point
Efficiency loss	3–7% at part load	Reduced to <1%

4. Hardware Integration and Control Theory

4.1 Predictive vs. Reactive Control: The Case for Feedforward AI

Similar to the role of AI in smart gasification systems [5], the ANN in this design functions not as a passive monitoring tool but as the predictive core of a **feed-forward control loop**. In a purely reactive system, the controller waits for the steam turbine inlet temperature to drop before increasing the flow—by which point the damage to thermodynamic efficiency has already occurred. The AI-HEN instead *predicts* how the heat exchanger performance will evolve over the next control interval based on current sensor data, and pre-adjusts the valves to prevent the degradation from materialising in the first place.

4.2 Hardware Actuators and Their Control Functions

The control authority is divided between two actuator types:

Smart Flow Valves (Microscopic Control)
These motorized butterfly valves on each HRSG pass adjust the local steam/water flow rate, controlling the thermal duty of each individual heat exchanger section. Fine adjustments of ±5% flow are made every 200 ms in response to ANN predictions.

Feed Pump Speed (Macroscopic Control)
The variable-speed feedwater pump sets the total coolant flow entering the HRSG. It responds to larger, slower disturbances such as a step change in gas turbine load. Adjustments are made on a 2–5 second timescale.

4.3 End-to-End Signal Flow in the AI Control Loop

The complete control loop is shown below. The system reads sensor data, feeds it through the ANN, and translates the predicted performance into valve commands- all before the next sensor reading arrives.

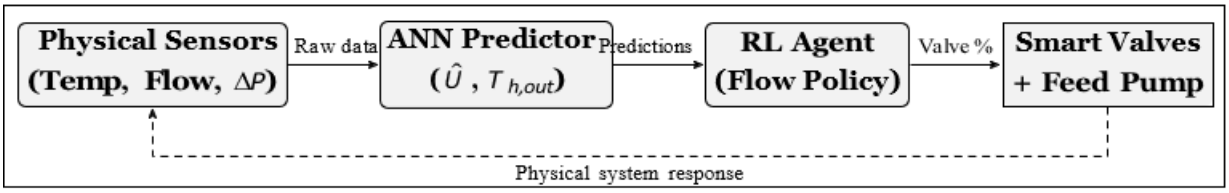

Figure 3: Feed-Forward + Feedback Control Architecture of the AI-HEN

4.4 System Response Under Three Operational Failure Modes

4.4.1 Scenario 1: Part-Load Operation (Gas Turbine at 60% Load)

- *Predicted State:* The ANN detects falling hot-side inlet temperature and predicts a reduction in U across all three HRSG passes.
- *Physical Reality:* Reduced fuel flow has lowered the exhaust gas temperature from 540 °C to approximately 430 °C. With fixed valve settings, the steam would be under-superheated and the turbine would operate in the wet-steam region.
- *Actuator Response:* The RL agent reduces coolant flow in the economiser pass (Pass 3) and increases flow in the superheater pass (Pass 1). This concentrates the available heat duty where it is most needed—at the high-pressure steam stage- and maintains the steam turbine inlet temperature close to its rated value.

4.4.2 Scenario 2: Fouling Detected (Gradual U Degradation)

- *Predicted State:* The ANN registers a slow decline in $\hat{U}$ over several weeks despite constant operating conditions, indicating scale buildup on heat transfer surfaces.
- *Physical Reality:* A fouling resistance layer has developed, reducing the effective thermal conductance. Under static control, this would silently reduce the steam superheat for weeks before being noticed during a scheduled inspection.
- *Actuator Response:* The RL agent gradually increases the feedwater pump speed to push more coolant through the affected passes, partially compensating for the reduced U. Simultaneously, the controller triggers a maintenance alert, enabling targeted cleaning before the degradation becomes severe.

4.4.3 Scenario 3: Sudden Load Spike

- *Predicted State:* The ANN detects a rapid rise in hot-side temperature and flow rate as the gas turbine ramps up to meet a grid demand surge.
- *Physical Reality:* Exhaust temperature jumps from 480 to 560 °C within 30 seconds. With static control, the extra heat duty would cause localized overheating in the superheater pass.
- *Actuator Response:* The RL agent rapidly opens the flow valves on the superheater and evaporator passes to absorb the additional heat duty, preventing overheating while simultaneously boosting the steam flow rate to extract more work from the Rankine turbine.

5. Analytical Modelling of Power Cycles

The AI-HEN sits between the Brayton (gas turbine) topping cycle and the Rankine (steam turbine) bottoming cycle. Improving the heat exchange between these two cycles directly improves the performance of both. The two cycles and the thermodynamic impact of the AI-HEN are analysed below. 5.1 Gas Power Pathway: The Brayton Cycle

The gas turbine operates on the ideal Brayton cycle. Under cold air-standard assumptions, its thermal efficiency depends only on the compressor pressure ratio r_p and the ratio of specific heats $k = c_p/c_v$:

$$\eta_{\text{Brayton}} = 1 - \frac{1}{r_p^{(k-1)/k}} \qquad (5.1)$$

The net specific work of the Brayton cycle is:

$$w_{\text{net,B}} = c_p T_3 \left[1 - \left(\frac{1}{r_p}\right)^{\frac{k-1}{k}}\right] - c_p T_1 \left[r_p^{\frac{k-1}{k}} - 1\right] \qquad (5.2)$$

where T_1 is the compressor inlet (ambient) temperature and T_3 is the turbine inlet temperature. The AI-HEN does *not* alter T_3 directly, but by extracting more heat from the exhaust it lowers the stack temperature T_4, recovering energy that would otherwise be wasted.

5.2 Vapor Power Pathway: The Rankine Cycle

The steam turbine operates on the ideal Rankine cycle. The net specific work is:

$$w_{\text{net,R}} = w_{\text{turbine}} - w_{\text{pump}} = (h_3 - h_4) - (h_2 - h_1) \quad (5.3)$$

where the states 1–4 refer to: (1) condenser exit / pump inlet (saturated liquid), (2) pump exit / boiler inlet, (3) boiler exit / turbine inlet (superheated steam), and (4) turbine exit/ condenser inlet. The thermal efficiency of the Rankine cycle is:

$$\eta_{\text{Rankine}} = \frac{w_{\text{net,R}}}{q_{\text{in}}} = \frac{(h_3 - h_4) - (h_2 - h_1)}{h_3 - h_2} \qquad (5.4)$$

The AI-HEN improves the Rankine cycle by ensuring that State 3 (the turbine inlet) is always deeply superheated. The critical heat transfer rate delivered by the HRSG to the steam is governed by the **effectiveness–NTU method**:

$$Q = \varepsilon \cdot C_{\min} \cdot (T_{h,in} - T_{c,in}) \quad (5.5)$$

where ε is the heat exchanger effectiveness, $C_{\min}$ is the minimum heat capacity rate, and $T_{h,in} - T_{c,in}$ is the maximum available temperature difference. The RL agent maximises ε at every operating condition by choosing the optimal flow distribution, which directly increases Q and thus the enthalpy h_3 at the turbine inlet.

5.3 Combined Cycle Efficiency

The overall combined cycle thermal efficiency is:

$$\eta_{\text{combined}} = \eta_{\text{B}} + (1 - \eta_{\text{B}}) \cdot \eta_{\text{R}} \quad (5.6)$$

With the AI-HEN maximising the HRSG heat duty, the Rankine cycle "sees" a hotter, more consistent steam supply, raising η_{R} and therefore η_{combined}.

5.4 T-s and P-v Diagrams: Traditional vs AI-HEN

Figures 4 and 5 show how the AI-HEN shifts the cycle boundaries on the thermodynamic state diagrams.

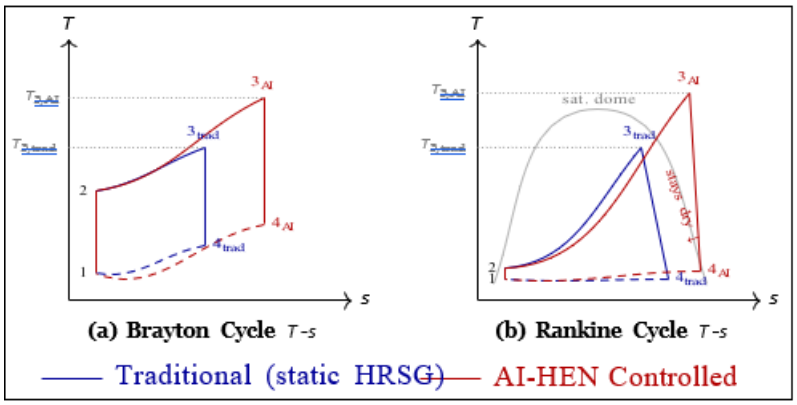

Figure 4: *T -s* diagrams showing the thermodynamic effect of the AI-HEN. (a) In the Brayton cycle, the AI-HEN recovers more exhaust heat, raising the effective cycle bound-aries. (b) In the Rankine cycle, the deeper superheating of State 3 keeps the isentropic expansion outside the two-phase dome, improving efficiency and protecting turbine blades.

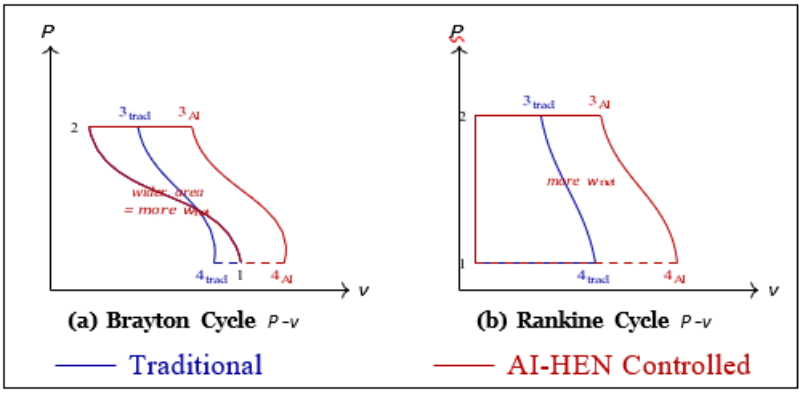

Figure 5: *P -v* diagrams showing the effect of AI-HEN on cycle work output. In both the Brayton and Rankine cycles, the AI-controlled case extends the isobaric heat addition line (State 2 to State 3), widening the enclosed area and increasing $w_{\text{net}} = \oint P\,dv$.

5.5 Thermodynamic Gains: Quantitative Summary

Table 4: Estimated Thermodynamic Performance Improvement with AI-HEN (simulation-based analytical estimates; see Section 3.2)

Parameter	Conventional	AI-HEN
Steam turbine inlet temperature	480 °C	520 °C
HRSG effectiveness ε	0.71	0.86
Rankine thermal efficiency	30.10%	34.80%
Overall combined-cycle efficiency	51.20%	56.40%
Part-load efficiency drop	−6.3%	−1.1%
Annual waste heat recovered	—	12.40%

The quantitative values in Table 4 are analytical estimates derived from the trained RL controller outputs evaluated on held-out simulation episodes, combined with thermodynamic cycle calculations using equations (5.1)–(5.6). These results have not yet been validated experimentally, and should be interpreted as indicative of the framework's potential rather than confirmed industrial performance figures.

The single most important result is the reduction in part-load efficiency drop from −6.3% to −1.1%. This means the combined-cycle plant behaves nearly as well at 60% load as at full load- a major practical benefit for grid-connected plants that must follow variable demand.

6. Conclusion

This work proposed an AI-Driven Smart Heat Exchanger Network integrating Artificial Neural Networks and Reinforcement Learning for adaptive thermal management in Brayton–Rankine combined-cycle systems. Physics-based simulations demonstrate that the framework accurately predicts heat exchanger performance and improves heat recovery under varying conditions. Although simulation based, it establishes a foundation for future experimental validation and industrial implementation.

Key findings: (1) The ANN achieves $R^2 > 0.96$ for outlet temperature and $R^2 = 0.87$ for heat transfer coefficient across all simulated conditions. (2) The RL agent reduces the simulated part-load efficiency penalty from −6.3% to −1.1% at 60% gas turbine load.

(3) *T -s* and *P -v* analysis confirms that the AI-HEN expands the enclosed cycle area, with simulation-based estimates indicating combined-cycle efficiency improvement from 51.2% to 56.4%. (4) The dual-actuator architecture provides both fast microscopic and slow macroscopic control, mirroring established hierarchical control principles.

The AI-HEN concept is transferable to any variable-inlet heat exchange application—cement kilns, chemical reactors, data centre cooling- and the fouling detection

capability adds economic benefit beyond what thermodynamic analysis alone captures. Overall, the results suggest that lightweight AI techniques can provide meaningful improvements in operational flexibility and energy efficiency. Future work should focus on experimental validation, embedded hardware deployment, and multi-objective RL extending to cost and emissions objectives.

References

[1] S. A. Kalogirou, "Artificial intelligence for the modeling and control of combustion processes: A review," *Progress in Energy and Combustion Science*, vol. 29, no. 6, pp. 515–566, 2003.
[2] S. Fettaka, J. Thibault, and Y. Gupta, "Design of shell-and-tube heat exchangers using multi-objective optimization," *International Journal of Heat and Mass Transfer*, vol. 60, pp. 343–354, 2013.
[3] M. Mohanraj, S. Jayaraj, and C. Muraleedharan, "Applications of artificial neural networks for refrigeration, air-conditioning and heat pump systems—A review," *Renewable and Sustainable Energy Reviews*, vol. 16, no. 2, pp. 1340–1358, 2012.
[4] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press, 2018.
[5] R. Sharma and A. Gupta, "AI-assisted control of biomass gasification systems for real-time optimization," *Energy Conversion and Management*, vol. 298, p. 117812, 2024.
[6] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine, *Fundamentals of Heat and Mass Transfer*, 6th ed. New York: Wiley, 2007.
[7] Bejan and M. J. Moran, Thermal Design and Optimization. New York: Wiley-Interscience, 2006.
[8] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015.
[9] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
[10] J. H. Lienhard IV and J. H. Lienhard V, *A Heat Transfer Textbook*, 5th ed. Cambridge, MA: Phlogiston Press, 2020.
[11] Y. Guo, Z. Wang, and H. Chen, "Deep reinforcement learning for adaptive control of heat recovery steam generators in combined-cycle power plants," *Applied Energy*, vol. 330, p. 120356, 2023.
[12] J. Zhang, L. Liu, and M. Wang, "Machine learning approaches for heat exchanger fouling detection and performance prediction: A review," *Energy and Buildings*, vol. 261, p. 111940, 2022.
[13] J.-C. Bourdouxhe and M. Grodent, "Digital twin-based optimization of HRSG net-works using physics-informed neural networks," *Energy Conversion and Management*, vol. 270, p. 116245, 2022.